

UC Berkeley

Controls and Information Technology

Title

Viability of dynamic cooling control in a data center environment

Permalink

<https://escholarship.org/uc/item/0wj7r61r>

Authors

Boucher, T.
Auslander, D.
Bash, C.
[et al.](#)

Publication Date

2006

Peer reviewed

Viability of Dynamic Cooling Control in a Data Center Environment

^{1,5}Timothy D. Boucher, ³David M. Auslander, ⁴Cullen E. Bash, ^{2,5}Clifford C. Federspiel, ⁴Chandrakant D. Patel

¹Lockheed Martin Space Systems Company, Sunnyvale, CA

²Federspiel Controls, El Cerrito, CA

³Dept. of Mechanical Engineering

University of California, Berkeley, CA

⁴Hewlett-Packard Laboratories, Palo Alto, CA

⁵Center for the Built Environment (formerly with)

University of California, Berkeley, CA

Phone 408.742.7474, timothy.d.boucher@lmco.com

Phone 510.642.4930, Fax 510.643.5599, dma@me.berkeley.edu

Phone 650.236.2748, Fax 650.857.7029, cullen.bash@hp.com

Phone 510.418.3392, Fax 510.236.1512, cf@federspielcontrols.com

Phone 650.857.7140, Fax 650.857.7029, chandrakant.patel@hp.com

ABSTRACT

Data center thermal management challenges have been steadily increasing over the past few years due to rack level power density increases resulting from system level compaction. These challenges have been compounded by antiquated environmental control strategies designed for low power density installations and for the worst-case heat dissipation rates in the computer systems. Current data center environmental control strategies are not energy efficient when applied to the highly dynamic, high power density data centers of the future. Current techniques control the computer room air conditioning units (CRACs) based on the return air temperature of the air – typically set near 20 C. Blowers within the CRACs are normally operated at maximum flow rate throughout the operation of the data center unless they are equipped with non-standard variable frequency drives. At this setting the blowers typically provide significantly more airflow than is required by the equipment racks to prevent recirculation and the subsequent formation of hot spots. This strategy tends to be overly conservative and inefficient. As an example air entering a given system housed in a rack undergoes a temperature rise of 15 C due to the heat added by the system. The return air control strategy strives to keep the entire room at a fixed temperature. Therefore in a typical data center the CRAC supply temperature, and hence the air entering the racks, is 13-15 C and the CRAC return is 20-22 C. At these settings the CRACs can consume almost as much energy as the computer equipment they are cooling [1] [2]. Experiments conducted by the authors using these CRAC settings show that nearly 0.7 W is consumed by the environmental control system for every 1 W of heat dissipated by the computer equipment in the authors' experimental facility indicating that the energy efficiency of standard data center environmental control systems is poor. This study examines several opportunities for improving thermal management and energy performance of data centers with automatic control. Experimental results are presented that demonstrate how simple, modular control strategies can be implemented. Furthermore, experimental data is presented that show it is possible to improve the energy performance of a data center by up to 70% over current standards while maintaining proper thermal management conditions.

INTRODUCTION

Data center thermal management challenges have been steadily increasing over the past few years due to rack level power density increases resulting from system level compaction [2][3]. These challenges have been compounded by antiquated environmental control strategies designed for low power density installations. The current state of the art in data center thermal management consists of a single sensory feedback signal located at the return of the computer room air conditioner (CRAC). This sensor provides a global indication of the heat being dissipated in the room and controls the temperature of the computer room air conditioner (CRAC) supply air according to the return air temperature setpoint. Typically the CRAC fan speed is fixed throughout operation. This mode of operation allows no local flexibility in how the cooling is delivered to the computers and there is no local state feedback information from different areas of the data center [3]. This style of operation is inefficient and not readily adaptable to changes in the environment. Without this flexibility, there is no potential to optimize the operation of the data center. Previous work in this area includes using CFD (computational fluid dynamic) models to create intricate mathematical models of the data center thermo-fluid dynamics to find energy-optimal layouts of the data center [4][5][6][7][8][9]. However, these methods are typically based on static data center behavior. If a data center had the sophistication to locally sense and actuate dynamically, then there would exist a great potential for increasing the first and second law thermodynamic efficiency [3]. To do this, the system needs:

- (i) A distributed sensor network to indicate the local conditions of the data center
- (ii) The ability to vary cooling resources locally
- (iii) Knowledge of how each variable affects the conditions of the data center

In theory, with these three requirements in place, a dynamic controller could be developed and implemented to control the conditions of the data center according to certain specifications [10]. Furthermore, this controller could automatically optimize the configuration with regards to minimum energy usage.

This study examines the third requirement listed above. Specifically, we consider how each available actuation variable affects the states of the data center. Several actuators are evaluated in order to determine their usefulness from a control perspective. To satisfy (ii), the actuators of interest in this study are variable CRAC supply temperature, variable speed CRAC fans, and variable opening plenum vent tiles. The variable supply temperature can be achieved either by using a variable capacity compressor (direct expansion units) or by modulating a chilled water supply valve (water cooled units), depending on the type of CRAC

unit. To regulate the CRAC fans speeds, a *variable frequency drive* (VFD) is used to vary motor speeds. For the variable opening plenum ventilation tiles, electronically actuated vent tiles have been developed by the authors and reported previously [10]. These vent tiles are equipped with a linear actuator and can electronically adjust a sliding damper mechanism that varies the effective opening of the vent. In order to satisfy requirement (i), sensory information is required to observe local data center environmental conditions. For this study, a network of temperature sensors is used to collect this information. Thermistors are placed on the inlet and outlet of the computer racks, and at inlets to particular vent tiles, to provide a distributed knowledge of temperature at critical locations in the data center. This provides the necessary information to indicate the local state of the data center.

The states that are of interest are the rack inlet air temperatures (T_i), the “*supply heat index*” metric (SHI), and the total power consumed by the CRAC units. The first is important because equipment inlet air temperatures must be regulated to within boundaries prescribed by the computer equipment manufacturers [11]. In order to keep the computer chips cool, the rack inlet air must be kept below a threshold reference temperature. This constraint is referred to as the “*thermal management*” constraint.

