Proceedings of the ASME 2011 International Mechanical Engineering Congress & Exposition
IMECE2011
November 11-17, 2011, Denver, Colorado, USA

IMECE2011-62506

MODELING AND CONTROL FOR COOLING MANAGEMENT OF DATA CENTERS WITH HOT AISLE CONTAINMENT

Rongliang Zhou, Zhikui Wang, Cullen E. Bash, Alan McReynolds
Sustainable Ecosystems Research Group, HP Labs,
Hewlett-Packard Company, 1501 Page Mill Road, Palo Alto, CA 94304-1126.
Email: {firstname.lastname}@hp.com

ABSTRACT

In traditional raised-floor data center design with hot aisle and cold aisle separation, the cooling efficiency suffers from recirculation resulting from the mixing of cool air from the Computer Room Air Conditioning (CRAC) units and the hot exhaust air exiting from the back of the server racks. To minimize recirculation and hence increase cooling efficiency, hot aisle containment has been employed in an increasing number of data centers. Based on the underlying heat transfer principles, we present in this paper a dynamic model for cooling management in both open and contained environment, and propose decentralized model predictive controllers (MPC) for control of the CRAC units. One approach to partition a data center into overlapping CRAC zones of influence is discussed. Within each zone, the CRAC unit blower speed and supply air temperature are adjusted by a MPC controller to regulate the rack inlet temperatures, while minimizing the cooling power consumption. The proposed decentralized cooling control approach is validated in a production data center with hot aisles contained by plastic strips. Experimental results demonstrate both its stability and ability to reject various disturbances.

1 INTRODUCTION

Due to the ever-increasing power density of the IT equipment, today's data centers consume tremendous amount of power. According to [1, 2], about a third to a half of data center total power consumption goes to the cooling system, and hence highly efficient cooling systems are indispensable to reduce the total cost of ownership and environmental footprint of data centers.

Figure 1 shows a typical raised-floor air-cooled data center with hot aisles and cold aisles separated by rows of IT equipment racks. The thermal requirements of IT equipment are usually specified in terms of the inlet air temperatures of the equipment [3]. The blowers of the Computer Room Air Conditioner (CRAC) units pressurize the under-floor plenum with cool air, which in turn is drawn through the vent tiles located in front of the racks in the cold aisles. Hot air carrying the waste heat from the IT equipment is rejected into the hot aisles. Depending on its design, the CRAC unit internal control can regulate the chilled water valve opening to track the given reference of Supply Air Temperature (SAT) or Return Air Temperature(RAT). The flow rate of the cool air supply can also be tuned continuously if a Variable Frequency Drive (VFD) is installed for each CRAC unit to vary the speed of its blowers.

In an open environment without aisle containment, air streams are free to mix. Most of the hot air in the hot aisles returns to the CRAC units, but a small portion of it might escape into the cold aisles from the top, the sides, or even the bottom of the racks and causes recirculation. The inlet air flow of the IT equipment is thus a mixture of cool air from the vent tiles in its vicinity and the recirculated hot air [4]. Recirculation can be also due to the reverse flows with certain IT equipment (some network switches, for example) of which the internal fans blow the hot exhaust air from the hot aisle into the cold aisle.

The recirculation of hot air into the cold aisle generates entropy and reduces the data center cooling efficiency [4]. In order to reduce the mixing of hot and cold air streams and hence improve the cooling efficiency of the original cold and hot aisle sep-

