



Local Temperature Control in Data Center Cooling: Part III, Application

Keke Chen, David M. Auslander, Cullen E. Bash, Chandrakant D. Patel
Enterprise Software and Systems Laboratory
HP Laboratories Palo Alto
HPL-2006-46
March 7, 2006*

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

Recent work has proposed a control algorithm to optimize local temperature distribution in data centers, based on the information of the correlation matrix between vent group configuration and rack inlet temperature [2] [3]. This paper presents the results of a case study in which this algorithm is applied to the appointed area of an experimental data center. Another algorithm which does not use the correlation matrix is also tested in the same area. These results indicate an improvement of rack inlet temperature distribution by using controllable vent tiles.

Local Temperature Control in Data Center Cooling:

Part III, Application

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

Recent work has proposed a control algorithm to optimize local temperature distribution in data centers, based on the information of the correlation matrix between vent group configuration and rack inlet temperature [2] [3]. This paper presents the results of a case study in which this algorithm is applied to the appointed area of an experimental data center. Another algorithm which does not use the correlation matrix is also tested in the same area. These results indicate an improvement of rack inlet temperature distribution by using controllable vent tiles.

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

INTRODUCTION

Data centers are playing a more and more important role in life with the development of computer technology. With the increasing power density in data centers and the exploding demands of centralized computing and data exchange, mechanical designers face more challenges in increasing the energy efficiency of data center cooling systems.

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.

Recent research has focused on designing and operating environmental control systems that properly provision cooling resources according to local heat loads. However, the situation at an individual computer location is usually not ideal. Cool air distribution depends on system layouts and flow patterns. It is likely that some computers will get more cool air than required, while others get less, instead of each computer getting the exact amount of cool air it re-

quires. As a result, rack inlet temperature distribution is not uniform. Thus, to ensure that all the computer equipment is safe, CRAC units must lower their temperature set-point resulting in some racks being over-cooled, which decreases the energy efficiency of system.

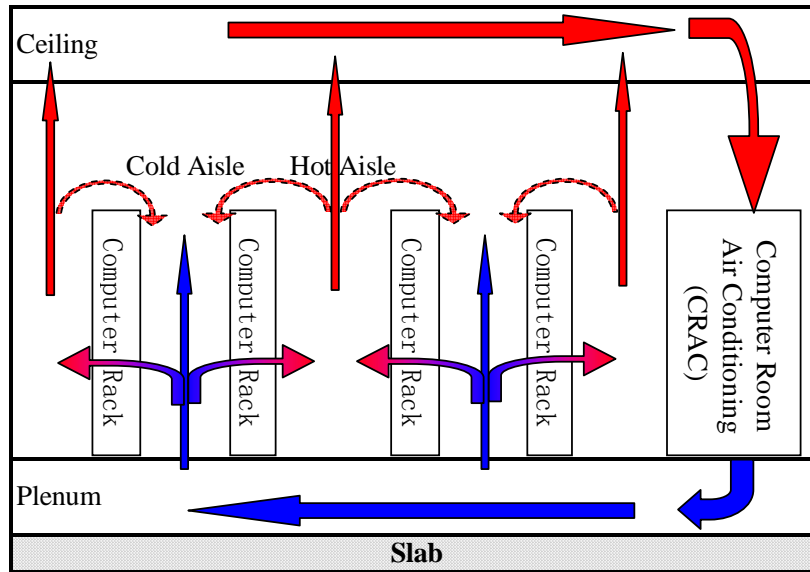


Figure 1: Schematic Section of a Data Center

Boucher et al. investigated global control strategies to optimize two convenient variable actuations of the CRAC units – supply air temperature setpoint and fan speed [1]. However, these variables are limited in their scope of influence: they can only affect the data center in a large scale, which makes it difficult to deal with local conditions distinctively. Previous experiments demonstrated that when temperature set-point was increased by 5.6°C , inlet temperature of most locations was increased by $3.5^{\circ}\text{C} \sim 4.0^{\circ}\text{C}$ in an experimental data center.

Researchers at Hewlett-Packard Laboratories introduced a “smart tile” – an actuated vent tile that can be commanded to vary its resistance to fluid flow. Previous work explored the relationship between vent group configuration and rack inlet temperature. The results indicated that, in the particular test area of the HP Data Center, that the correlation matrix between vent group configuration and rack inlet temperature had good stability and reproducibility. Current work will apply the local temperature control algorithm based on the correlation matrix [2]. Also, another algorithm which does not use the correlation matrix is proposed and tested in the same area.

EXPERIMENTATION

1) Layout

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 of the data center. The Research area is separated from the other two areas by a wall between the Research area and Grizzly area, and moveable curtains between the Research area and Central area and corresponding dampers under the

plenum. There is nothing to separate the Grizzly area and Central area physically. Figure 3 illustrates the two rows of racks in the Grizzly area. There is at least one temperature sensor in front of each rack's inlet.