SHI is a non-dimensional metric that indicates the local magnitude of hot and cold air mixing [3]. The definition of the supply heat index is given in equation (1). Figure 1 is a visualization of airflow surrounding a rack of equipment in a data center. The blue streams represent cool air supplied by the vent tiles that gets pulled into the rack. The red streams represent hot exhaust air exhausted from the rack. With reference to Fig. 1, dQ is the amount of hot air expelled into the data center that is pulled back into the intake of a rack prior to exiting the data center (by entering a CRAC for example). Numerous factors can cause this recirculation of hot air and lead to high values of dQ . Inadequate vent tile flow rate and pressure imbalances above the floor are well known causes [6] [12]. This hot air mixes with the air from the plenum vent and results in elevated inlet temperatures which reduces computational compaction and causes CRACs to be operated at lower temperatures and higher flow rates, thus reducing CRAC operational efficiency. In Eq. (1) Q is the heat load from the equipment racks and $(Q + dQ)$ is an artificial load that includes the effect of the recirculation. The resulting ratio is scalable and can be used to quantify recirculation at the system, rack or data center level. Additionally, it can be shown that the ratio of heat loads can be reduced to the ratio of temperature differences in Eq. 1 where T_i , T_o and T_v refer to rack inlet, outlet and vent temperatures respectively [3]. SHI can be measured with either temperature sensors or the heat loads in the ratio can be estimated using various means (power sensors, mass flow sensors combined with temperature sensors, etc.). Further detail on the use of SHI for numerical and experimental analysis of data center performance can be found in the literature [12] [13]. In this study SHI is estimated with temperature sensors only.

$$SHI = \frac{dQ}{Q + dQ} = \frac{(T_i - T_v)}{(T_o - T_v)} \quad (1)$$

In the present study SHI is specific to each height position on the racks. T_i is the inlet temperature and T_o the corresponding outlet temperature at the outlet of the same rack at the same height. T_v is the air temperature from the adjacent plenum vent.

Lastly, the power consumed by the CRAC units is important to consider because, ultimately, this is the state of the system directly impacts operational cost and is therefore desirable to minimize. This power metric is composed of two values, the power required to reduce the temperature of the air and the power required to move the air with the fans.

Experimental results from a raised floor data center are presented here via variation of three actuation variables: CRAC supply air temperature, CRAC airflow rate (S_{fan}), and ventilation tile openings. It is shown that each of these actuators play an important and distinct role in the proper distribution of cooling resources throughout a dynamic data center environment. Furthermore, the potential that these actuators have to control the system is discussed and control methods are proposed. Finally, it should be noted that these results were obtained from a production data center that, although similar in design to typical raised floor installations, may contain attributes that impact the results. Care should be taken when attempting to generalize the findings.

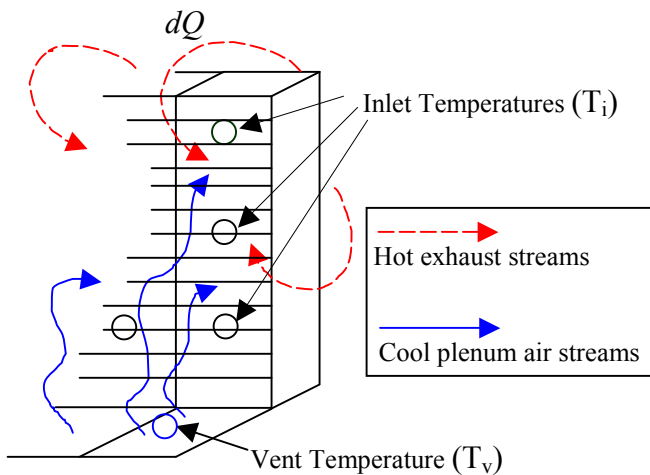


Figure 1 : Visualization of re-circulation flow

EXPERIMENTATION PROCESS

Airflow in a data center is a complex process. Various studies have considered this process numerically as previously discussed. Very few experimental investigations have been performed in working data centers [4] [12] [14]. In the present study, a set of experiments is designed to observe whether or not the trends of the major states of the physical system due to changes in the actuators are consistent and predictable. If so, it can be concluded that it is feasible to control that state of the system using the proposed actuators. These studies are done by performing a series of parametric experiments where only one actuator is varied at a time. The primary goal of these experiments is to quantify and observe the thermal behavior of the data center with respect to all of the actuation variables. This is useful in order to decouple the complicated effects of the data center and observe any subsequent trends. Figure 2 is a schematic of a typical raised floor data center where cool air from the CRACs pressurizes a raised floor plenum and from there is distributed to racked equipment via ventilation tiles placed on the surface of the raised floor. Experiments for this study were conducted at a similar facility located at Hewlett-Packard Laboratories in Palo Alto, CA [10]. The data center is used to support IT equipment for the facility and is therefore a working data center. A small section of the data center can be isolated from the rest of the room via a movable barrier above the floor and manually operable dampers below the floor. In this manner experiments can be conducted in the test section of the room without affecting IT equipment housed elsewhere.




Figure 2 : Typical raised floor data center configuration

Figure 3 displays the floor plan of the experimental section of the data center. All experiments described subsequently were performed in this section of the data center. This area has two CRAC units and two rows of computer racks, row F and G. The CRAC units are Liebert model DE412WUAAEI and are cooled by chilled water. Supply air temperature is controlled via a 3-way valve connected to each unit. The CRAC air movers are each controlled via ABB variable frequency drives (Model ACH400). Table I lists combined CRAC flow rate as a function of VFD speed where VFD speed is listed as a percentage of maximum (S_{fan}). Flow rate was measured using a vane-type anemometer at 18 locations at the intake of each CRAC. Uncertainty is estimated to be $\pm 10\%$ of the reading. The cold aisle consists of a 3 by 10 grid of plenum vents with adjustable

sliding dampers which supply cold air to the inlets of the racks. Vent tile flow rate is approximately $0.35 \text{ m}^3/\text{s}$ (750 CFM) per tile when all tiles are fully open and each CRAC is providing the maximum flow rate. Variation in flow rate among the tiles is 8% with all tiles fully open and 3% with all tiles 50% open resulting from a more uniform plenum pressure distribution. Tile flow rate was measured with a flow hood manufactured by Shortridge Instruments (model CFM-88) with an uncertainty of $\pm 3\%$ of the reading. Additionally, eight racks of variable heat loads, labeled “Resistive Loads” in Fig. 3, are deployed to enable control over room-level power dissipation and distribution. Each resistive load is capable of generating 9.9 kW in 3.3 kW increments and requires $0.57 \text{ m}^3/\text{s}$ (1200 CFM) of airflow at full load. The remaining racks, labeled “Comp. Racks” in Fig. 3, house computer servers that dissipate a maximum of 5 kW and require $0.26 \text{ m}^3/\text{s}$ (550 CFM) of airflow resulting in a temperature rise of 15 C across the racks. Rack power dissipation is measured with RMS power meters manufactured by Power Measurement (models 7330 and 6200). Temperature sensors (Precon thermistors model ST-R3R) are attached to the inlet and outlet of each rack at varying heights. Sensors accuracy is $\pm 0.3 \text{ C}$ according to the manufacturer.

