



Local Temperature Control in Data Center Cooling: Part I, Correlation Matrix

Keke Chen, David M. Auslander, Cullen E. Bash, Chandrakant D. Patel
Enterprise Software and Systems Laboratory
HP Laboratories Palo Alto
HPL-2006-42
March 6, 2006*

data center cooling,
smart cooling,
local temperature
control, vent tile,
correlation matrix

A data center is a computer room containing a dense aggregation of commodity computing, networking and storage hardware mounted in industry standard racks. With the evolution of microprocessor fabrication technology and the increasing demand of internet, power density has been growing from the chip level to the data center level. As a result, mechanical designers face the challenge of handling heat dissipation efficiently in a data center. This paper demonstrates the possibility of optimizing local temperature distribution using vent tiles. In the proposed experiments, the relationship between rack inlet temperature and vent tile configuration is analyzed, and used to demonstrate the effectiveness of a local control algorithm through a control simulation.

Local Temperature Control in Data Center Cooling:

Part I, Correlation Matrix

¹Keke Chen, ²David M. Auslander, ³Cullen E. Bash, ³Chandrakant D. Patel

¹Center for the Built Environment

University of California, Berkeley, CA

²Dept. of Mechanical Engineering

University of California, Berkeley, CA

³Hewlett-Packard Laboratories, Palo Alto, CA

Phone 510.642.2720, keke@berkeley.edu

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

ABSTRACT

A data center is a computer room containing a dense aggregation of commodity computing, networking and storage hardware mounted in industry standard racks. With the evolution of microprocessor fabrication technology and the increasing demand of internet, power density has been growing from the chip level to the data center level. As a result, mechanical designers face the challenge of handling heat dissipation efficiently in a data center. This paper demonstrates the possibility of optimizing local temperature distribution using vent tiles. In the proposed experiments, the relationship between rack inlet temperature and vent tile configuration is analyzed, and used to demonstrate the effectiveness of a local control algorithm through a control simulation.

Key Words: data center cooling, smart cooling, local temperature control, vent tile, correlation matrix

INTRODUCTION

A data center is a machine room containing computing, networking and storage hardware that provide useful services. Over the past decade, the power dissipated per unit area of a computer chip has increased by a factor of ten; heat dissipation of the microprocessor has also gone up by an order of magnitude. On the other hand, increasing demand by consumers for commodity computing pushes computer manufacturers to produce thinner systems. Nowadays a rack can accommodate 40 thin systems, with more possible when utilizing bladed architectures. If each computer system dissipates 300 W, a single rack will dissipate 12 KW. A 100,000 ft² data center can house 5000 of these racks. Then the total heat dissipation in such a data center will be 60MW, and an additional 30MW of power will be needed to cool the system. At \$100/MWh, the total cost of the data center would be around \$52.6 million per year to power the racks and an additional \$26.3 million per year in cooling cost. Clearly, efficient cooling for data centers is of great interest for both environmental energy conservation and cost savings. It also becomes a major challenge for mechanical engineers.

Figure 1 depicts a typical data center with an under-floor air distribution system. Computer racks are arranged in rows, separated by supply-cool-air zones (cold aisles) and return-hot-air zones (hot aisles). Computer Room Air Conditioning (CRAC) units deliver chilled air into under-floor plenums, push it upward through perforated tiles, feed it through the computer racks, and thus remove the heat that the computers generate.

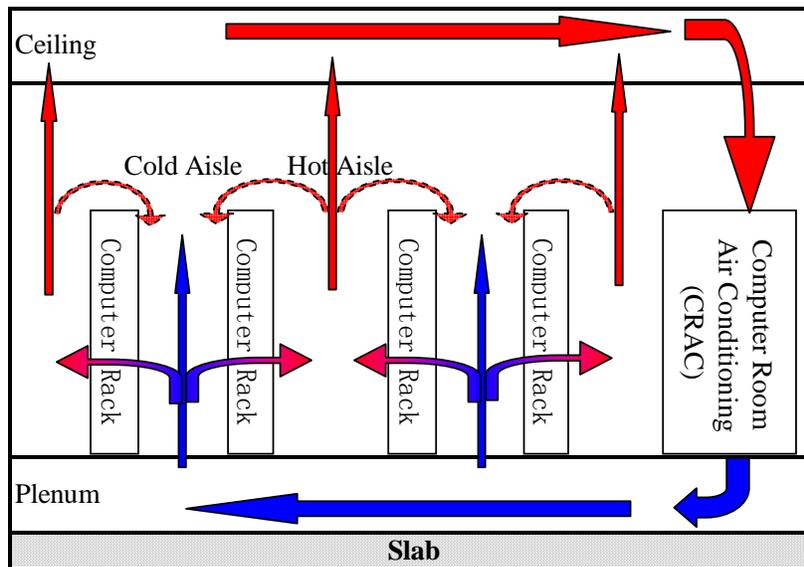


Figure 1: Schematic Section of a Data Center

To increase the energy efficiency of a data center cooling system, CRAC units have to deliver proper cooling resources for each computer system according to its heat load. However, the situation at an individual computer location is usually not ideal. Cooling-resource distribution depends on system layouts and boundary conditions. It is more likely that some computers will get more cooling resources than required, while others get less. This phenomenon is reflected by an uneven rack inlet temperature distribution. The inlet temperature of a computer system must be lower than the manufacturer's specification for the computer. If the inlet temperature is higher than the specification, the computer doesn't get enough cooling resource resulting in a danger of overheating. To satisfy all of the computer equipment in a data center, CRAC units have to ensure that the maximum inlet temperature of computer systems is within their manufacturers' specifications.