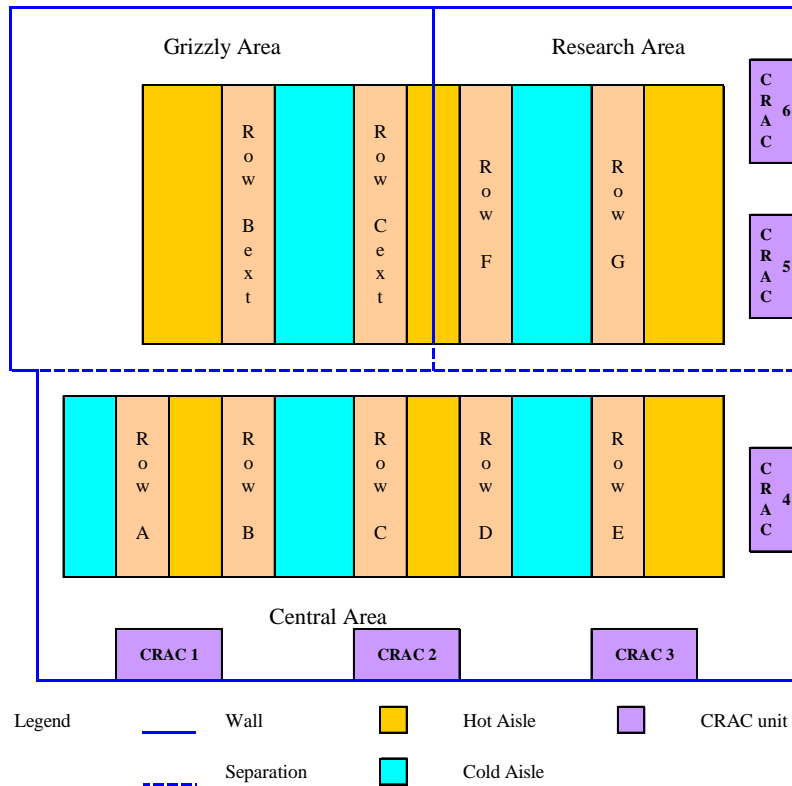


Figure 2: Plan View of the Data Center at HP Labs – Palo Alto

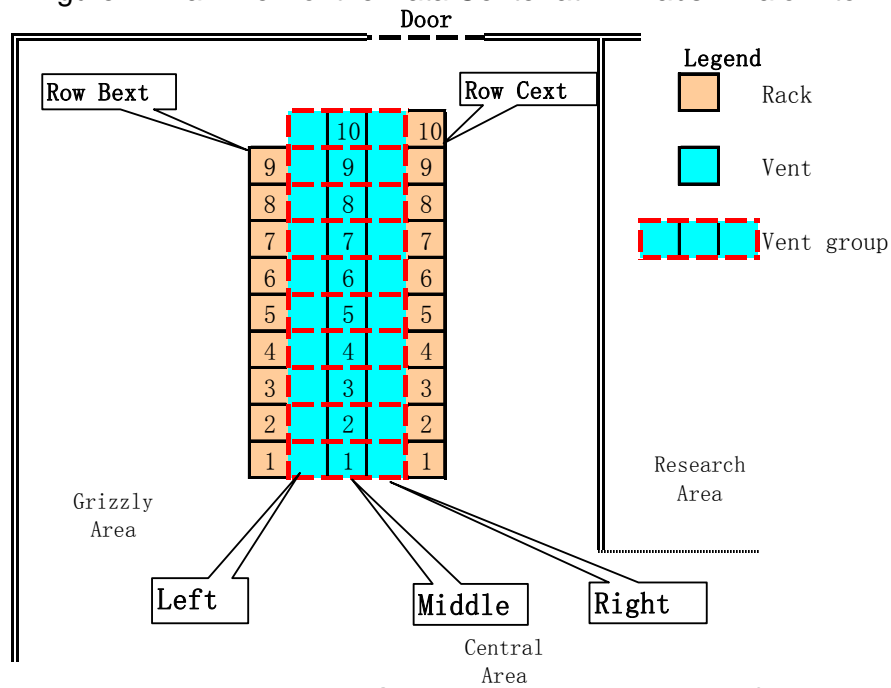
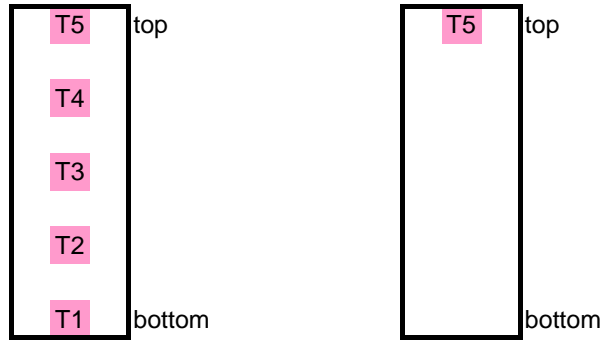


Figure 3: The Layout of Grizzly Area in the HP Data Center



Rack Bext1~Bext9, Cext1~Cext6 Rack Cext7~Cext10

Figure 4: Location of Temperature Sensor at Each Rack (facing the racks)

2) Correlation Matrix

Previous work has analyzed the relationship between vent tile configuration and rack inlet temperature. It verified the good reproducibility of the correlation matrix in the Grizzly area and demonstrated the possibility of implementing a local control algorithm based on such information [2].

Consider a simple case with only 3 racks and 3 vent groups in Figure 5, with each rack having one temperature sensor.




Figure 5: A Simple Cold Aisle

Assume that rack inlet temperatures are as 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 temperatures for two new base cases. Assume the correlation matrix in Table 2 can be reproduced. Therefore we can predict the

inlet temperatures after closing any vent group. With different objectives, we can achieve 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 those 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)

Above is a small example demonstrating how the local control algorithm based on the correlation matrix works. We implemented it in the Grizzly area of the HP Data Center. In addition to presuming reproducibility, we also assumed that the correlation matrix would still be valid with other base cases. Thus, when we implemented this local control algorithm, we forbid two adjacent vent groups from closing together. Otherwise the correlation matrix could be significantly different. Figure 6 is the flow chart of this local control algorithm.

One of the problems in the data center is hot air recirculation (as shown in Figure 1), resulting in upper sections of racks usually having higher inlet temperatures than lower ones. A successful approach should be able to make inlet temperature distributions more uniform so that we can increase the global temperature set-point. In detail, decreasing maximum temperatures is of more direct interest than increasing minimum temperatures. We tried to decrease the maximum temperature in the cold aisle with this algorithm as described below and in Figure 6..

First of all, collect the inlet temperature from all the sensors of the cold aisle. Locate the hottest spot. Predict the temperature distribution after manipulating each vent group using the correlation matrix. If there are one or more vent groups that can decrease the maximum temperature, manipulate the one that can produce the greatest maximum temperature reduction. Otherwise nothing will be changed and the local temperature distribution will not be optimized any further.