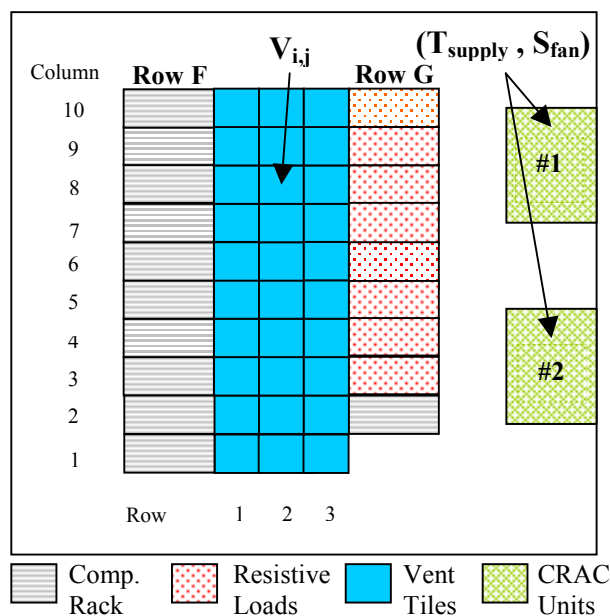


Figure 3 : Floor plan of the experimental data center area

VFD Speed (% max)	Flow Rate (m^3/s)	Flow Rate (CFM)
100	10.73	22738
90	9.50	20125
80	8.26	17511
75	7.34	15551
70	6.41	13591
60	5.30	11238

50	3.82	8102
40	2.84	6011
30	1.97	4182
25	0.86	1830

Table I: VFD Speed vs. Combined CRAC Flow Rate

CRAC Supply Temperature

The impact of CRAC supply temperature variation on rack inlet temperatures is investigated while CRAC fan speed and vent positions are held constant. The goals of this experiment are two-fold. First to show whether or not buoyancy effects in the air are negligible in the dynamics of the data center, and second to see if we are able to exert “control” over the rack inlet temperatures by varying the CRAC supply temperatures.

The density of air is a function of temperature. When hot air exists in the presence of cooler air, convective flows can form as a result of buoyancy forces [15]. These buoyant flows, if induced and large enough, may have an impact on the bulk flow patterns in the room and thereby negatively affect control. There is evidence from the study of office underfloor air distribution systems that the temperature distribution in a room is not dependant upon supply air temperature indicating that buoyancy forces are not affected by changes in supply air temperature [16]. This result is expected to carry over to data centers where buoyancy induced flow is far overshadowed by forced convective flows. If the buoyancy effect is negligible, the flow patterns around the racks should be independent of CRAC supply air temperature for a given CRAC fan speed setting (S_{fan}). In this situation the rack inlet temperatures should vary linearly with the CRAC supply temperature.

The second goal of this experiment is to see if we are able to control the rack inlet temperatures by varying only the CRAC supply temperatures and to determine the magnitude of this control. Again, the characteristics of these results may depend on CRAC fan speed. For example, there may exist a fan speed such that the inlet temperature requirements cannot be met even with the lowest possible CRAC supply temperature. These experiments will give an idea of limitations regarding inlet temperature regulation with CRAC supply air temperature.

CRAC Fan Speeds

This experiment investigates the effect on the system as a result of varying CRAC fan speed (S_{fan}). The hypothesis of the present study is that CRAC fan speed has a significant effect on the re-circulation flows in the data center, which are thought to

be a major cause of inefficiency. If a predictable correlation between CRAC fan speed setting and SHI can be found, this implies that fan speed can be used to control the SHI metric, and can therefore be utilized to minimize operational inefficiencies.

Ventilation Tile Positions

Since ventilation tiles are the actuators closest to the equipment racks, they potentially have the greatest impact on the distribution of local cooling resources. All the vent tiles in the test data center are equipped with manually adjusted dampers that allow for variable flow resistances. It is useful to know the properties and effects of these vent tiles when analyzing a static configuration. But more importantly, understanding the vent tiles can help in formulating a control scheme in which the vent tiles can be controlled using local state feedback.

The benefit of implementing a controllable tile scheme is that it enables a controller to affect local conditions in the cold aisle. This is where the CRAC fan speed actuation is limited. The fan speed can only change the conditions for a large portion of the plenum from which the effects are too widespread to effectively control a local region without impacting neighboring areas. As with CRAC fan speed, it is believed that increasing the flow into the cold aisle should decrease the amount of hot air infiltration from the hot aisle and vice versa.

The primary purpose of this experiment is to change the openings of groups of vent tiles to observe the subsequent effects on the rack inlet temperatures. Specifically, the relationship between variable vent grouping size to the local effect on thermal management is considered. The secondary purpose of this study is to compare and contrast different configurations of vent actuations. Each configuration is being defined here by the number of vent rows that are altered simultaneously.

The economics of automatic vent tile control are important to consider. Controllable vent tiles will necessarily be more costly than uncontrollable vents. It is therefore important to understand how many controllable vents are required to significantly impact the environment. In order to study this, two different configurations of vent tiles are employed with each configuration containing a different number of controllable vents. As shown in Figure 4, *method 1* considers all three rows as one. Therefore, for each vent column there are 3 vents, one for each vent row, which are all coupled together. For example, if vents 1-3, as shown on Figure 4, were closed, a total of 9 vents would be closed (rows 1-3, columns 1-3). In *method 2*, the middle row of vents are controllable and the outside two rows are kept fully open representing static tiles. So, if vents 1-3 were closed in method 2, only vents 1-3 in row 2 would be closed. All experiments performed are done with a constant CRAC supply temperature of 18.3 °C and an S_{fan} of 50% (30Hz) for both CRACs, while the vent tile openings were adjusted.