Uneven airflow distribution across the arrays of perforated tiles can result in air recirculation from hot aisles to cold aisles. Thus rack inlet temperature distribution is not uniform. If the maximum inlet temperature increases, CRAC units have to reduce their temperature set-points or increase their fan speed, which increases the energy consumption of the cooling systems. Therefore, we want to make airflow distribution across the perforated tiles and the rack inlet temperature distribution uniform. Once the maximum inlet temperature is lower, the energy consumption will also be lower. This is a significant saving in cost.

Either decreasing the maximum inlet temperature of the computers or increasing their minimum inlet temperature can make temperature distribution more uniform, but decreasing the maximum temperature is of more direct interest than increasing the minimum temperature. Remotely actuated vent tiles, or Smart Tiles, can be constructed to vary local airflow rates and manipulate the temperature distribution across a row of racks. However, the utility of such a device is not well understood. Their impact on the distribution of local cooling resources must be clarified before using “smart tiles” in data centers.

EXPERIMENTATION

1) Research Area

Figure 2 shows the plan view of the HP Data Center. The Grizzly area, Central area and Research area in the plot represent different sections. The research area can be separated from the other two areas by a wall between the Research area and Grizzly area. There are curtains between the Research area and Central area and corresponding dampers under the plenum, but there is nothing to separate the Grizzly area and Central area physically.

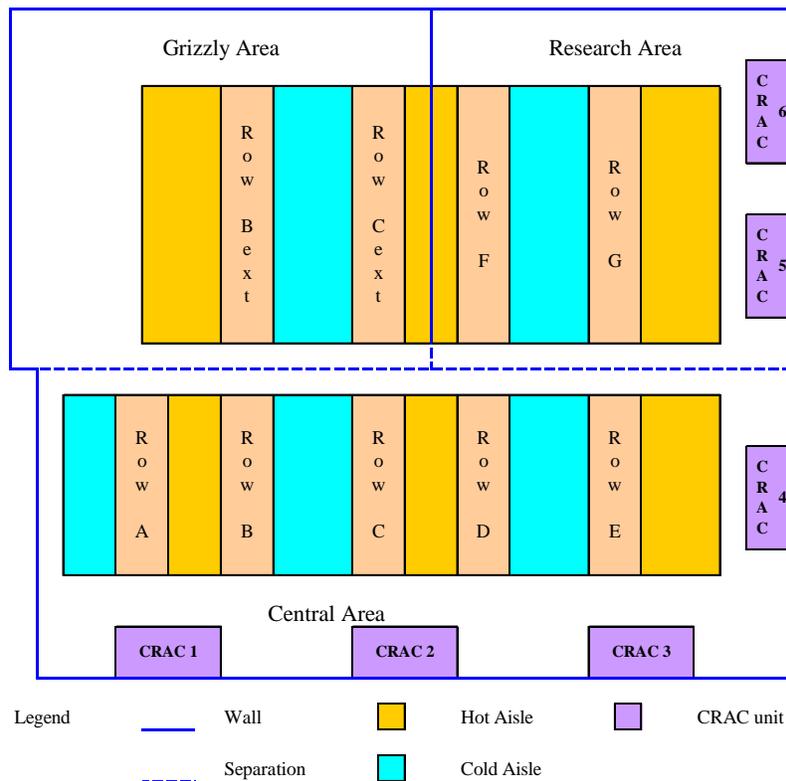


Figure 2: Plan View of the Data Center in HP Lab

Figure 3 illustrates the two rows of racks in the Research area of the HP Data Center. Each row consists of several computer racks, with a temperature sensor installed on top. We conducted the following experiments in this area while it was completely separated from other areas. Only rack inlet temperature and airflow from each vent tile were available in the Research area.

- Define the base case as a situation in which all the vent tiles are completely open. In the experiments dealing with the behavior of a single vent tile, close

each of the vent tiles in the 3 columns one by one, following the base case. In each specific situation, measure the rack inlet temperatures and airflow from the other open vent tiles.

- Define three vent tiles in the same row as a group of vent tiles. In the experiments dealing with the behavior of a vent group, close each of the vent groups one by one, following the base case. Measure the rack inlet temperatures and airflow from the other open vent tiles.