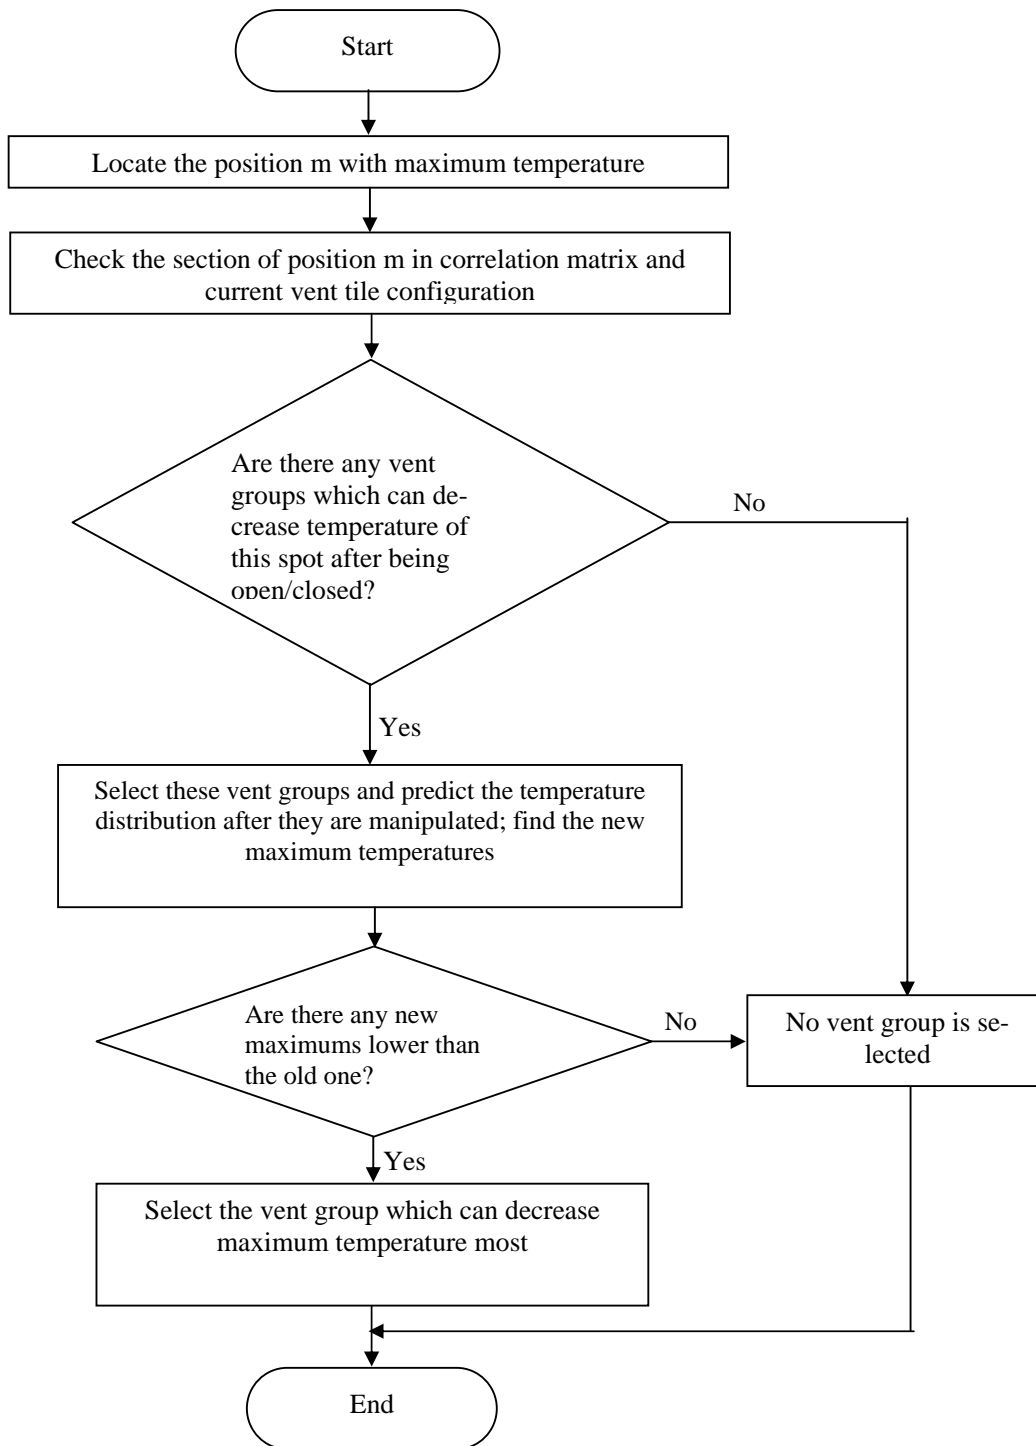


Figure 6: Flow Chart of Local Control Algorithm (Correlation Matrix)

3) Vent Group Family

Another algorithm is based on the process of identifying the vent group family for each rack. Recall that our objective is to make the inlet temperature distribution more uniform. Therefore, we want to maximize the supply airflow for the location with maximum temperature and minimize it for the location with

minimum temperature. For a specific rack, the supplier can include several groups of nearby vent tiles. In addition, the quantity of included vent groups for each family also varies from rack to rack. Thus, we define the vent group family as the closest two or three vent groups for each rack. The following list includes some considerations when manipulating vent groups to accomplish these objectives. Assume the maximum temperature exists at rack m, and minimum temperature exists at rack n.

- If rack m/n happens to be at the end of the cold aisle, i.e. rack Cext 10, at most two vent groups belong to its vent group family; otherwise three vent groups belong to that family.

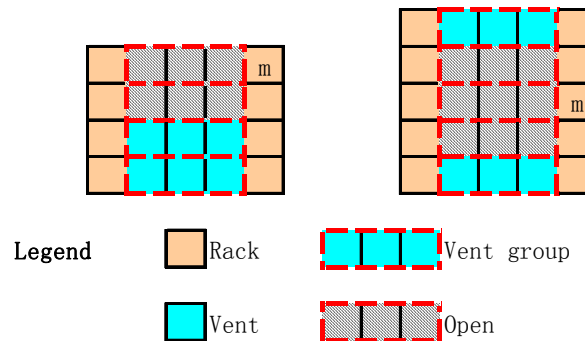


Figure 7: Relationship between a Rack and Its Vent Group Family

- If there are some overlaps between the vent group families for rack m and n, and these vent groups cannot be divided evenly for each rack, decreasing the maximum temperature always has higher priority, as shown in Figure 8.
 - a. The first plot represents the general situation in which vent group families do not overlap, and we can open three vent groups for rack m and close the other three vent groups for rack n.
 - b. In the second plot, vent group families are completely the same. Since maximum temperature has the higher priority, we open those three vent groups.
 - c. In the third plot, there are two overlapped vent groups. However, the layout is symmetric and we can open two vent groups while closing two.
 - d. In the last plot, there is one overlapping vent group and the layout is not symmetric. Again, we prioritize the maximum temperature reduction by opening three vent groups and closing two vent groups.

In a large data center with multiple cold aisles, the vent group family of a rack can even include vent groups in other cold aisles. However it is reasonable to assume that the algorithm is still valid in such a system since two cold aisles are usually not very close to each other. Figure 9 summarizes the algorithm with a flow chart. This algorithm was implemented in the Grizzly area of the HP Data Center and its effects were compared with that of the correlation matrix.




Figure 8: Control Schematic for Various Situations

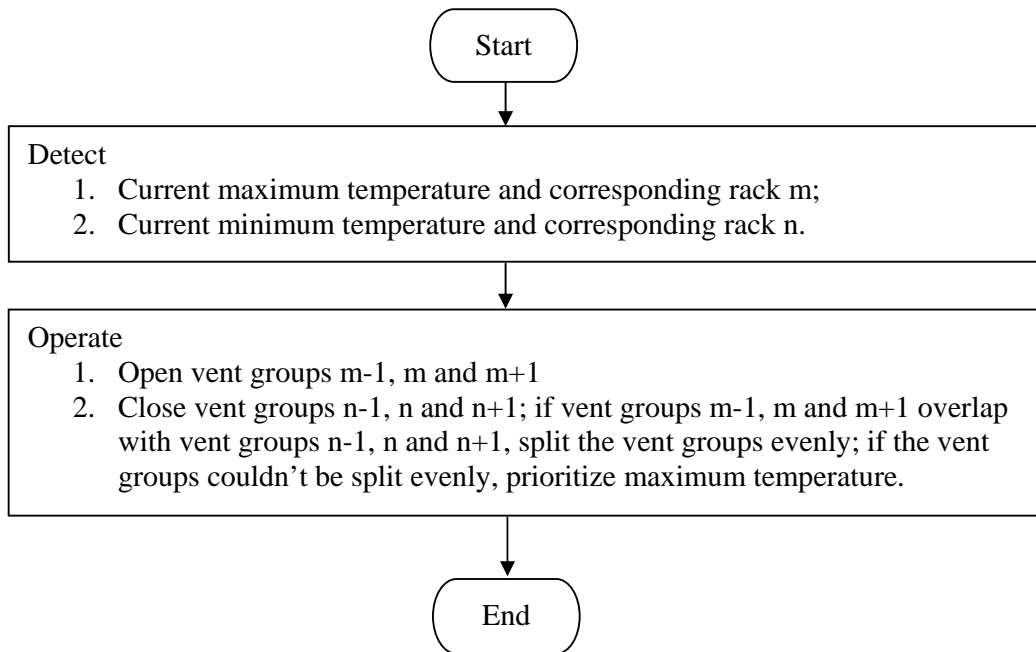


Figure 9: Flow Chart of Local Control Algorithm (Vent Group Family)

RESULTS/DISCUSSION

1) Correlation Matrix

Table 4 lists some of the results from the experiments of implementing the control algorithm based on the correlation matrix in the Grizzly area, and Table 5 summarizes the meaning of each column in Table 4. All the CRAC units were using the same settings: temperature setpoint is 15.56°C and VFD is 85%.