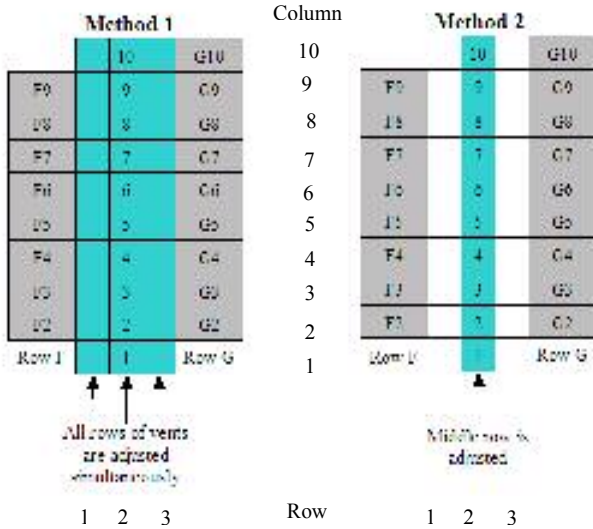


Figure 4 : Mapping of vents used in each method

In this investigation, vent columns 4-7 (from Figure 4) were moved from 100% open to 0% open. Each data point represents an inlet temperature location (top or middle of the rack inlets), which has a corresponding vent position and a ΔT . The corresponding vent position is the column (from Figure 3 and 4) of the vent that is adjacent to the rack inlet in question. The corresponding ΔT is defined as follows:

$$\Delta T \equiv (T_i)_{4to7} - (T_i)_{open} \quad (2)$$

The temperatures are recorded from each sensor when all the vents are completely open ($(T_i)_{open}$). Then after plenum vents 4-7 are completely closed, the inlet temperatures are recorded again ($(T_i)_{4to7}$). The difference between these two temperatures define ΔT .

Power Consumption

Along with thermal management, the power consumption of the environmental control system is important to consider as it directly impacts the data center operational costs. Due to the multi-variant nature of the problem, it is difficult to study the effect on power due to each individual actuation. In spite of this, experiments are conducted that allow for the profiling of the data center in terms of its energy usage. Results of this profiling provide insight into operational efficiency and can be compared with the energy consumption of standard data center environmental control methods.

There are two processes that consume power, the power to cool the air in the CRAC, and the power to move the air through the CRAC and, subsequently, the data center. The former can be estimated by using one of two methods, depending on whether the CRAC units are chilled water cooled or are cooled by direct expansion. The chilled water units in the experimental data center are supplied by a chiller that serves a broader facility and, therefore, is not sensitive to changes in individual CRACs. Consequently, for chilled water systems, power consumption is estimated via the heat load balance measured using the difference in chilled water supply and return temperatures and the chilled water flow rate for each CRAC. An estimated COP for the central cooling cycle can then be used to approximate changes in compressor power for the central chiller as a function of CRAC supply air temperature. However, if the CRAC units are direct expansion, the power consumed by the compressor can be directly measured using a power meter. Fan power is obtained by measuring the power drawn by the CRAC fans using a power meter incorporated into each VFD. Therefore, for any given values of CRAC and vent settings, the corresponding steady state operating power demands can be estimated.

From a facility management perspective, there are two goals for the operation of a data center:

- Supply computer systems with sufficiently cool air (thermal management constraint)
- Minimize power consumption

To define exactly what the thermal management constraint is, a desired reference temperature, T_{ref} , is defined to be the upper limit for the rack inlet temperatures. For a given system, T_{ref} is dictated by the thermal specifications of the computer systems. Therefore, as long as the hottest rack inlet temperature is less than or equal to T_{ref} , the thermal management constraint is satisfied.

RESULTS AND DISCUSSION

CRAC Supply Temperature

Rack inlet temperature as a function of CRAC supply temperature for CRAC fan speeds of 33% and 66% are shown in Figures 5a and 5b respectively. Each of the curves represents the measurements for a specific temperature sensor placed at the inlet near the top of a rack in Row G of Fig. 3. CRAC supply temperatures were measured in the under floor plenum near the exhaust of each CRAC and then averaged. At lower flow rates the CRAC units have difficulty maintaining setpoints resulting in slightly different supply temperatures for different CRAC fan speeds. The tests were conducted on different days under different heat

load conditions and should not be compared against each other. All of the sensors exhibit a monotonically increasing response to increases in CRAC supply temperature, and the degree of this increase is a function of sensor (or rack) location and CRAC flow rate. The magnitude of the response is generally less than the corresponding change in supply temperature. Although there is some support for this result in the literature on underfloor air distribution plenums in office environments, additional work is required to fully explain this behavior [16]. Additionally, the racks located near the end of the row (G9 and G10) generally exhibit a higher temperature than those toward the interior due to recirculation of hot exhaust air around the end of the row [3].