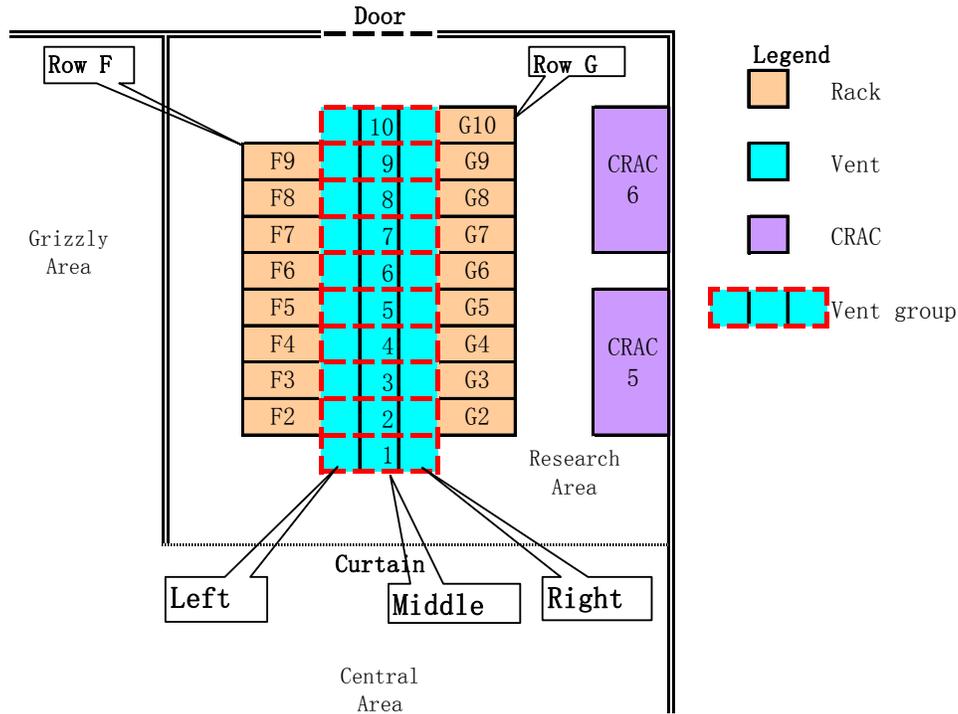


Figure 3: The Layout of Research Area in HP Data Center

2) Grizzly Area

We conducted similar experiments in the Grizzly area, where rack-power and inlet temperature was available.

- Define the base case as occurring when all the vent tiles are completely open. Close each group of vent tiles in the same row one by one, following the base case. Sample rack inlet temperatures and rack power every minute during the period. Airflow from each tile is not measured. Keep the research area isolated from the other areas in the HP Data Center for this series of experiments.

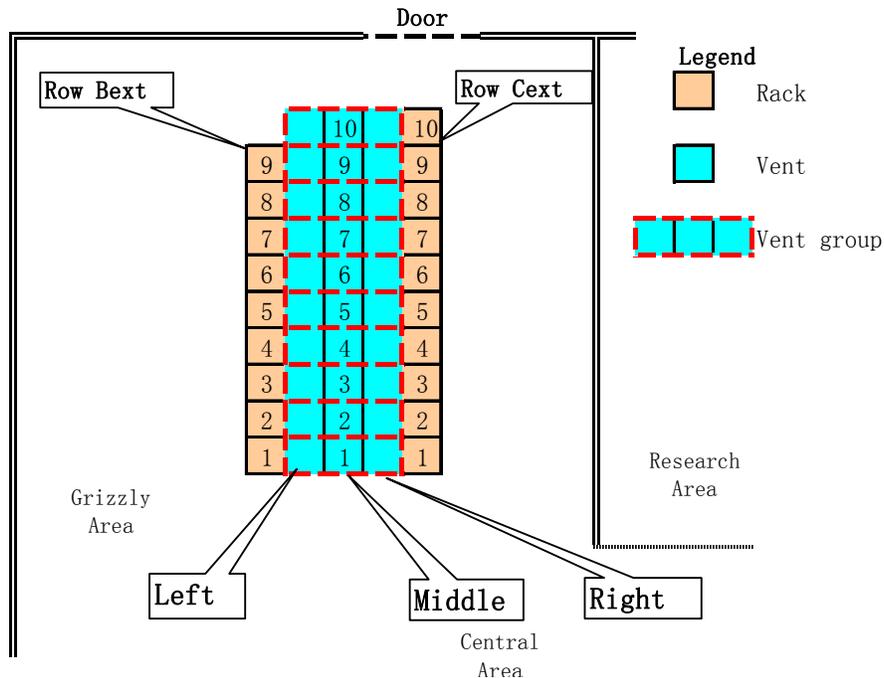


Figure 4: The Layout of Grizzly Area in HP Data Center

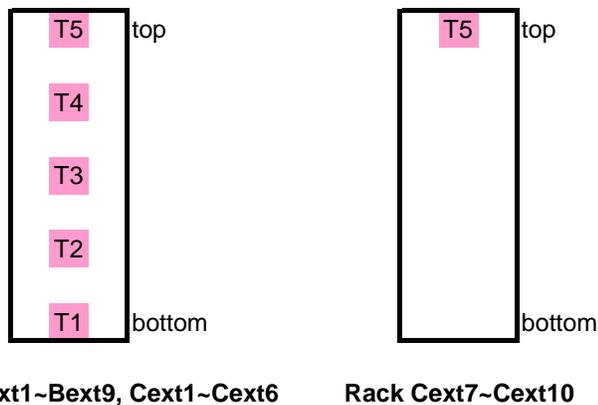


Figure 5: Location of Temperature Sensor at Each Rack (facing the racks)

3) Correlation Matrix

After collecting the results of the above experiments in the Research area and Grizzly area, we can generate correlation matrices between vent behavior and rack inlet temperature. Following is a small example demonstrating how the local control algorithm, based on a correlation matrix, works. Consider a simple case with only 3 racks and 3 vent groups in Figure 6, with each rack having one temperature sensor.

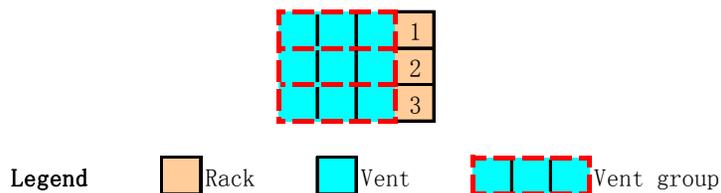


Figure 6: A Simple Cold Aisle

Assume that rack inlet temperature measured in experiments is as shown in Table 1.

	Base case	Close V1	Close V2	Close V3
Rack 1	21	24	20	20
Rack 2	22	20.5	24	22
Rack 3	23	21	22	25.5

Table 1: Rack Inlet Temperatures of the Experiments (Unit: °C)

Table 2 shows the corresponding correlation matrix, where values indicate the inlet temperature variation of each rack when closing a group of vent tiles. For example, from the base case, the inlet temperature of rack 1 will increase by 3°C when we close vent group 1.