Step	Reset?	Original vent group status										Maximum			Minimum			Operation	
		V1	V2	V3	V4	V5	V6	V7	V8	V9	V10	T(°C)	Rack	Loc	T(°C)	Rack	Loc	Open	Close
1	Y	1	1	1	1	1	1	1	1	1	1	25.07	Bext9	T1	18.39	Bext1	T4	-	5
2		1	1	1	1	0	1	1	1	1	1	25.56	Bext9	T1	18.23	Bext5	T4	-	7
3	Y	0	0	1	1	1	1	1	1	1	1	25.48	Bext9	T1	18.52	Bext6	T5	-	5

4		0	0	1	1	0	1	1	1	1	1	25.81	Bext9	T1	18.19	Bext5	T4	-	7
5	Y	1	1	1	1	1	1	1	1	1	1	27.03	Cext7	T5	18.34	Bext1	T4	-	10
6		1	1	1	1	1	1	1	1	1	0	26.96	Cext7	T5	18.31	Cext5	T3	-	7
7		1	1	1	1	1	1	0	1	1	0	27.36	Cext7	T5	18.37	Bext1	T4	-	-
8	Y	1	1	1	1	1	1	1	1	0	0	27.37	Cext7	T5	18.26	Cext5	T3	-	7
9		1	1	1	1	1	1	0	1	0	0	26.32	Bext9	T5	18.15	Bext1	T4	9	-
10		1	1	1	1	1	1	0	1	1	0	25.54	Bext9	T1	18.33	Bext1	T4	-	5
11		1	1	1	1	0	1	0	1	1	0	25.03	Bext9	T1	18.37	Bext1	T4	-	-
12		1	1	1	1	0	1	0	1	1	0	25.25	Bext9	T1	18.40	Bext2	T3	-	-
13	Y	0	0	0	0	1	1	1	1	1	1	25.54	Bext9	T1	18.40	Bext6	T5	-	8
14		0	0	0	0	1	1	1	0	1	1	25.48	Bext9	T1	18.41	Bext6	T5	-	6

Table 4: Results from Experiments of Testifying the Correlation Matrix

Column	Description
Step	The order of each operation, during which system reaches the steady state and control algorithm is implemented once
Reset?	Whether the original vent group status is reset manually or following the previous step after operation (“Y” means vent groups are manually reset; blank means they are the same as those after last operation)
Original vent group status	“1” represents 100% opening ratio “0” represents 0% opening ratio
Maximum	The information of maximum temperature of this period
Minimum	The information of minimum temperature of this period
T	Temperature
Rack	The rack that this temperature (max/min) belongs to
Loc	The specific location of sensor that detects the extreme temperature
Operation	Information given by the control algorithm about how to manipulate the vent groups

Table 5: Summary of Columns in Table 4

Most of time maximum temperature didn't respond to the operation given by the algorithm, except in the tests that included steps 8, 9, and 10. Since this experiment was done 3 months later, after the correlation matrix was achieved, we believe that some adjustments in the plenum have already changed the correlation matrix. When we talk about the reproducibility of the correlation matrix, we only assume that it is stable within similar boundary conditions. It is also reasonable that the correlation matrix will change after major modifications in the system.

Minimum temperature was largely unaffected by the algorithm as well. The value of minimum temperature was very close to 18°C, although the location of minimum temperature varied. The supply air temperature set-point was also different than the minimum temperature. We discovered that since supply temperature was set at a really low value, 15.56°C, the CRAC units had started the dehumidification process and thus the actual supply temperature

was higher. In later experiments we set the temperature at a higher value and kept monitoring the real-time data of supply temperature to avoid this problem.

Generally speaking, the control algorithm based on the correlation matrix works. However, since the operation of the data center cannot remain unchanged at all times, we should endow the algorithm with some capacity to update the correlation matrix by following the experimental procedures. Also, the algorithm should be capable of receiving feedback from the system. If the results of one operation are counter-productive (i.e. the temperature increases instead of decreases), the algorithm must have the capability of reversing this operation and trying another. Automatic recalibration of the correlation matrix should also be included. Finally, further research concerning the correlation matrix is required to verify the effectiveness of the corresponding local control algorithm.

2) Vent Group Family

We used different settings (Table 6) to test the algorithm of vent group family. Results are shown in Table 7.

	Tset (°C)	VFD (%)			
		Step 1~8	Step 9~10	Step 11~19	Step 20~23
CRAC 1	20	80	85	70	90
CRAC 2	20	80	85	70	90
CRAC 3	20	80	85	70	90
CRAC 4	20	85	85	70	90
CRAC 5	21.67	70	85	70	90
CRAC 6	22.22	70	85	70	90

Table 6: Experiment Settings for Control Algorithm (Vent Group Family)

Step	Reset?	Original vent group status										Maximum			Minimum			Operation	
		V1	V2	V3	V4	V5	V6	V7	V8	V9	V10	T(°C)	Rack	Loc	T(°C)	Rack	Loc	Open	Close
1	Y	1	1	1	1	1	1	1	1	1	1	28.62	Cext8	T5	20.10	Bext6	T5	-	5,6
2		1	1	1	1	0	0	1	1	1	1	29.48	Cext8	T5	20.11	Bext2	T3	-	1,2,3
3		0	0	0	1	0	0	1	1	1	1	27.00	Cext7	T5	20.17	Bext6	T5	-	-
4		0	0	0	1	0	0	1	1	1	1	26.83	Cext7	T5	20.17	Bext6	T5	-	-
5	Y	0	0	0	1	1	1	1	1	1	1	25.47	Cext6	T2	20.00	Bext6	T5	-	-
6		0	0	0	1	1	1	1	1	1	1	25.51	Cext6	T2	20.07	Bext8	T5	-	8,9
7		0	0	0	1	1	1	1	0	0	1	27.75	Cext10	T5	20.07	Bext6	T5	9	5,6,7
8	Y	0	0	0	1	0	1	1	1	1	1	26.34	Cext6	T2	19.99	Bext6	T5	5	-
9	Y	0	0	0	1	0	1	1	1	1	1	27.28	Cext6	T2	20.10	Bext6	T5	5	-
10		0	0	0	1	1	1	1	1	1	1	26.26	Cext8	T5	19.90	Bext6	T5	-	5,6
11	Y	1	1	1	1	1	1	1	1	1	1	28.65	Cext8	T5	20.14	Bext5	T4	-	4,5,6
12		1	1	1	0	0	0	1	1	1	1	29.70	Cext8	T5	20.11	Bext2	T3	-	1,2,3
13		0	0	0	0	0	0	1	1	1	1	28.98	Cext8	T5	20.54	Bext7	T5	-	7
14		0	0	0	0	0	0	0	1	1	1	32.49	Cext8	T5	20.56	Bext8	T3	7	-
15	Y	1	1	1	1	1	1	1	1	1	1	28.01	Cext8	T5	20.16	Bext8	T5	-	-
16	Y	0	0	0	1	1	1	1	1	1	1	25.54	Cext8	T5	20.07	Bext8	T5	-	-
17	Y	0	0	0	0	0	1	1	1	1	1	27.38	Cext2	T5	20.14	Bext6	T5	1,2,3	6,7

18	Y	0	0	0	0	1	1	1	1	1	1	1	27.31	Cext1	T5	20.10	Bext6	T5	1,2	5,6,7
19	Y	0	0	0	1	1	1	1	1	1	1	1	25.50	Cext6	T2	20.06	Bext8	T5	-	8,9
20	Y	1	1	1	1	1	1	1	1	1	1	1	28.63	Cext8	T5	20.12	Bext2	T3	-	1,2,3
21		0	0	0	1	1	1	1	1	1	1	1	26.50	Cext10	T5	20.04	Bext6	T5	-	5,6,7
22	Y	0	0	0	0	1	1	1	1	1	1	1	26.74	Cext6	T2	20.05	Bext6	T5	-	-
23	Y	0	0	0	1	0	1	1	1	1	1	1	27.42	Cext6	T2	20.10	Bext6	T5	5	-

Table 7: Implementation of Control Algorithm (Vent Group Family)

The algorithm was tested in the following operations: 1~3, 5~6, 9, 11~13, 20. Other operations were set manually for comparison. Checking the trend of the maximum and minimum temperature during these operations, we will understand the effects of this algorithm.