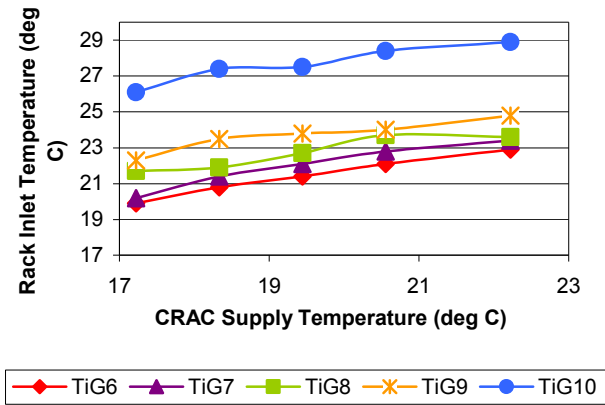


Figure 5a : Rack inlet temperatures for 33% VFD

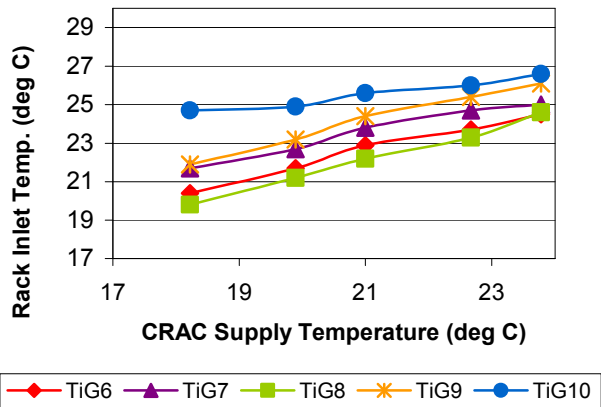


Figure 5b : Rack inlet temperatures for 66% VFD

This monotonic dependence of rack inlet temperature on CRAC supply temperature, although expected, is positive in terms of control. A clear and predictable trend exists between this actuation and the rack inlet temperatures. Furthermore, it is shown that CRAC flow rate does not impact the general monotonic shape of this dependence. This means that CRAC supply temperature is well suited for controlling rack inlet temperature regardless of flow rate. As discussed previously, the control that we need over inlet temperatures is to maintain thermal management. Therefore, a feedback control loop can be employed between the inlet

temperatures and the CRAC supply temperatures, with the error defined as the difference between the maximum inlet temperature $(T_i)_{\max}$ and T_{ref} .

However, there may exist ranges of CRAC fan speeds and/or vent positions that make it impossible for the CRAC supply temperature alone to be able to satisfy thermal management. This possibility must be considered when controlling the CRAC fan speeds and vent positions in combination with supply air temperature. The controller must be able to detect this infeasibility and alleviate the problem by changing the settings of these other two actuators accordingly.

CRAC Fan Speed

Figure 6 is a chart of the Supply Heat Index vs. CRAC blower speed at various rack inlet sensor positions corresponding to Fig. 3. S_{fan} values are uniformly varied for the two CRACs from 30% to 95%. Variance in steady state SHI with S_{fan} for eight different rack inlet sensor locations in row G is indicated in the figure. For this experiment the CRAC supply temperatures are 18.3°C (65°F) and all vent tiles are 100% open. The resistive heat loads are set to 9.9 kW (racks G9 and G10) and 6.6 kW (racks G3 – G8). Uncertainty of the SHI calculation is less than ± 0.03 from a root sum of the squares analysis. It is expected that SHI is inversely proportional to CRAC fan speed indicating that recirculation, or dQ , will increase as fan speed decreases. The results from Figure 6 indicate that this trend is in fact the case for a subset of the test range. Specifically, at less than 60% VFD speed this inverse trend holds for most racks. At 60% VFD speed the CRACs are circulating approximately 5.2 m³/s (11,000 CFM) of air through the data center. From a caloric analysis of the racks the combined air flow rate through the racks is 5.5 m³/s (11,600 CFM). As the total CRAC flow rate (and, thus, vent flow rates) fall below that being pulled through the racks, the total flow rate is considered under-provisioned and additional air is pulled by the racks from alternate areas of the room causing recirculation to necessarily increase. As total CRAC flow rate is increased beyond the level of equilibrium with the racks there is an over-provisioning of air-flow and further increases do not significantly impact SHI and, indeed, waste energy. This is evident by a flattening of the SHI curve for most sensor locations beyond the 60% mark in Fig. 6. It can be further observed in Figure 6 that the local SHI for some points (G3 and to a lesser degree G4) remain largely unaffected by flow rate. A flow hood manufactured by Shortdridge Instruments (model CFM-88) was used to verify that vent tile flow rates are independent of location in this data center (when all vents are 100 % open), therefore this trend is not a function of flow conditions under the floor but can be explained rather by flow and heat dissipation conditions above the floor. Specifically, Racks G3 and G4 are located near the end of Row G in an area of the data center that dissipates less heat than the rest of the room. As such, there is less hot air for those racks to scavenge. Additionally, they require about 33% less air flow than racks G9 and G10 located at the opposite end of the row - about 0.42 m³/s (900 CFM) compared to 0.64 m³/s (1350 CFM) further

contributing to reduced values of SHI at lower CRAC flow rates. Further, note that in conditions of extreme underprovisioning ($S_{fan} < 40\%$), with the exception of G10, SHI increases sequentially within the margin of experimental uncertainty from rack locations G3 to G9 indicating that hot air is clustered around the upper portion of the room. Other flow phenomena above the floor likely contribute to this result as well indicating that sensor/actuator characterization is important to the development of any data center control system.

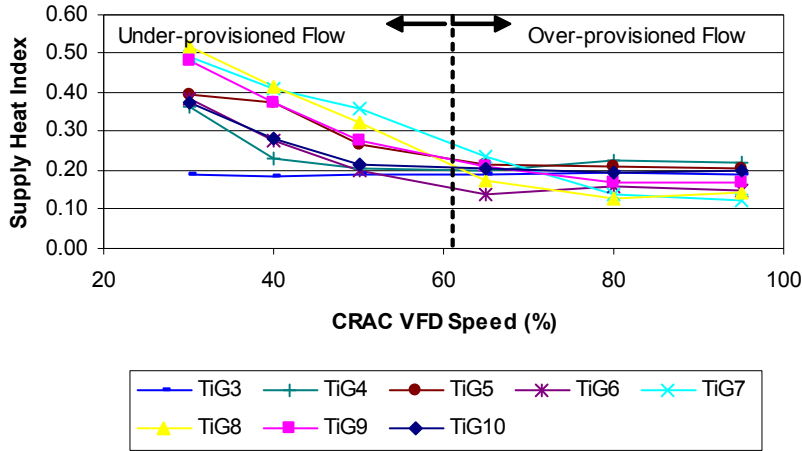


Figure 6 : Local SHI vs. CRAC VFD

From a control perspective, one of the preliminary ideas was to control SHI (i.e. re-circulation) using this actuation. However, based on the results here, a control strategy of this kind may not be ideal. The non-linearity in the trend would make it difficult for a simple linear regulation controller to perform well. This type of control would not be guaranteed to actuate the system in the correct fashion unless the range of control were limited to the linear response range (i.e. under-provisioning). Non-linear control techniques, however, might prove satisfactory. Using neural networks or fuzzy logic control, for example, a controller could be implemented that could use information that it collects in real time. Using this information along with a way to ascertain whether it is producing positive or negative results, the system can “learn” how to control the non-linear behavior. The only pre-requisite would be that the behavior be non-chaotic and repeatable. The downside of this is that any changes made to the system would require a re-education of the controller. Another option would be to use CFD models to simulate the details of the CRAC fan speed influence. However, this method is time-consuming and may be insensitive to changes in system configuration unless the models are re-run to accommodate these changes. Other, non-model based approaches that do not require learning could be considered as well. For example, the CRAC fan speed could be optimized via minimization of the total CRAC power consumption. A searching algorithm could be implemented that would search through various S_{fan} values and

use the corresponding power values in order to intelligently make new guesses for S_{fan} . Further explanation of this concept is forthcoming.