	Close V1	Close V2	Close V3
Rack 1	3	-1	-1
Rack 2	-1.5	2	0
Rack 3	-2	-1	2.5

Table 2: Correlation Matrix of a Simple Cold Aisle (unit: °C)

Table 3 lists the inlet temperature for two new base cases. Assume the correlation matrix in can be reproduced. Therefore, we can predict the inlet temperature after closing any vent group. With different objectives, we can reach different operations. For example, if we want to decrease the maximum temperature, we do nothing for base case 1, and close vent group 3 for base case 2. Notice that based on the assumption of reproducibility, the correlation matrix can only produce one operation. Its operational ability is questionable in all base cases except when all vent tiles are open.

	Base #1	Prediction			Base #2	Prediction		
		Close V1	Close V2	Close V3		Close V1	Close V2	Close V3
Rack 1	22	25	21	21	24	27	23	23
Rack 2	24	22.5	26	24	22	20.5	24	22
Rack 3	21	19	20	23.5	21	19	20	23.5
Maximum	24	25	26	24	24	27	24	23.5
Minimum	21	19	20	21	21	19	20	22
Difference	3	6	6	3	3	8	4	1.5

Table 3: Original Temperatures and Predictions (Unit: °C)

RESULTS/DISCUSSION

1) Rack Inlet Temperature Distribution

Figure 7 is the rack inlet temperature distribution in the Grizzly area. The difference between the maximum and minimum temperature is around 6.5°C.

Local distinctions can't be affected by global settings. Figure 8 indicates the inlet temperature difference between two steady states with different tempera-

ture set-points. The set-point was increased by 5.6°C. As a result, the inlet temperature of most locations was increased by 3.5°C ~ 4.0°C.

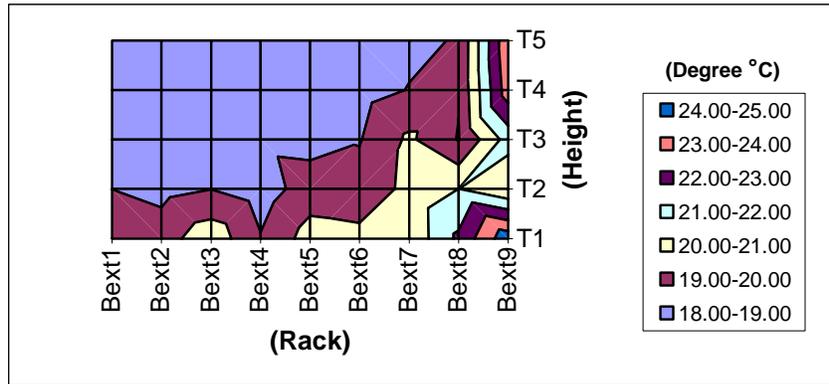


Figure 7: Temperature Contour of Row Bext in the Grizzly Area

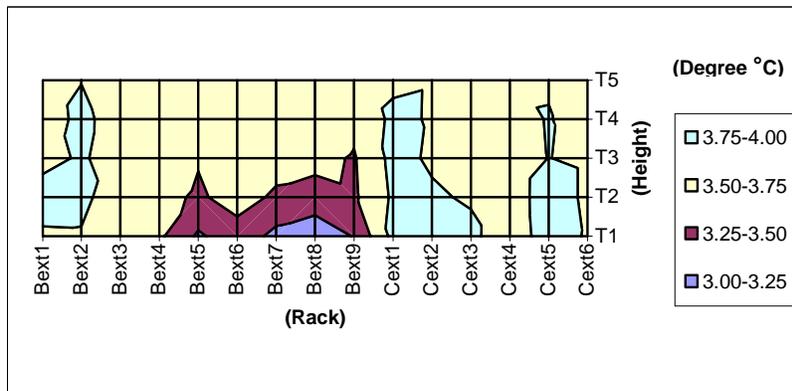


Figure 8: Contour of Temperature Difference between Two Statuses

2) Flow Analysis

(1) Single Vent Tile

Table 4 shows the total airflow of a cold aisle, corresponding to the experiments of a single vent tile in the research area. “m, l, r” represent the middle, left and right column of vent tiles respectively (Figure 3). Each value represents the total flow rate in a specific case, i.e. the value of one cell (“Close 9” and “Flow(r)”) is 14625, which means the total airflow from all the open vent tiles is 14625 CFM when we close vent tile r9.

The airflow of each vent tile was measured by flow hoods (backpressure compensated air balanced system model CFM-88), whose accuracy is about $\pm 3\%$ of reading ± 7 CFM from 100 to 2000 CFM. Turbulence also results in some variations in flow rate measurements, thus we estimate the uncertainty of airflow measurements to be $\pm 10\%$ of readings.

Coefficient of Variance (CV) is one of the indicators showing variation levels.

Its value is given by $CV = \frac{\text{standard error}}{\text{mean}}$. A higher CV implies a larger variation

among data. The CVs of total airflow are very small and far lower than the estimated uncertainty even though we include base cases in the calculation.

Thus total airflow won't be affected significantly by closing a single vent tile.

Table 5 lists the CV of airflow from the vent tiles in different ranges in the single vent experiments. There are three sections, each of which represents one series of experiments. For instance, 'Close middle' is the series of experiments closing one vent tile in the middle column, and 'Close 10' means closing the corresponding vent tile in row 10. When we close vent tile m10, the CV of airflow from the vent tiles in the left column is 4.5%, and the CV of airflow from the rest of the open vent tiles in all the three columns is 7.5%. These values correspond to the sub-columns "Left" and "Total". Most CVs are smaller than the estimated uncertainty. Especially, the small CV in sub-column "Total" indicates that the airflow distribution is quite uniform among all the open vent tiles. More importantly, the CV of each sub-column doesn't change much with a single vent tile closed. Therefore, it is not evident that closing one vent tile can change the airflow distribution from the other open vent tiles in the same column or even the whole cold aisle.