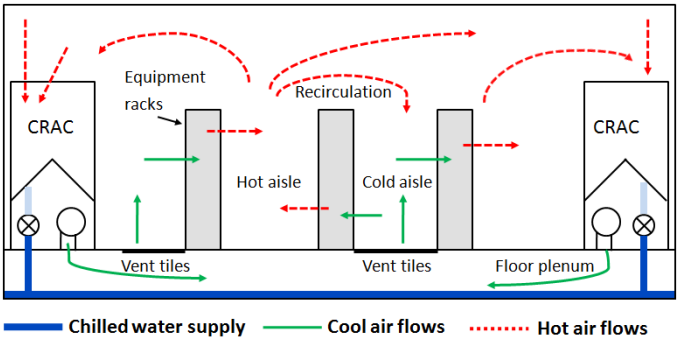

FIGURE 1. TYPICAL RAISED FLOOR DATA CENTER

aration scheme developed in [5], various modifications have been proposed. Most of the improvement efforts are centered around building a mass and heat transfer boundary between the hot and cold air streams and eliminating recirculation. Variations of the alternating cold and hot aisles configuration such as cold aisle containment, in row cooling with hot aisle containment, overhead cooling, and ducted hot air return path are studied through simulation in [6]. In an experimental study, 40% CRAC units blower power savings are achieved by establishing a ducted hot air return path between the IT equipment and the CRAC units in a production data center owned by Oracle Corporation [7]. To the authors' knowledge, hot aisle containment has been adopted in some data centers, but its modeling, simulation, and experimental explorations are still scarce in the literature.

While improvements in the cooling infrastructure can help enhance the cooling efficiency, its effectiveness is limited if the CRAC units are poorly controlled. Cooling control based on the return air temperature of the CRAC units, for example, could result in gross overprovisioning of the cooling resources. In order to maximize the cooling power savings, real-time thermal status monitoring of the IT equipment and cooling actuation with fine time and space granularities are essential. Real-time sensing capability, such as the extensive temperature sensor network introduced in [8], ensures timely response to thermal anomalies of the IT equipment. Coordinated tuning of the CRAC units blower speeds and SAT, on the other hand, seeks to minimize the power consumption of both CRAC units and chiller plants. The authors' recent work [9], for example, demonstrates that significant cooling power savings can be achieved through the coordinated zonal (CRAC units blower speeds and SAT) and local cooling actuation (adaptive vent tiles). The holistic modeling and control approach presented in [9] is validated on a small portion of a production data center, leaving cooling control system design for the entire data center unexplored.

In this paper, we develop a dynamic model for data center cooling management and design decentralized model predictive controllers (MPC) for the multiple CRAC units within the data center. The model characterizes the dynamics of the rack inlet temperatures as the results of both the mixing of the cool air flows from the CRAC units and the mixing of the cool air and hot air. The hot aisle or cold aisle contained environment are special cases of data centers described by the model. Based on the influences of CRAC units on the rack inlet temperatures, a data center is partitioned into different zones, each of which contains certain number of CRAC units and IT equipment. The CRAC units of each zone can significantly affect the inlet temperatures of the IT equipment within that zone, but not outside the zone. The MPC controller of each partitioned zone coordinates the CRAC units blower speed and supply air temperature (SAT) to maintain the rack inlet temperatures within the partition below the specified temperature thresholds, while minimizing the cooling power consumption. The decentralized control system structure lowers the risk of controller failure and is scalable to large scale data centers.

The other sections of this paper are organized as follows. Section 2 derives the dynamic cooling models using the energy and mass balance principles. In Section 3, we present the decentralized control system structure and controller design. Section 4 focuses on the controller implementation details and experimental results. The paper is concluded in Section 5 with discussion on the future work.

2 DYNAMIC COOLING SYSTEM MODELING

In this section, we derive simplified models from the basic mass and energy balance principles to characterize the complex mass and energy flows within the raised-floor air-cooled data centers.

2.1 Air Streams Mixing at Rack Inlet

In the open environment, air flow coming into the IT equipment inlet is a mixture of the cool air from the CRAC units (through the vent tiles) and the recirculated hot (exhaust) air that escapes into the cold aisle. In hot aisle contained environment, although significantly reduced, recirculation could still exist because of imperfect containment, as happened in the data center for our experiments, and the reverse flows from some network switches mentioned earlier. The internal fans of these network switches draw hot air from the hot aisle for cooling, reject the even hotter air (up to $40°C$) into the cold aisle, and cause the inlet temperatures of neighboring IT equipment to rise. For generality, we choose to include the effects of recirculation in the modeling, and we can easily set the corresponding item to zero for perfectly contained environment without reverse flows.