(1) Maximum Temperature

The algorithm works well in decreasing the maximum temperature in general. Steps 4, 10 and 21 had lower maximum temperatures than step 1, 9, and 20, respectively.

Step 14 had a higher maximum temperature. However, the algorithm did suggest opening vent group 7 and going back to the status of step 13. Notice that the maximum temperature of step 13 was close to that of step 11. In this situation we could only say that the algorithm had no improvement.

Only step 7 showed an increasing trend, comparing to step 5. However, the algorithm was not completed: at step 7, it was trying to open V9 again. Since decreasing the maximum temperature has higher priority than increasing the minimum temperature, we didn't foresee a much higher temperature. However, the algorithm could make the system remain at the worse situation, as in the relationship of step 13 and 14.

Table 8 compares the actual maximum temperatures and the maximum temperatures used in the algorithm (maximum among all the sensors except T1).

Label	Max (Overall)			Max (Exclude T1)			Min(Overall)			Min (Exclude T1)		
	T(°C)	Rack	Loc	T(°C)	Rack	Loc	T(°C)	Rack	Loc	T(°C)	Rack	Loc
1	28.62	Cext8	T5	28.62	Cext8	T5	20.10	Bext6	T5	20.10	Bext6	T5
2	29.48	Cext8	T5	29.48	Cext8	T5	20.11	Bext2	T3	20.11	Bext2	T3
3	27.38	Bext9	T1	27.00	Cext7	T5	20.17	Bext6	T5	20.17	Bext6	T5
4	27.36	Bext9	T1	26.83	Cext7	T5	20.17	Bext6	T5	20.17	Bext6	T5
5	27.06	Bext9	T1	25.47	Cext6	T2	20.00	Bext6	T5	20.00	Bext6	T5
6	26.95	Bext9	T1	25.51	Cext6	T2	20.07	Bext8	T5	20.07	Bext8	T5
7	27.75	Cext10	T5	27.75	Cext10	T5	20.07	Bext6	T5	20.07	Bext6	T5
8	27.21	Bext9	T1	26.34	Cext6	T2	19.99	Bext6	T5	19.99	Bext6	T5
9	27.28	Bext9	T1	27.28	Cext6	T2	20.10	Bext6	T5	20.10	Bext6	T5
10	27.81	Cext6	T1	26.26	Cext8	T5	19.90	Bext6	T5	19.90	Bext6	T5
11	28.65	Cext8	T5	28.65	Cext8	T5	20.14	Bext5	T4	20.14	Bext5	T4

12	29.70	Cext8	T5	29.70	Cext8	T5	20.11	Bext2	T3	20.11	Bext2	T3
13	28.98	Cext8	T5	28.98	Cext8	T5	20.54	Bext7	T5	20.54	Bext7	T5
14	32.49	Cext8	T5	32.49	Cext8	T5	20.56	Bext8	T3	20.56	Bext8	T3
15	28.01	Cext8	T5	28.01	Cext8	T5	20.16	Bext8	T5	20.16	Bext8	T5
16	27.19	Bext9	T1	25.54	Cext8	T5	20.07	Bext8	T5	20.07	Bext8	T5
17	27.38	Cext2	T5	27.38	Cext2	T5	20.14	Bext6	T5	20.14	Bext6	T5
18	27.31	Cext1	T5	27.31	Cext1	T5	20.10	Bext6	T5	20.10	Bext6	T5
19	27.17	Bext9	T1	25.50	Cext6	T2	20.06	Bext8	T5	20.06	Bext8	T5
20	28.63	Cext8	T5	28.63	Cext8	T5	20.12	Bext2	T3	20.12	Bext2	T3
21	28.51	Cext6	T1	26.50	Cext10	T5	20.04	Bext6	T5	20.04	Bext6	T5
22	27.28	Cext6	T1	26.74	Cext6	T2	20.05	Bext6	T5	20.05	Bext6	T5
23	27.42	Cext6	T2	27.42	Cext6	T2	20.10	Bext6	T5	20.10	Bext6	T5

Table 8: Maximum and Minimum Temperature in the Cold Aisle

In most situations the maximum temperature was located at T5 – the highest temperature sensor of a rack, which indicated that the major factor of inlet temperature variation was hot-air recirculation. There were some exceptions in the location of maximum temperature. Some of them were at T2, which is the second lowest sensor. This is because at the bottom of each rack is a network switch, whose outlet is facing the cold aisle side. Thus the discharged hot air will make the temperature at T1 really high. Since there is no other equipment at the bottom of each rack that requires cool air, it was safe even though our control algorithm didn't consider T1 temperature sensor in this experiment. However, when T2 happened to be the location of maximum temperature and it was not in a rack at the end of a cold aisle, such as the case of step 6, the hot spot stemmed from the network switch underneath. It can become a bottleneck in this algorithm.

(2) Minimum Temperature

During the whole experiment, minimum temperature was approximately 20°C at several locations. This was consistent with the set-point of supply air temperature in Table 6. In general, as long as the supplied air is not totally blocked, inlet temperatures at some locations will be close to the set-point. Thus, we have to be cautious with the concept of “increasing the minimum temperature” in the algorithm. Unless we reduce the flow rate of supply air and aggravate the recirculation for every location, or increase the temperature set-point of the CRAC units, we cannot increase the minimum temperature among all the locations. Otherwise, trying to close the vent groups in the family corresponding to the rack with minimum temperature, in order to “increase minimum temperature”, can probably only increase the inlet temperature of this particular rack; it is possible that there are other racks with similar minimum temperatures.

However, this doesn't mean that trying to “increase the minimum temperature” is useless. Among those operations (1~3, 5~6, 9, 11~13, 20) where the control algorithm took actions, step 1, 11 and 20 started with the status that all the vent tiles were completely open. At the first step (1, 11 and 20) we cannot do any-thing about the maximum temperature because all vent tiles are open.

However, we can take some actions regarding the minimum temperature. These actions actually decrease the minimum temperature of a specific rack rather than that of the whole cold aisle. But they may indirectly increase the airflow from other vent tiles and thus decrease the maximum temperature. So the only reason for closing a vent group family in order to increase the minimum temperature is actually to decrease the maximum temperature. Thus, what we need to do is to identify the vent group family, corresponding to the minimum temperature, which can effectively decrease the maximum temperature.

In the experiments, only step 21 had lower maximum temperature than step 20; step 12 and step 2 both had higher maximum temperature than step 11 and 1, respectively. We know the loss of flow rate by closing several vent groups may evenly redistribute among the rest of the open vent groups in the same cold aisle; or it may partly transfer to cold aisles nearby. So we have no accurate information about which racks will be benefited by closing some specific vent groups. From the comparison of step 1 and 2, 11 and 12, we find that closing V4, V5 and V6 didn't increase the airflow of vent groups contributing to Cext8, where the maximum temperature was located. However, comparing step 1 and 5, 11 and 16, we notice that closing V1, V2 and V3 could decrease the maximum temperature in Cext8 effectively.