Case	Flow(m)	Flow(l)	Flow(r)
Base	14264	15210	15210
Close 10	14040	14314	14554
Close 9	14181	14285	14625
Close 8	14246	14284	14611
Close 7	14120	14488	14634
Close 6	14065	14463	14495
Close 5	14100	14289	14579
Close 4	13913	14432	14602
Close 3	14044	14602	14442
Close 2	14009	14489	14562
Median	14083	14448	14591
Average	14098	14486	14631
Standard Error	109	277	212
Coefficient of variance	0.80%	1.90%	1.40%

Table 4: Total Flow Rate of the Cold Aisle in the Research Area (Unit: CFM)

	Close middle				Close left				Close right			
	Left	Middle	Right	Total	Left	Middle	Right	Total	Left	Middle	Right	Total
Base	5.9%	5.9%	6.7%	8.5%	3.6%	7.0%	5.7%	7.5%	3.6%	7.0%	5.7%	7.5%
Close 10	4.5%	6.0%	4.9%	7.5%	4.8%	6.4%	5.8%	7.4%	4.3%	7.4%	5.9%	8.3%
Close 9	3.3%	8.1%	5.1%	8.4%	3.4%	9.0%	5.3%	8.5%	1.8%	7.7%	5.7%	7.9%
Close 8	3.7%	8.1%	6.8%	8.2%	2.9%	9.0%	5.6%	8.6%	3.0%	7.6%	5.3%	7.6%
Close 7	5.1%	7.9%	6.2%	8.7%	2.9%	8.1%	6.4%	9.0%	2.6%	9.2%	5.8%	8.6%
Close 6	2.2%	7.0%	6.3%	7.4%	3.0%	10.0%	6.4%	9.2%	2.9%	8.2%	4.7%	7.8%
Close 5	4.0%	7.3%	4.6%	7.7%	3.8%	7.3%	5.3%	7.5%	2.8%	8.0%	4.4%	7.8%
Close 4	4.9%	8.8%	5.5%	8.5%	4.6%	7.5%	5.0%	8.2%	3.4%	7.7%	5.3%	7.4%
Close 3	3.1%	5.6%	5.0%	7.6%	2.7%	6.4%	6.1%	7.8%	2.8%	7.0%	4.7%	6.9%
Close 2	2.5%	5.7%	4.0%	7.5%	3.8%	7.0%	3.8%	7.4%	3.8%	7.5%	5.3%	7.5%

Table 5: Coefficient of Variance for Experiments of Single Vent Tile

(2) Vent Groups

Table 6 shows the statistics of total airflow from the vent group experiments in the Research area. We conducted two series of experiments with different base cases (opening ratio =100%; 50%). The variability of total airflow is still much smaller than the estimated uncertainty, which indicates that closing one vent group has negligible influence on total airflow.

Table 7 lists the CV of the different categories in vent group experiments. For example, when we close vent group 10, the CV of airflow from the rest of the open vent tiles in the left and middle column, the right column, and all three columns are 5.7%, 8.4% and 7.1%, respectively. In the vent group experiments, we measured the airflow of two adjacent vent tiles in the left and middle column together. When calculating the CV of airflow from the vent tiles in these two columns, we assumed that the airflow from the left column was the same as that from the middle column. In the experiments with the base case at 100% opening ratio, some values of CV are large. Clearly a vent group has more influence on airflow distribution than one vent tile. With the opening ratio decreasing in the base case and the flow resistance increasing, this influence diminishes.

Case	Flow (CFM)	
	Base=100%	Base=50%
Base	9754	8484
Close 10	9422	8001
Close 9	9533	7973
Close 8	9425	8016
Close 7	9643	8038
Close 6	9693	8042
Close 5	9491	8013
Close 4	9503	7922
Close 3	9385	7887
Close 2	9400	7918
Close 1	9581	8001
Median	9503	8001
Average	9530	8027
Std Err	124	160
CV	1.30%	2.00%

Table 6: Airflow Statistics of Vent Group Experiments (Unit: CFM)

Although we cannot observe significant variation in the airflow distribution when one vent tile or vent group is closed, rack inlet temperature may change significantly. A single vent or vent group won't affect the pressure distribution under a plenum much since the supply air velocity is high in the plenum and closing several vent tiles doesn't change the flow resistance of the cooling system a lot. However, rack inlet temperature may be sensitive to vent configurations due to the slower air velocity and lower pressure above the plenum. Closing one or several vent groups could affect rack inlet temperature when the closed groups incorporate enough of the vent tiles contributing to the cool-

ing of a specific rack. Control algorithms for rack inlet temperature have to identify these relationships.