To determine the effects of both cool air and recirculated hot air on the rack inlet temperature, consider a small control volume in the proximity of the rack inlet with mass m and temperature T, as shown in Fig. 2. Cool and recirculated hot air flows with mass

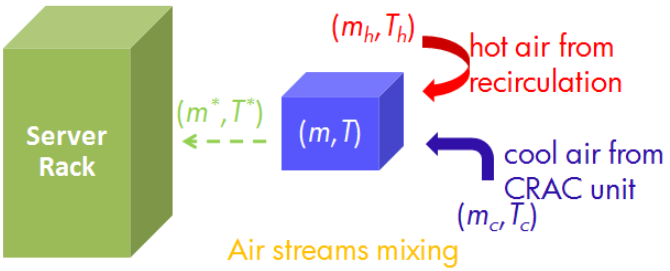

FIGURE 2. AIR MIXING AT THE RACK INLET

and temperature (m_c, T_c) and (m_h, T_h) enter the control volume, mix well with the air (m, T) already in the volume, leave the control volume altogether and enter the rack inlet with total mass m^* and temperature T^*. Based on mass balance principle,

$$m^* = m + m_c + m_h, \tag{1}$$

and from energy balance principle,

$$m^* h^* = mh + m_c h_c + m_h h_h, \tag{2}$$

in which h stands for enthalpy. Within the typical data center operation temperature range, air can be approximated as an ideal gas with $dh = c_p dT$, and the constant-pressure specific heat capacity c_p can be assumed constant.

Combining Eqn.(1) and (2), it can be found that the temperature change ΔT of the air within the control volume before and after the mixing is:

$$\Delta T \triangleq T^* - T = \frac{m_c(T_c - T)}{m + m_c + m_h} + \frac{m_h(T_h - T)}{m + m_c + m_h}. \tag{3}$$

Equation (3) reveals that the influence of cool and recirculated hot air on rack inlet temperature can be mainly captured by $m_c(T_c - T)$ and $m_h(T_h - T)$, respectively. This seemingly very simple insight is consistent with the physical intuition and also provides guidance to unite the CRAC unit SAT and VFD control as we will see later.

2.2 Cool Air Distribution From CRAC Units to Rack Inlets

While the temperature of the recirculated hot air is beyond direct control, the cool air delivered to the rack inlets can be adjusted through tuning of SAT and VFD of the CRAC units.

In raised-floor data centers, the pressure difference below and above the floor drives the cool air flow toward above the floor through the vent tiles. Assuming that the air density change is negligible during the normal operation of CRAC units, the total cool air flow $\dot{m}_{CRAC}$ delivered by the blower of each CRAC unit can be determined by the fan law [10]:

$$\dot{m}_{CRAC} = K \cdot VFD, \tag{4}$$

in which VFD stands for the speed of the blower in the percentage of its maximum. The coefficient K may vary with each CRAC unit and can be either provided by the manufacturer or determined through experiments.

The cool air flow, after leaving the CRAC unit blowers and traveling through the under-floor plenum, is distributed through the vent tiles. For multiple CRAC units deployment, the cool air flowing into a rack inlet could come from all the CRAC units:

$$\dot{m}_c = \sum_{j=1}^{N_{CRAC}} b_j \cdot \dot{m}_{CRAC,j} = \sum_{j=1}^{N_{CRAC}} b_j \cdot K_j \cdot VFD_j, \tag{5}$$

in which N_{CRAC} is the number of CRAC units and b_j quantifies the cooling air contribution from the j^{th} CRAC unit to a specific rack inlet. The values of b_j can vary depending on the amount of cool air provided by the CRAC unit to the rack inlet.

2.3 Dynamic Rack Inlet Temperature Model

In this section, we incorporate the effects of SAT tuning on the rack inlet temperatures and derive the discrete dynamic rack inlet temperature model.

It can be seen from Eqn.(3) that both cool and recirculated hot air contribute to the rack inlet temperature change ΔT. Since the recirculated hot air flow is beyond direct control, we can lump its effect into a time-varying term C and simplify Eqn.(3) as:

$$T^* - T = \frac{\dot{m}_c \Delta t (T_c - T)}{m + \dot{m}_c \Delta t + m_h} + C, \tag{6}$$

in which Δt is the length of the sampling interval. Substitute Eqn.(5) into Eqn.(6) and replace T^* and T with rack inlet temperatures at time steps $k+1$ and k, respectively, we have after further simplification:

$$\begin{aligned} T(k+1) =& T(k) + \{ \sum_{j=1}^{N_{CRAC}} g_j \cdot [SAT_j(k) - T(k)] \cdot VFD_j(k) \} \\ &+ C(k), \end{aligned} \tag{7}$$

in which g_j quantifies the combined influences of VFD and SAT tuning of the j^{th} CRAC unit, and also lumps the effects of parameters $b_j, K_j, \Delta t$ together with the nonlinearity associated with $\dot{m}_c$ in Eqn.(6).