Plenum Vents

Results of the vent tile experiments are shown in Figures 7 and 8. Figure 7 plots rack inlet temperature data for the sensors at the top of the row G racks according to test method 1 with vents 4-7 closed. The data show that the vents can produce predictable trends in the inlet temperature values. When vents are closed, it causes not only the inlet temperatures adjacent to the vents to increase, but it causes the inlet temperatures of nearby locations in the aisle to decrease. This result can be explained as follows. When a plenum vent is closed, it restricts all airflow through that vent. This lack of air supply causes the rack inlets to be susceptible to hot air infiltration because it causes a lower pressure region to form at the rack inlet. This in turn causes rack temperatures to rise. Additionally, for a given CRAC fan speed, there is a zero net flow into the plenum. Therefore, the rate at which air is entering from the CRACs is equal to the sum of the air flow exiting the vents. So, disregarding the effect on the fan operating point and CRAC flow rate due to vent tile manipulation, when some of the vents are closed more air is forced out of the remaining vents that are open in order to maintain this net zero flow. This direct increase in flow supplies more cooling air to nearby racks and, in general, lowers their temperature (though the degree to which a given rack is affected depends upon the fluid movement and pressure distribution above the floor).

Although CRAC fan speed and vent tile restriction both are used to vary flow rate through ventilation tiles, their relative effect on the data center fluid dynamics is very different. Changing S_{fan} causes changes to the global flow patterns in the plenum, which are more or less transmitted to the equipment racks located in the region of influence of the CRAC, though this transmission is highly location dependant as evident from Fig. 6. On the other hand, the plenum vent tile is a local actuation. It doesn't have as much effect on global flow patterns in the data center and is much more influential on the flow provided to individual racks in close proximity to the manipulated vent.

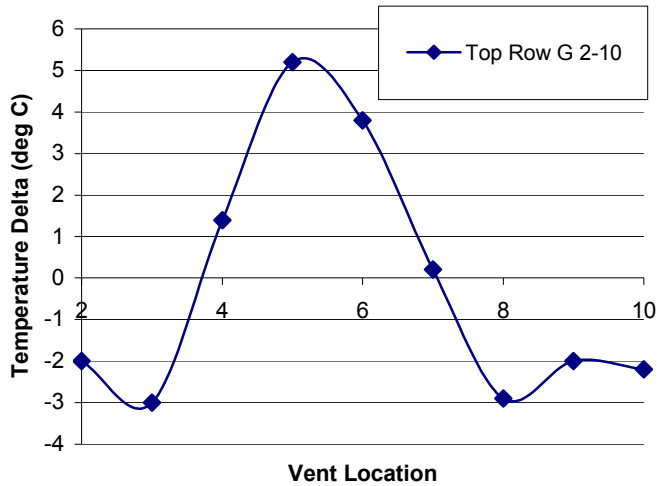


Figure 7 : Method 1 rack inlet delta temperatures

Figure 8 compares the two experimental methods of vent tile actuation. Data from temperature sensors located at varying elevations on each rack are displayed. A 4th order polynomial curve is fitted to all of the data points for each method using all of the available inlet temperature sensors (top, middle, and bottom height level of row G). This gives an average ΔT distribution with respect to vent location for each method. Results indicate that for the method that involves only one row of vents (*method 2*), the corresponding influence on the inlet temperatures is attenuated. For the point in the middle of the closed vent group, the temperature delta for *method 2* is about 40% that of *method 1*. It is also interesting to see that the shape of the curve is similar indicating that both methods produce the same qualitative effect, but the strength of this effect is a function of the number of controllable vents implemented. The size of the groups of vents is also an issue with the consistency of these results diminishing, and trends in the data becoming less clear, as the group size gets smaller.

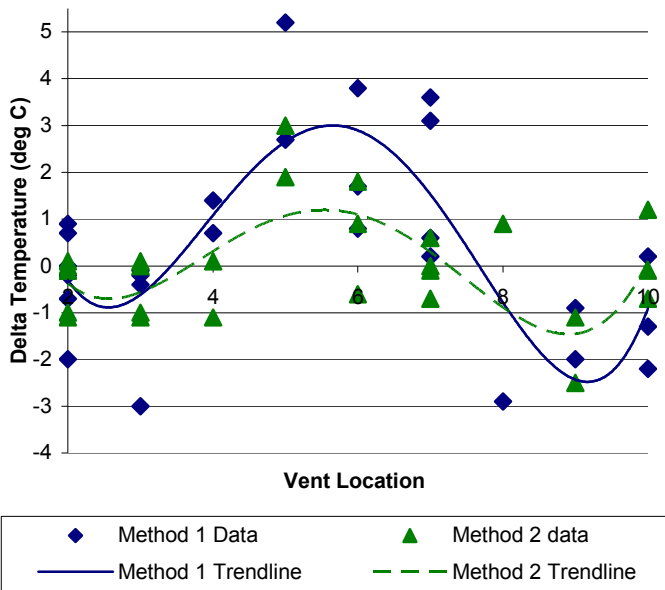


Figure 8 : Comparison of Methods 1 and 2

The results indicate that vent tiles can exhibit significant control locally ($\sim 3^{\circ}\text{C}$) and that they could be useful in a control framework. Also, results show that control of even a limited number of tiles may still be beneficial for the data center operation in terms of economic efficiency since vent tile adjustment is less expensive than either increasing CRAC flow rate or decreasing the supply air temperature. It is, therefore, feasible that a local control scheme could be applied to vent tiles in order to control the local rack inlet conditions. However, vent tile actuation by itself may not provide effective control given the non-monotonic response with rack location. Rather, it may be better to consider vent tile control in the context of a larger control strategy. For example, vent tiles could be controlled in coordination with CRAC supply temperature to help efficiently satisfy thermal management. CRAC supply temperature control can regulate the hottest inlet temperature to T_{ref} , but cannot affect the relative distribution of the rack inlet temperatures locally. This means that all the other rack inlet temperatures, aside from the hottest one, are over-cooled in any particular CRAC region of influence. The vent tile control could tune the flow distribution within the cold aisles to provide more uniform rack inlet temperature and enable a subsequent increase in CRAC supply temperature. This would yield a more energy efficient operating point. However, as is evident from the scatter in the data of Fig. 8, vent tile influence on nearby racks can be difficult to predict and must be carefully considered when formulating an appropriate control strategy.