	Base=100%			Base=50%		
	Left+Middle	Right	Total	Left+Middle	Right	Total
Base	3.8%	7.7%	5.9%	5.7%	5.1%	5.6%
Close 10	5.7%	8.4%	7.1%	4.5%	5.7%	5.1%
Close 9	7.6%	13.4%	11.1%	5.8%	4.9%	5.5%
Close 8	9.4%	6.4%	7.8%	3.3%	5.8%	4.8%
Close 7	10.9%	7.5%	9.2%	5.1%	4.2%	5.4%
Close 6	9.4%	6.4%	7.8%	4.2%	3.5%	4.1%
Close 5	6.9%	4.7%	5.8%	3.9%	2.8%	3.5%
Close 4	8.2%	4.8%	6.6%	4.1%	4.1%	4.0%
Close 3	6.3%	5.2%	6.5%	4.1%	2.9%	3.4%
Close 2	5.5%	4.1%	5.3%	2.6%	2.2%	2.4%
Close 1	7.6%	3.1%	6.0%	3.2%	3.6%	3.4%

Table 7: Coefficient of Variance for Vent Group experiments

3) Correlation Matrix

(1) Research Area

▪ Single Vent

Previous airflow data indicated that neither a single vent nor a vent group has strong influence on airflow distribution in the Research area. Here we conducted a series of experiments to generate a correlation matrix and discovered the relationship between a single vent and rack inlet temperature.

Precon thermistors (model ST-R3R) are used to measure the rack inlet temperature in the Research area. Their accuracy is $\pm 0.3^{\circ}\text{C}$, and A/D resolution is about 0.2°C . So the uncertainty of temperature measurements is about 0.5°C (0.9 F). A threshold of 0.5°C is therefore chosen to determine if a change in temperature can be related to the manipulation of vent tiles. Any temperature variation larger than the threshold is significant and results from the variation of vent configurations. In the following tables containing the correlation matrices, the cells in light orange indicate that closing a vent will increase the rack inlet temperature by more than 0.5°C , and the cells in turquoise indicate that closing a vent will decrease the rack inlet temperature by more than 0.5°C .

The behavior of a single vent tile in these three columns varies a lot. To validate the basic assumption of reproducibility, we selected some vent tiles (r2, m2, l2, m4 and r5) to repeat the experiments. Table 8 and 9 are the results of the original and repeated experiments.

Each cell in the correlation matrices belongs to one of the three categories: positive ($x > 0.5^{\circ}\text{C}$), negative ($x < -0.5^{\circ}\text{C}$) or negligible ($-0.5^{\circ}\text{C} < x < 0.5^{\circ}\text{C}$). If an "original" cell and its "repeated" cell fall into the same category, they are consistent; otherwise they are inconsistent. Table 10 shows that only about half of the cells are consistent in both of the experiments. We notice that most consistent cells are "negligible", which cannot be used in the control algorithm, because "negligible" cells mean that closing these vent tiles only has negligible effects on the corresponding inlet temperature.

	Close r2	Close m2	Close l2	Close m4	Close r5
F9	-0.32	-0.01	-1.73	1.29	0.07
F8	-0.17	0.58	-2.12	1.21	-0.31
F7	-0.88	-0.17	0.00	-0.20	-0.56
F6	-0.45	-0.31	-0.10	-0.24	-0.35
F5	-0.59	-0.74	-0.08	-0.24	-0.45
F4	0.95	-0.45	2.46	-0.79	0.48
F3	-0.48	-0.60	2.58	-0.84	-0.04
G10	0.27	0.94	-0.32	0.52	-0.34
G9	1.39	0.90	-0.70	0.76	-0.22
G8	0.09	0.10	-0.19	-0.04	-0.24
G7	0.18	0.18	-0.46	0.01	-0.10
G6	0.13	0.17	-0.03	0.13	0.34
G5	-0.08	-0.11	-0.27	-0.07	0.31
G4	-0.11	0.11	0.32	0.55	0.03
G3	0.08	-0.13	1.16	0.27	0.25
G2	0.43	-0.46	-0.27	-0.70	0.55

Table 8: Part of Correlation Matrix of Single Vent Experiments (I)

	Close r2	Close m2	Close l2	Close m4	Close r5
F9	-0.94	-0.74	-0.41	-0.87	-0.66
F8	-0.73	-0.80	-0.58	-0.48	-0.50
F7	-0.24	-0.24	-0.24	-0.24	-0.31
F6	-0.29	-0.29	-0.29	-0.29	-0.29
F5	-0.50	-0.50	-0.50	-0.50	0.22
F4	-0.14	-0.10	0.10	-0.11	-0.48
F3	0.03	0.08	0.48	-0.17	-0.36
G10	-0.22	-0.04	-0.21	0.20	0.15
G9	-0.58	-0.66	-0.48	-0.56	-0.68
G8	0.06	0.03	-0.06	0.03	0.00
G7	-0.75	-0.75	-0.75	-0.56	-0.56
G6	0.32	0.32	0.32	0.21	0.36
G5	-0.04	-0.07	-0.17	-0.18	-0.11
G4	-0.13	0.00	-0.13	0.24	0.01
G3	0.00	0.11	0.08	0.27	-0.08
G2	0.11	-0.06	-0.01	-0.08	-0.20

Table 9: Part of Correlation Matrix of Single Vent Experiments (II)

We have learned that manipulating a single vent tile barely changes the air-flow distribution within the same cold aisle. Sometimes a single vent tile can affect some other vent tiles. But such a relationship in airflow is loose and unstable, so is the relationship between a single vent tile and inlet temperature.

Also, even after the system reaches a steady state, inlet temperature could fluctuate instead of keeping constant. The temperature fluctuation at each sensor can be different, which may result in a larger error when we calculate the difference between two temperature readings in order to generate the correlation matrix.

	Consistence	Inconsistence
Close r2	10	6
Close m2	10	6
Close l2	9	7
Close m4	6	10
Close r5	10	6
Total	45	35
Percentage	56%	44%

Table 10: Reproducibility of Correlation Matrices in Single Vent Experiments

- Vent Groups

The behavior of a vent group in the Research area was explored. Table 11 lists the different settings used by all the CRAC units in the experiments. “Tset” is the temperature set-point; “VFD” is the percentage of fan speed, ranging from 0% to 100%; “Base Opening” is the opening ratio of vent tiles in the base case; “Repeat?” indicates whether the experiments with this setting were repeated and if the answer is yes, “Partial?” indicates whether the experiments were partially repeated or completely repeated. Table 12 indicates that the correlation matrix from the vent group experiments has better reproducibility than that from the single vent experiments.