The vector form of Eqn.(7) for multiple rack inlet temperatures is:

$$\overline{T}(k+1) = \overline{T}(k) + \overline{F} + \overline{C}, \tag{8}$$

in which

$$\begin{aligned}
&\overline{T} = [T_1, T_2, \cdots, T_{N_T}]^T, \\
&\overline{F} = [F_1, F_2, \cdots, F_{N_T}]^T, \\
&F_i = \sum_{j=1}^{N_{CRAC}} g_{i,j}[SAT_j(k) - T_i(k)]VFD_j(k), \qquad 1 \le i \le N_T, \\
&\overline{C} = [C_1, C_2, \cdots, C_{N_T}]^T,
\end{aligned}$$

and N_T is the number of rack inlet temperatures of interest.

3 DECENTRALIZED CONTROLLER DESIGN

Data center cooling management can be formulated as an optimal control problem, in which the total cooling power is minimized in response to the dynamic IT workload while the rack inlet temperatures are maintained at or below the specified thresholds. The temperature thresholds are not necessarily uniform across the entire data center but are dependent on the different functions, such as computing, storage, and networking, that the IT equipment serves. Service contracts of the IT workload hosted in the IT equipment also affect the temperature thresholds.

3.1 Control System Structure

The multiple CRAC units within a data center, together with hundreds or even thousands of rack inlet temperatures to be regulated, form a complex large scale system. Cooling control of this system using a centralized controller is computationally intensive, and the robustness and reliability associated are also big concerns.

With the observation that each CRAC unit within a data center has its own zone of influence, and significantly affects only the inlet temperatures of the nearby racks, decentralized control can be utilized to replace the commonly used centralized controller design. After associating each CRAC unit with the rack inlet temperatures on which it has significant effects, the entire data center can be partitioned into a number of zones. Each zone contains one CRAC unit and a number of rack inlet temperatures that the CRAC unit has control over. In traditional decentralized control for large scale industrial systems, the subsystems do not share inputs or outputs [11]. Inlet temperature of a specific rack within a data center, however, could be significantly affected by more than one CRAC units. Naturally, the different zones of a data center after partition can have overlapping outputs of rack inlet temperatures. The inputs of the various zones, which are the CRAC unit VFD and SAT, are not shared among different zones since each zone has only one CRAC unit and in the current paper we assume the zones do not communicate with each other. In the decentralized control scheme, the simplified model for a rack inlet temperature T in the j^{th} zone is:

$$T(k+1) = T(k) + g_j \cdot [SAT_j(k) - T(k)] \cdot VFD_j(k) + C(k). \tag{9}$$

Compared with Eqn.(7), the influences of CRAC units other than the one inside zone j are captured by $C(k)$ as well.

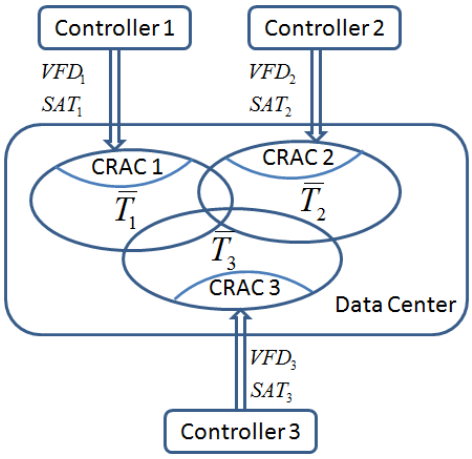

FIGURE 3. DECENTRALIZED CONTROL SYSTEM STRUCTURE

Figure 3 illustrates the proposed decentralized control system structure for a data center with three CRAC units, in which $\overline{T}_i (i = 1,2,3)$ are vectors containing rack inlet temperatures of the i^{th} zone. Note that each decentralized controller regulates the blower speed VFD and supply air temperature SAT of one zone, and neighboring zones could have overlapping rack inlet temperatures. Note also that, by sharing rack inlet temperatures between neighboring zones, the decentralized cooling system has better robustness compared with a non-overlapping partition. In the event of controller or CRAC failure in one zone, the temperature rise of that specific zone will be detected by the adjacent zones with rack inlet temperature shared, and these neighboring zones will respond with increased cooling resources provisioning

to maintain the shared rack inlet temperatures below the specified thresholds.