Power Consumption

Given any arbitrary, but satisfactory, CRAC fan speed and vent tile distribution, there will be a unique corresponding CRAC supply temperature such that the hottest steady state rack inlet temperature will be T_{ref} . However, in general, a set of CRAC fan speeds and vent tile distributions exist such that there is no CRAC supply temperature that will satisfy thermal management subject to the operational constraints of the CRAC. This set of points define the set of unsatisfactory operating points. Using this structure, and allowing the CRAC supply temperature to vary according to the thermal management constraint, variation of power consumption due only to variation in S_{fan} can be investigated. This simplifies the analysis because it decomposes the multi-variant system into a single variable – CRAC fan speed. Figure 9 shows two examples of how CRAC power consumption changes with S_{fan} . At each data point on the plot, S_{fan} is held constant and the CRAC supply temperature is modulated until thermal management is satisfied with $T_{ref} = 25$ C. There are two curves, each for a different configuration. Configuration 1 is defined by all the computer racks (Racks F2-F9 and G2) and five of the resistive loads (Racks G3 and G5 at 6.6 kW each, and G6, G8, and G10 at 9.9 kW each) being active. Total power dissipation for configuration 1 is 56 kW with the active computer racks lightly loaded. Configuration 2 consists of all computer racks and all of the resistive loads (Racks G3, G4, G5 and G7 at 6.6 kW each, and G6, G8, and G10 at 9.9 kW each) active. Total power dissipation is 86 kW. Uncertainty in the power estimation is ± 2 kW for each configuration. Figure 9 shows a maximum CRAC power consumption equal to 60% of the total rack heat load at 100% CRAC fan speed for Configuration 1. The minimum CRAC power consumption occurs at 60% CRAC fan speed with CRAC power consumption equal to 24% of the total heat load. Configuration 2 exhibits a maximum of 52% total load at 100% CRAC fan speed and a minimum of 38% at 80% CRAC fan speed. The reduction in efficiency for configuration 2 is likely attributed to increased recirculation (dQ) due to higher rack air flow rates. The single data point on the graph marked “Config 1 – 21.5C Return” shows the CRAC power consumption for configuration 1 when the data center is operated according to standard practice with the return air temperature of each CRAC kept at 21.5 C and the CRAC fan speeds fixed at 100%. CRAC power consumption in this operational mode is 68% of the total heat load. Comparing the minimum of configuration 1 with the standard CRAC return air temperature control method, it is possible to reduce CRAC power consumption by 65% and still maintain thermal management in this data center.

Notice in Figure 9 that the configuration with the higher compute load leads to a higher power consumption of the cooling system for all fan speeds. In a dynamic data center environment, this curve will also be dynamic. As the power loads change on the computational equipment the power curves will change shape and shift up and down. Note that the location of the minimum for each curve occurs at different VFD speeds. This is a result of the fact that the dynamics of the data center change as the power dissipation and rack air flow rates are modified. This dependence of the CRAC power consumption profile on CRAC

flow rate subject to the constraints of thermal management adds complexity from a controls perspective because a method is required that can adapt to a dynamic optimal point. It is also important to point out that there are S_{fan} values that don't have a corresponding power consumption value because the thermal management constraint was unattainable. These are unsatisfactory S_{fan} values as defined previously.




Figure 9 : Steady-state power consumption vs. VFD

Figure 9 shows two examples of a “power profile” that can be obtained for a data center. This result shows that for each data center there will be a curve that will characterize the power usage of the environmental control system as a function of S_{fan} , while constrained by thermal management. There will be a minimum power value of this curve, whether it is an interior minimum or boundary minimum. Corresponding to this minimum will be the optimal S_{fan} value. This is the optimal CRAC fan speed required to minimize energy consumption. Using optimization theory, an algorithm can be developed that can “search” for this optimal S_{fan} intelligently. The challenge is to develop a self-optimizing system that is robust, that does not require an accurate or complex model of the data center, and that can adapt to changing conditions and different data centers.

For complex, multi-CRAC data centers, the same strategy can be applied, only on a CRAC by CRAC basis. Each CRAC unit uses the power it consumes as the cost function for the optimization algorithm. Therefore, several of these optimizations can be performed simultaneously to form a de-centralized optimization strategy. The repeatability and uniqueness of the resulting optimal solution is an issue of further study, but the method itself is scalable to large complex data centers. Ongoing research in this area will be discussed by the authors in subsequent reports.

CONCLUSIONS AND NEXT STEPS

Over the course of this study three types of actuators were evaluated and tested in a production data center; namely, CRAC supply temperature, CRAC fan speed, and plenum vent tile openings. The results of this evaluation indicate the following for the data center under test:

- CRAC supply temperatures have an approximately linear relationship with rack inlet temperatures, but the slopes differ slightly with respect to rack location and CRAC flow rate.
- Fan speed has a predictable effect on the Supply Heat Index when the flow from the CRACs is underprovisioned. Overprovisioned CRAC flow has little or no impact on SHI.
- Vent tile opening has a direct effect on local rack inlet temperatures in close proximity to the vent and an inverse relationship on racks farther away. The magnitude of the effect is a function of the number of vent tiles that are manipulated.
- There is a concave up relationship between fan speed and total CRAC power consumption when thermal management is being enforced. Total CRAC power consumption was shown to be reduced by 65% while maintaining thermal management when compared with standard data center thermal management techniques.