Setting	Tset (°C)	VFD (%)	Base Opening (%)	Repeat?	Partial?
1	15.56	60	100	Yes	Yes
2	15.56	40	100	Yes	No
3	15.56	40	50	No	-
4	15.56	80	100	No	-

Table 11: Settings of Experiments of Vent Groups

	Setting 1		Setting 2	
	consistence	inconsistence	consistence	inconsistence
Close 1	14	2	11	5
Close 2	-	-	12	4
Close 3	-	-	10	6
Close 4	-	-	11	5
Close 5	-	-	10	6
Close 6	16	0	10	6
Close 7	14	2	11	5
Close 8	15	1	11	5
Close 9	13	3	10	6
Close 10	14	2	7	9

Total	86	10	103	57
Percentage	90%	10%	64%	36%

Table 12: Reproducibility of Correlation Matrices in Vent Group Experiments

The correlation matrix at setting 1 is more reproducible, but it was only partially repeated. The correlation matrix at setting 2 was completely repeated and with lower reproducibility. In general, a vent group has more predictable influence than a single vent tile and the correlation matrix is more stable.

The comparisons in Table 12 are based on the correlation matrices calculated by single temperature measurements in the steady state, so the influence of temperature fluctuation can be large. Table 13 compares two correlation matrices calculated by averaging the inlet temperature at the steady states, assuming the transient states would last 10 minutes for each situation.

	consistence	inconsistence
Close1	13	3
Close2	15	1
Close3	12	4
Close4	11	5
Close5	14	2
Close6	13	3
Close7	9	7
Close8	11	5
Close9	10	6
Close10	12	4
Total	120	40
Percentage	75%	25%

Table 13: Reproducibility Check (Setting 2; Average Temperature)

The reproducibility of the correlation matrix is improved by using average temperature. About 25% of all the cells are inconsistent. This may be due to other factors. For example, some temperature sensors in the Research area were later found to have low accuracy. Their accuracy was $\pm 0.3^{\circ}\text{C}$; however, this drifts significantly when it is not frequently calibrated .

(2) Grizzly Area

In the Grizzly area, we conducted vent group experiments with setting 2 (Table 11) and with all VDSs at 85% rather than the 40% indicated in the table. Time series analysis shows that there is no strong relationship between rack power and inlet temperature. Thus we can think of significant variations in inlet temperature as a function of vent tile configurations during the experiments. We calculated the correlation matrices by averaging the temperature in steady states, assuming the transient state lasts 10 minutes. Table 14 shows that the reproducibility of the correlation matrix in the Grizzly area is much better than that in the Research area.

There are two factors which account for the improved reproducibility of the correlation matrix in the Grizzly area.

1. Layout. The supply air in the Grizzly area mainly comes from CRAC 1~3, while that in Research area comes from CRAC 5 and 6. Thus the Grizzly area is further away from the CRAC units than the Research area. Correspondingly, it has less turbulent flow under the floor. Thus the influence of vent groups on inlet temperatures is more stable and reproducible, which make it possible to use the correlation matrix for local temperature control.
2. Reliable sensors. 1-wire sensors (Dallas Semiconductor, model DS18S20) were deployed in the Grizzly area. Their accuracy is $\pm 0.5^{\circ}\text{C}$ from -10°C to $+85^{\circ}\text{C}$, and the A/D resolution is about 0.2°C . Since the 1-wire sensors in Grizzly area were new and had been calibrated recently, they were more reliable than the sensors in the Research area.

	All sensors		Top sensors	
	consistence	inconsistence	consistence	inconsistence
Close1	77	2	19	0
Close2	73	6	16	3
Close3	77	2	19	0
Close4	78	1	19	0
Close5	72	7	17	2
Close6	73	6	16	3
Close7	75	4	17	2
Close8	77	2	18	1
Close9	77	2	19	0
Close10	70	9	15	4
Total	749	41	175	15
Percentage	95%	5%	92%	8%

Table 14: Reproducibility of Correlation Matrices (Setting 2; Grizzly Area)

4) Control Simulation

We have calculated the correlation matrices from vent-group experiments in the Grizzly area and proved the reproducibility of the correlation matrix is good in this area. Before we experimentally implement the control algorithm of correlation matrix, it is better to run simulations first for the following purposes:

1. Take advantage of the simulation to generate any possible temperature distributions in the cold aisle;
2. Learn the effects of this algorithm quickly.

However, there are some drawbacks of control simulations:

1. The simulation performs only one actuation starting from each base case. In each simulation the program randomly generates the inlet temperature distributions of base cases. Based on this distribution, the algorithm generates the new vent tiles configuration.
2. No information about the real effects. It doesn't simulate the physics of the system, such as temperature and airflow.

Thus, we only use this control simulation to demonstrate how the control algorithm responds to different inlet temperature distributions of the base cases.