Table 8 also compares the actual minimum temperatures and the minimum temperatures used in the algorithm. Because of the existence of network switches at the bottom of each rack, although T1 is the nearest position to the vent tiles, it's not likely for the minimum temperature to exist in T1. As mentioned above, the minimum temperature should be close to the temperature set-point of the CRAC units. As long as not all the vent tiles are closed, the location of minimum temperature is not unique. This is very critical to the algorithm. At step 5, no action was conducted. The temperature data indicated that the status of step 5 and 6 were close to each other. However the algorithm detected that the minimum temperature had moved and made an inappropriate decision. Thus, the control algorithm should have the capacity to deal with multiple locations having the same minimum temperature.

(3) Learning Ability

This algorithm has to become smarter! It is based on the assumption that each rack is supplied with cool air by the nearest two or three vent groups, which doesn't completely match reality. When implementing the algorithm in a specific cold aisle, we need to give it some learning capacity, so that it can continuously adjust the assumptions to reflect reality.

a. Temperature history

The algorithm should always make use of temperature history, and have the ability to reverse inappropriate operations. For example, after the operation of step 1, the algorithm detected a higher maximum temperature at the same location. However, it took the action of closing the vent groups of minimum temperature, instead of reversing the previous inappropriate actions. Step 3 reached a lower maximum temperature than step 1 and 2, and the status remained unchanged at step 4. Step 5 had a lower maximum temperature than

step 4, which was achieved by reversing the inappropriate actions of step 1. Clearly, if the algorithm always compared the current status and previous status, it could have a correct priority in manipulating the vent groups.

b. Multiple minimum temperatures

The algorithm should be able to deal with multiple minimum temperatures. Maximum temperature is affected by recirculation and other factors, and it's unlikely for several locations to have the same maximum temperature. However, as long as not all the vent tiles are shut off, we can detect the same minimum temperature at several locations, and the value of minimum temperature will fluctuate around the supply temperature set-point. The algorithm will have multiple choices to close vent groups for minimum temperatures. A trial-and-error test is required to select the vent groups for maximum benefit.

c. Self-learning

The algorithm should be able to learn the characteristics of a specific cold aisle by itself. This capacity can help the algorithm reach to an optimal settings quickly. By comparing step 1 to 4, we learn that the racks close to other cold aisles are less influenced by manipulating vent groups, i.e. Bext1 and Cext1 might even be supplied cool air by other cold aisles: even though vent groups 1 to 3 were closed, inlet temperatures of these two racks never reached the maximum value across the cold aisle. On the other hand, the racks on the other ends of a cold aisle (i.e. Cext8 ~ Cext10) might be more sensitive to more vent tiles, i.e. besides vent group 7, 8 and 9, rack Cext8 could be influenced by vent group 7 or even 6. Step 2 had a higher maximum temperature than step 1, while step 5 had a lower maximum temperature. It seemed that V5 and V6 contributed to rack Cext8; V1 to V3 didn't; and closing V1 to V3 could indirectly contribute to rack Cext8. After learning from such histories, when the algorithm detects the maximum temperature at Cext8 and multiple minimum temperatures at Bext6 and Bext2 again, it should try to close V1, V2 and V3 first, rather than closing V5 and V6. Some similar cases are step 11 and 12, step 15 and 16, step 20 and 21. This capacity can also be summarized as an identification process of vent group family for each rack. Similar to the process for setup and recalibration of the correlation matrix, this algorithm starts from an initial vent group family for each rack in a specific cold aisle. During the control process it continues to update the family for each rack..

(4) Flow Domain

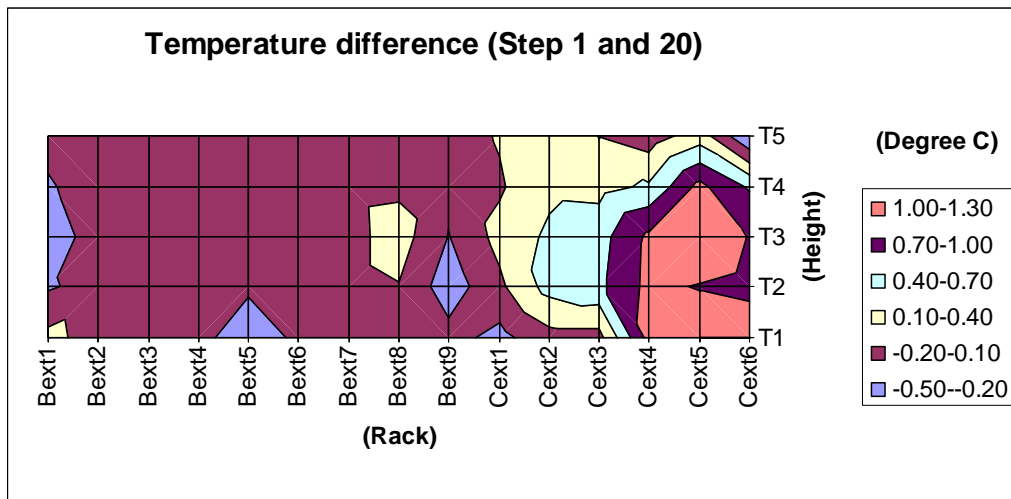
Table 9 compares the maximum and minimum temperature with same configuration of vent groups but different VFD settings. For convenience, we name the four different settings used in the experiments as "S1" to "S4" according to Table 6.

Label	Setting	Original vent group status										Maximum			Minimum		
		V1	V2	V3	V4	V5	V6	V7	V8	V9	V10	T(°C)	Rack	Loc	T(°C)	Rack	Loc
1	S1	1	1	1	1	1	1	1	1	1	1	28.62	Cext8	T5	20.10	Bext6	T5
11	S3											28.65	Cext8	T5	20.14	Bext5	T4
15												28.01	Cext8	T5	20.16	Bext8	T5

20	S4												28.63	Cext8	T5	20.12	Bext2	T3											
5	S1												25.47	Cext6	T2	20.00	Bext6	T5											
6													25.51	Cext6	T2	20.07	Bext8	T5											
10	S2	0	0	0	1	1	1	1	1	1	1	1	26.26	Cext8	T5	19.90	Bext6	T5											
16	S3												25.54	Cext8	T5	20.07	Bext8	T5											
19													25.50	Cext6	T2	20.06	Bext8	T5											
21	S4												26.50	Cext10	T5	20.04	Bext6	T5											
8	S1												0	0	0	1	0	1	1	1	1	1	1	26.34	Cext6	T2	19.99	Bext6	T5
9	S2																							27.28	Cext6	T2	20.10	Bext6	T5
23	S4	27.42	Cext6	T2	20.10	Bext6	T5																						
18	S3	0	0	0	0	1	1	1	1	1	1	1	27.31	Cext1	T5	20.10	Bext6	T5											
22	S4												26.74	Cext6	T2	20.05	Bext6	T5											
3	S1	0	0	0	1	0	0	1	1	1	1	1	27.00	Cext7	T5	20.17	Bext6	T5											
4													26.83	Cext7	T5	20.17	Bext6	T5											

Table 9: Comparison of Temperature Distributions at Different Settings

As illustrated before, there should be multiple locations with the same minimum temperature and the one listed in the table is only one of them. We notice that the maximum temperature could be different in both value and location while VFD settings were changed. Some cases indicated small variations in maximum and minimum temperature while they were using different settings, such as step 1 and 20; on the contrary, some cases had relatively large variations while they were using the same settings, such as step 11 and 15. We compared the contours of temperature difference at each location in Figure 9. Apparently, the second contour indicates that the status of step 11 was very close to that of step 15, while the first one suggests the temperature distribution of step 1 was quite different than that of step 20. Thus, the algorithm is proved to be effective under these different settings. Compared to the correlation matrix that is only valid within a relatively stable setting, this algorithm is more robust.