The results of this paper were produced by operating multiple CRAC units in unison. Additional work is required to understand the impact of independent CRAC control on these results and whether or not the results translate to alternative data center layouts and designs. Nevertheless, the results from the current study demonstrate the feasibility of control via each of the three primary actuators. Using the results, methods are proposed for future control methodologies that have the potential to significantly increase the operational efficiency of data center environmental control systems. Overall, significant control potential exists using these actuators allowing a dynamic control scheme to be a feasible method to optimally distribute cooling resources.

ACKNOWLEDGMENTS

Thanks very much to Ratnesh Sharma of Hewlett-Packard Laboratories who helped out greatly during the experiments by making sure the computing infrastructure was working properly as well as for offering technical advice and assistance in other areas. Also, thanks to Amip Shah and Keke Chen of UC Berkeley for various advice and helping to gather data.

NOMENCLATURE

<i>T</i>	Temperature (°C or °F)
<i>S</i>	Speed (% of maximum)
<i>SHI</i>	Supply heat index
CRAC	Computer Room Air Conditioner
COP	Coefficient of Performance (for refrigeration cycle)

Subscripts

<i>Ref</i>	Reference
<i>i</i>	Rack Inlets
<i>o</i>	Rack Outlets
<i>v</i>	Plenum vents
<i>supply</i>	CRAC supply
<i>fan</i>	CRAC fans
<i>open</i>	Case when vents are all completely open
<i>x to y</i>	Case when vent rows x-y are shut completely
<i>max</i>	Maximum value

REFERENCES

- [1] Friedrich, R., Patel, C.D., 2002, "Towards Planetary Scale Computing – Technical Challenges for Next Generation Internet Computing," *THERMES 2002*, Santa Fe, NM.
- [2] The Uptime Institute, "Heat density trends in Data Processing, Computer systems and Telecommunications Equipment", White Paper issued by *The Uptime Institute*, 2000

- [2] C.D. Patel, "A vision of energy aware computing – from chips to data centers", *The International Symposium on Micro-Mechanical Engineering*, ISMME2003-K15, December, 2003
- [3] Bash, C.E., Patel, C.D., Sharma, R.K., "Efficient Thermal Management of Data Centers – Immediate and Long-Term Research Needs", *Intl. J. HVAC&R Res.*, Vol.9, No.2, pp., Apr 2003
- [4] C.D. Patel, C.E. Bash, C. Belady, L. Stahl, D. Sullivan, "Computational fluid dynamics modeling of high compute density data centers to assure system inlet air specifications", *Proceedings of IPACK'01 – The PacificRim/ASME International Electronics Packaging Technical Conference and Exhibition*, Kauai, HI, July 2001
- [5] C.D. Patel, R.K. Sharma, C.E. Bash, A. Beitelmal, "Thermal Considerations in Cooling Large Scale High Compute Density Data Centers", *ITherm 2002 – 8th Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems*, San Diego, CA, May 2002
- [6] Schmidt, R., Cruz, E. "Cluster of High Powered Racks within a Raised Floor Computer Data Center: Effect of Perforated Tile Flow Distribution on Rack Inlet Air Temperatures", *Proceedings of IMECE*, November 15-21, 2003.
- [7] Kang, S., Schmidt, R., Kelkar, K., Patankar, S., "A Methodology for the Design of Perforated Tiles in Raised Floor Data Centers Using Computational Flow Analysis", *IEEE-CPMT Journal*, Vol. 24, No. 2, pp.177-183, June 2001.
- [8] Rambo, J. D., Joshi, Y. K., 2003, "Multi-scale Modeling of High Power Density Data Centers," *Proceedings of the International Electronic Packaging Technical Conference and Exhibition (IPACK2003-35297)*, Maui, HI.
- [9] Shah, A. J., Carey, V. P., Bash, C. E., Patel, C. D., 2004, "An Exergy-Based Control Strategy for Computer Room Air-Conditioning Units in Data Center," *Proceedings of the 2004 International Mechanical Engineering Congress and Exposition (IMECE2004-61384)*, Anaheim, CA.
- [10] Patel, C.D., Bash, C.E., Sharma, R.K, Beitelmal, A and Friedrich, R.J., "Smart Cooling of Datacenters", *Proceedings of IPACK'03 – The PacificRim/ASME International Electronics Packaging Technical Conference and Exhibition*, Kauai, HI, IPACK-35059, July 2003
- [11] *Thermal Guidelines for Data Processing Environments*, ASHRAE, 2004.
- [12] R.K. Sharma, C.E. Bash, C.D. Patel, A. Beitelmal, "Experimental investigation of design and performance of data centers", *ITherm 2003 – 9th Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems*, Las Vegas, NV, June 2004
- [13] Blackmore, B., "Sequential Optimization of Data Centers with CFD", *Proceedings of IMAPS Advanced Technology Workshop on Thermal Management for High-Performance Computing and Wireless Applications*, October 24, 2003.

- [14] Schmidt, R., Karki, K., Kelkar, K., Radmehr, A., Patankar, S., "Measurements and Predictions of the Flow Distribution through Perforated Tiles in Raised Floor Data Centers", Proceedings of IPACK '01, The Pacific Rim/ASME International Electronic Packaging Technical Conference and Exhibition, Kauai, Hawaii, July 8-13, 2001.
- [15] Gebhart, B., Sammakia, B., Jularia, Y., Mahajan, R.L., "Buoyancy Induced Flows and Transport", *Taylor and Francis*, 1988
- [16] Webster, T., F. Bauman, J. Reese, and M. Shi. 2002. "Thermal stratification performance of underfloor air distribution (UFAD) systems." *Proceedings of Indoor Air 2002*, Monterey, CA, 30 June - 5 July 2002.

FIGURE AND TABLE CAPTIONS

Figure 1 : Visualization of re-circulation flow

Figure 2 : Typical raised floor data center configuration

Figure 3 : Floor plan of the experimental data center area

Figure 4 : Mapping of vents used in each method

Figure 5a : Rack inlet temperatures for 33% VFD

Figure 5b : Rack inlet temperatures for 66% VFD

Figure 6 : Local SHI vs. CRAC VFD

Figure 7 : Method 1 rack inlet delta temperatures

Figure 8 : Comparison of Methods 1 and 2

Figure 9 : Steady-state power consumption vs. VFD

Table I: VFD Speed vs. Combined CRAC Flow Rate