There are two correlation matrices from repeated experiments. With the following rules, we reach the final correlation matrix for the control algorithm:

1. Use the consistent cells from both experiments, which have significant influence on temperatures, in the control algorithm. For these cases, we simply average the influence in both experiments.
2. For other consistent cases which do not have significant influence on temperatures, we still average the influence in both experiments but they won't be used in the control algorithm.
3. For inconsistent cases, we use the worst case from the two experiments in the final correlation matrix to make the correlation matrix conservative. They can't be used in the control algorithm yet.
4. Some inconsistent cases have been intentionally placed in the category of consistence and used in the control algorithm: if one cell shows positive influence, its counterpart shows negligible influence but it's close to the threshold 0.5°C . We define the secondary threshold as 0.05°C . If the counterpart is between 0.45°C to 0.5°C , we consider both cells are consistent and thus their average value will be used in the control algorithm. Similar rules are applied for the cases when one cell shows negative influence and the other shows negligible influence.

After we have constructed the correlation matrix, we simulate the temperature distribution across the racks, based on the previous experiment data. The temperature of each sensor is uniformly distributed within the range we observed in both experiments. Such simulated inlet temperature distributions may not represent the real temperature distribution. However, our objective is to demonstrate the control algorithm under any given temperature distribution, rather than focusing on the realistic temperature distribution.

The following are the rules for the control algorithm based on the correlation matrix:

1. Maximum temperature. Scan all the temperature sensors, and locate the maximum temperature. Check correlation matrix to see if there is any vent group(s) which can be used to lower the temperature at this location. Predict the temperature distribution after closing each of these vent groups, and select the vent group which can decrease the maximum temperature by a larger amount than any other vent group(s).
2. Minimum temperature. Use a concept similar to that described above, but instead of trying to lower the maximum temperature, increase the minimum temperature. Select the vent group which can increase the minimum temperature by a larger amount than any other vent group.

1000 independent simulations were generated for two algorithms towards the maximum temperature and minimum temperature respectively. In Table 15, two major columns "Towards Maximum" and "Towards Minimum" correspond to the two algorithms; each cell corresponds to the number of different cases, i.e. in the algorithm of maximum temperature, there are 877 cases out of 1000 simulations, in which maximum temperatures are success-fully decreased; there are 119 cases where minimum temperatures are in-creased; there are 766 cases where the difference between the maximum and minimum temperature is decreased.

	Towards Maximum		Towards Minimum	
	Yes	No	Yes	No
Maximum Decrease?	877	123	64	936
Minimum Increase?	119	881	77	923
Variation Decrease?	766	234	65	935
Operation?	877	123	77	923

Table 15: Results from Control Simulation (Correlation Matrix)

From the base case, it is more effective to decrease the maximum temperature than increasing the minimum temperature. Notice that only after the maximum inlet temperature is decreased, can we set the temperature set-points of the CRAC units higher and reduce energy consumption. So the correct order for the control algorithm is to decrease the maximum temperature first, then increase the minimum temperature without increasing the maximum temperature.

Since the simulated temperature distribution couldn't represent the real temperature distribution in the data center, the next step is to implement the control algorithm in the data center experimentally.

CONCLUSION

The issue of energy efficiency has become more and more important for data center cooling. We have experimentally explored rack inlet temperature distribution and analyzed the correlation matrix between vent tile configuration and rack inlet temperature. The results indicate that:

- Airflow distribution across vent tiles in the Research area of the experimental data center is quite uniform.
- Closing a single vent tile only shows negligible influence on rack inlet temperature and airflow distribution most of the time. Even though sometimes the influence is significant, it is unlikely to be stable and reproducible.
- Closing a group of vent tiles can change rack inlet temperature and airflow distribution to some extent. What's more, this influence is more stable and reproducible.
- In the control simulation, the control algorithm based on the correlation matrix of vent-group experiments is satisfying. Specifically, it effectively decreases the maximum temperature across the racks.

The results have demonstrated the possibility of using vent tiles as additional actuators for local temperature control and system optimization. The next step is using this information to implement a control algorithm for rack inlet temperature in the real data center.

ACKNOWLEDGMENTS

Many thanks to Jack Cattolico and Wayne Mack of Hewlett-Packard Laboratories for assistance with experiments and data collection. Also thanks to Ratanesh Sharma who helped a lot with temperature network and data logging. Finally, thanks to Kathe Gust for reviewing the report.

REFERENCES

- Bash, C. E., Patel, C. D., Sharma, R. K., 2003, "Efficient Thermal Management of Data Centers – Immediate and Long-Term Research Needs", HVAC&R Research, Vol.9, No.2, Apr 2003
- Boucher, T. D., Auslander, D. M., Bash, C. E., Federspiel, C. C., Patel, C. D., 2004, "Viability of Dynamic Cooling Control in a Data Center Environment"
- Friedrich, R., Patel, C. D., 2002, "Towards Planetary Scale Computing – Technical Challenges for Next Generation Internet Computing", THERMES 2002, Santa Fe, New Mexico
- Koplin, E. C., 2003, "Data Center Cooling", ASHRAE Journal, March 2003, p.46 - p.53
- Patel, C. D., Bash, C. E., Sharma, R. K., Beitelmal, A and Friedrich, R. J., "Smart Cooling of Data Centers", Proceedings of IPACK'03, The Pacific Rim/ASME International Electronic Packaging Technical Conference and Exhibition, Kauai, Hawaii, USA, IPACK-35059, July 2003
- Patel, C. D., Bash, C. E., Belady C., Stahl L., Sullivan D., 2001, "Computational Fluid Dynamics Modeling of High Compute Density Data Centers to Assure System Inlet Air Specifications", Proceedings of IPACK'01, The Pacific Rim/ASME International Electronic Packaging Technical Conference and Exhibition, July 8-13, 2001, Kauai, Hawaii, USA
- Sharma, R. K., Bash, C. E., Patel, C. D., Friedrich, R. J., Chase, J. S., 2003, "Balance of Power: Dynamic Thermal Management for Internet Data Centers"