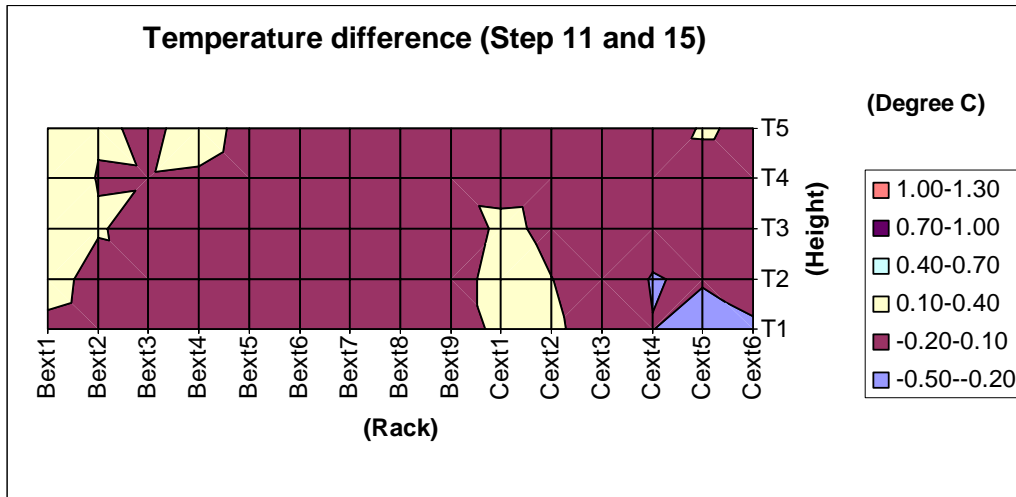


Figure 9: Contours of temperature difference from two comparisons

3) Data Center Cooling Scheme

Based on the above discussion, a modified algorithm of vent group family can be summarized by a flow chart as shown in Figure 10. First of all, collect inlet temperature readings from all the sensors in the cold aisle and locate the maximum and minimum temperatures. If the maximum temperature is larger than that in previous states, we have to reverse the last operation. Otherwise, open the vent groups in the family of the rack with the maximum temperature; try to close the vent groups in the families of the racks with the minimum temperature, and then select to close the vent groups of the family that can furthest lower maximum temperature. The flow chart in Figure 10 represents one operation, and the process continues until no more vent groups are manipulated. Then the inlet temperature distribution is optimized.

Table 10 lists the meaning of variables used in Figure 10.

Variable	Description
$TMAX_t$	Maximum temperature at time t
$TMIN_t$	Minimum temperature at time t
m	The rack with maximum temperature at time t
n_1, n_2, \dots, n_k	The racks with maximum temperature at time t
x_1, \dots, x_{kx}	Vent group family of the rack with maximum temperature at time t-1
y_1, \dots, y_{ky}	Vent group family of the rack with minimum temperature at time t-1

Table 10: Descriptions of Variables in Figure 10

Figure 11 suggests the interaction between the global control algorithm and the local control algorithm within a data center cooling system.

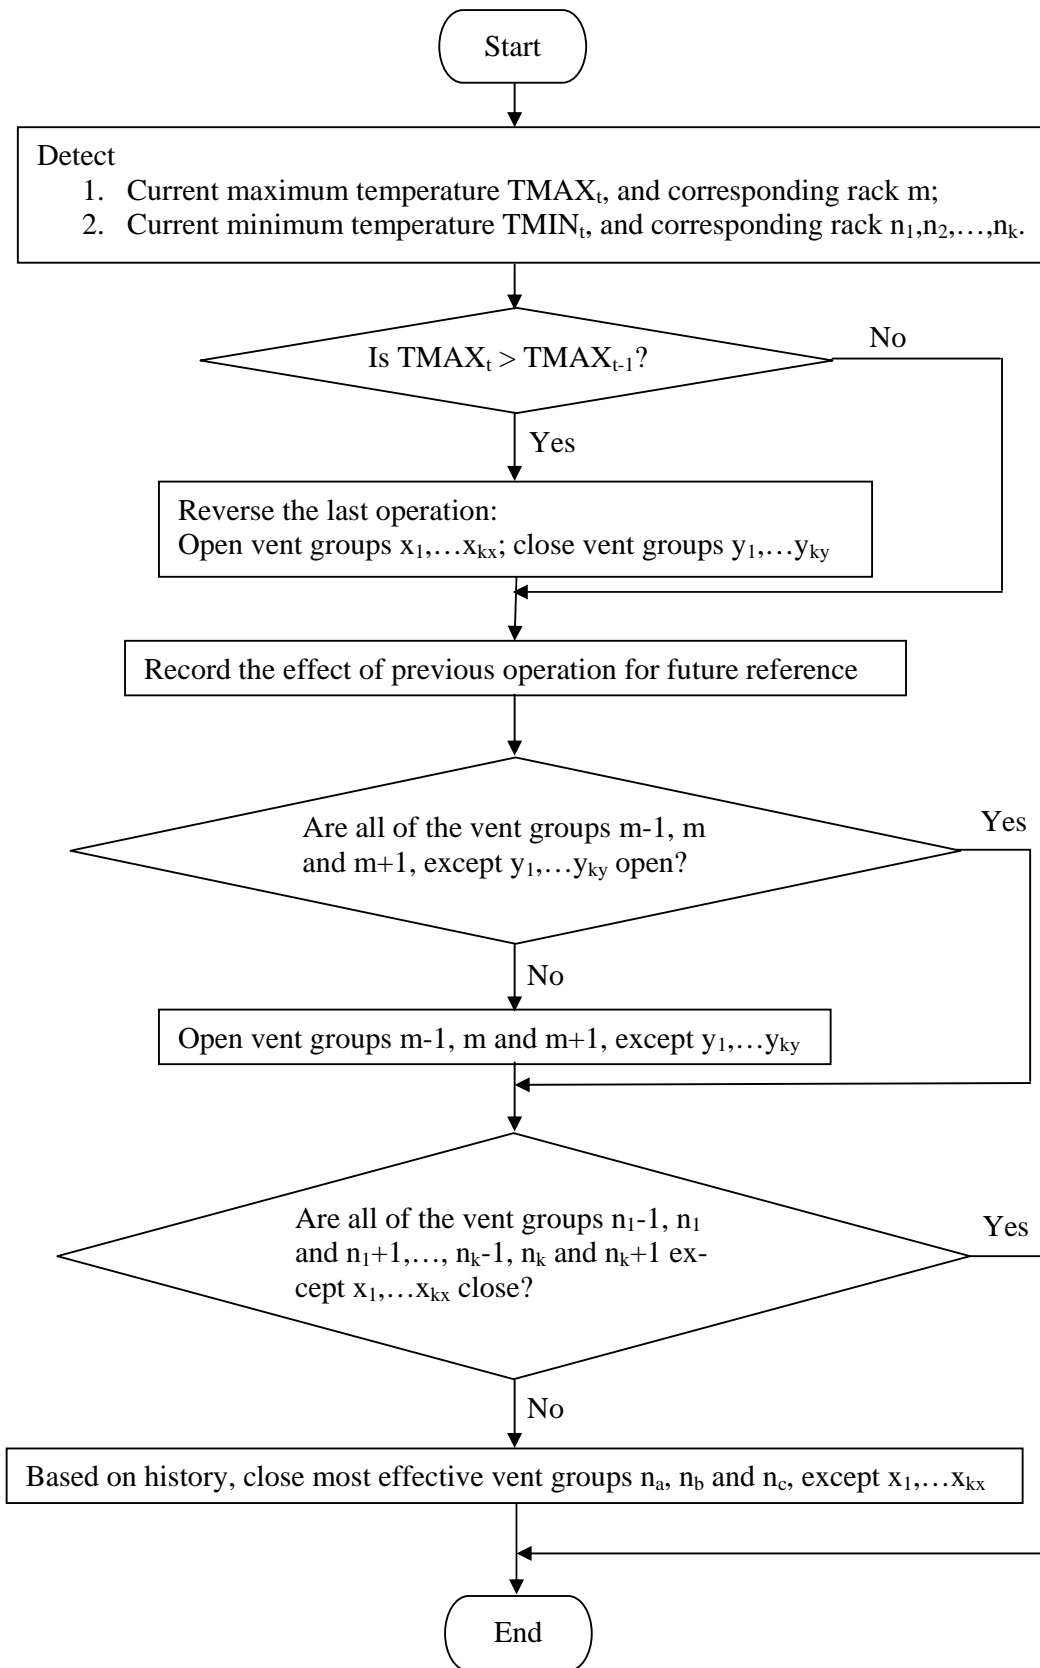


Figure 10: Revised flow chart of local control algorithm (vent group family)

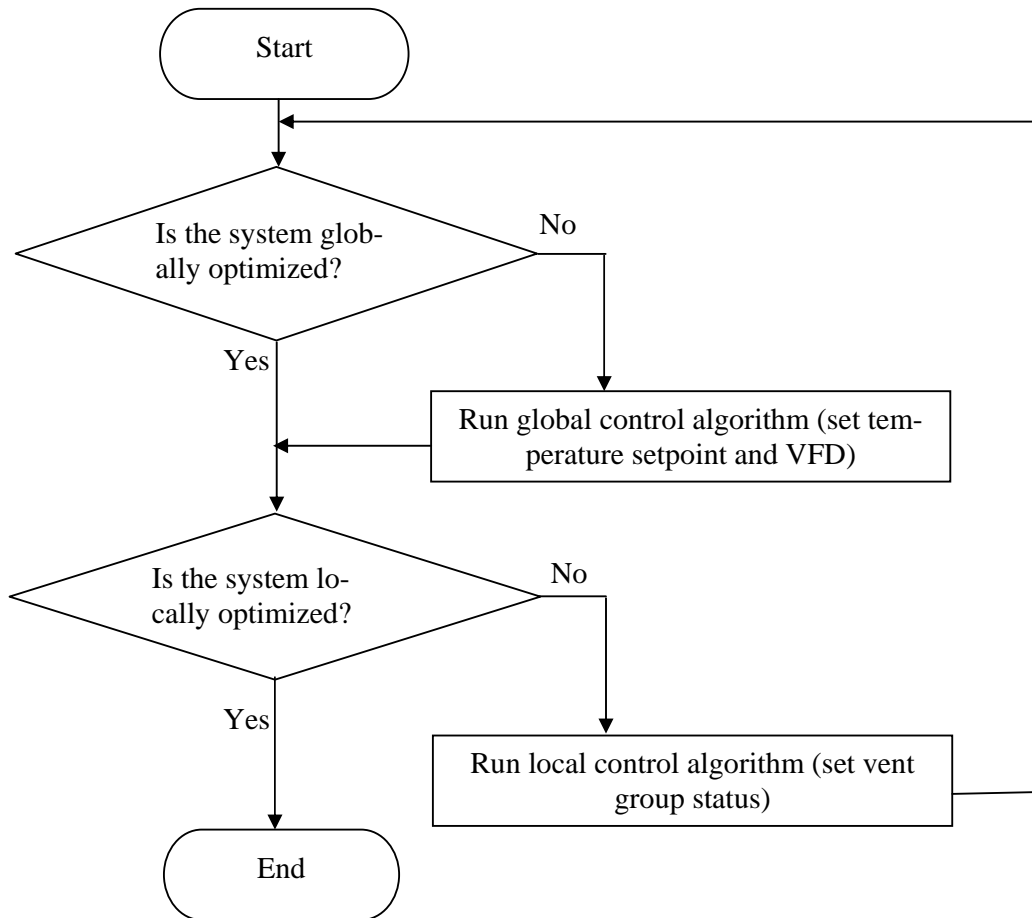


Figure 11: Flow Chart of Cooling Control System

We should pay attention to the following issues:

1. Figure 11 only indicates a complete cycle of system control. It ends with the status that both global and local optimization have been achieved.
2. One cycle doesn't equal ultimate optimization. Due to the characteristic of the data center (i.e. dynamic heat load) or other factors (i.e. ongoing construction) the optimal operation point can be different. The system should have the capability to check the status of system optimization from time to time.
3. In Figure 11, global control is "rough" control, and local control is "fine" control. Only one algorithm (either the local algorithm or the global algorithm) can be running during any given period so that these algorithms won't interfere with each other.
4. During any iteration, it doesn't matter which algorithm runs first. Figure 11 starts with the global control algorithm with the following consideration: if the system experiences an emergency, i.e. temperatures of some locations increase quickly, it is more reasonable to adjust the temperature setpoint and VFD of the CRAC units to deal with this situation. The capacity of local control is limited, depending on the interaction of temperature distribution and vent group configurations.

CONCLUSION

Two control algorithms addressing local rack inlet temperature control are proposed and implemented experimentally in the HP Data Center.

- The algorithm based on the correlation matrix between vent group configuration and inlet temperature requires preliminary experiments.
- Since the correlation matrix highly depends on the boundary conditions of flow domain, it requires recalibration whenever any major adjustment takes place in the data center.
- Further experiments are required to test the correlation matrix under different base cases.
- The algorithm based on vent group family doesn't require any preliminary experiments.
- The algorithm based on vent group family is effective and robust.
- The algorithm based on vent group family can be optimized by learning and adjusting the vent group families during the control process.

In general, both results have indicated the possibility of using vent tiles as additional actuators for local temperature control and system optimization. A further step is to test the effect of the local temperature control algorithms within the whole cooling scheme. Within a big picture of data center cooling, we can learn how much energy savings the local control algorithms can achieve.

ACKNOWLEDGMENTS

Many thanks to Ratnesh Sharma of Hewlett-Packard Laboratories who helped a lot with the temperature network and data logging. Thanks to Kathe Gust for reviewing the report.

REFERENCES

- 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"
- Chen, K., Auslander, D. M., Bash, C. E., Patel, C. D., 2005, "Local Temperature Control in Data Center Cooling: Part I, Correlation Matrix"
- Chen, K., Auslander, D. M., Bash, C. E., Patel, C. D., 2005, "Local Temperature Control in Data Center Cooling: Part II, Statistical Analysis"
- Friedrich, R., Patel, C. D., 2002, "Towards Planetary Scale Computing – Technical Challenges for Next Generation Internet Computing", THERMES 2002, Santa Fe, New Mexico
- 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
- 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"