3.2 MPC Controller Design within Each Zone

Figure 4 shows the MPC controller structure within each zone of the partitioned data center. The two cooling knobs available to the controller are the CRAC unit SAT and VFD. The effects of these cooling actuators on the rack inlet temperatures are captured by the models in the previous section. The objective function of the MPC controller is set up to reflect the total power usage of the cooling system. By comparing the rack inlet temperature measurements $\overline{T}$ with the temperature threshold $\overline{T}_{ref}$, the MPC controller automatically seeks the optimal CRAC unit settings in response to the dynamic IT workload.

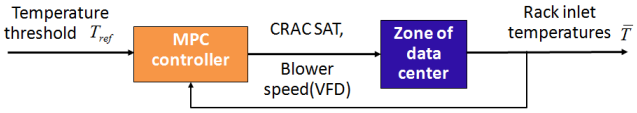

FIGURE 4. MPC CONTROLLER WITHIN EACH ZONE PARTITION

The MPC controller is formulated as a constrained minimization problem:

$$J(VFD,SAT) = \sum_{i=0}^{hu-1} \{\{VFD^3(k+i)\}_{R_{VFD}} + \{-SAT(k+i)\}_{R_{SAT}} + |\Delta VFD(k+i)|_{W_{VFD}} + |\Delta SAT(k+i)|_{W_{SAT}}\}$$

subject to:

$$\begin{aligned} VFD_{min} &\leq VFD(k+i) \leq VFD_{max}, \\ SAT_{min} &\leq SAT(k+i) \leq SAT_{max}, \\ \overline{T}(k+j+1) &\leq \overline{T}_{ref}, \end{aligned}$$

for all $0 \leq i \leq hu-1$ and $0 \leq j \leq hp-1$.

In the constrained optimization above, hu and hp are the control horizon and prediction horizon respectively, with $hu \leq hp$. Increment of control actions $\Delta VFD(k+i)$ and $\Delta SAT(k+i)$ are defined as:

$$\begin{aligned} \Delta VFD(k+i) &= VFD(k+i) - VFD(k+i-1), \\ \Delta SAT(k+i) &= SAT(k+i) - SAT(k+i-1). \end{aligned}$$

Beyond the control horizon, VFD and SAT remain constant from time step $hu-1$ to time step $hp-1$.

The objective function J penalizes the total cooling power consumption, and the rate of change of cooling actuation is also penalized for the purpose of system stability. The CRAC unit blower power increases along with VFD^3 according to the fan laws, and it is also assumed that the chiller power consumption increases linearly as the CRAC unit SAT decreases. R_{VFD} and R_{SAT} are appropriate weights on the blower power of the CRAC units and the thermodynamic work of the chiller plant.

Among the optimization constraints, $\overline{T}_{ref}$ is the rack inlet temperature threshold. Cooling control inputs including the blower speed VFD and supply air temperature SAT are constrained by their respective operational limitations. It is found through experiments, for example, that in most cases it is not desirable to turn a CRAC unit off even if its load is very low since doing so will significantly change the air flows within the data center while resulting in insignificant power savings.

In order to obtain smoother cooling control action, a deadband can be applied to the maximum temperature violation $T_{vio\_max} = max(\overline{T} - \overline{T}_{ref})$ within the zone, such that the control action does not change as long as T_{vio_max} stays within the deadband. The size of the deadband, however, needs to be chosen carefully for each individual zone without significantly compromising the system transient performances.

4 EXPERIMENTAL SETUP AND RESULTS

The proposed decentralized controller is implemented and evaluated through experiments in a production data center. The hot aisles of this data center are contained by plastic strips hanging from the ceiling. We present part of the experimental results in this section.

4.1 TestBed

Figure 5 shows the experimental data center with 10 rows of racks and 8 CRAC units. All of the racks hosted are fully instrumented with 5 temperature sensors in the front and another 5 in the back at different heights. The red dash-dot lines delineate the hot aisles isolated from the rest of the data center by walls, ceiling, and plastic strips hanging from the ceiling. Row F and G are separated from row Aext, Bext, and Cext by a wall above the floor and dampers along the wall in the underfloor plenum. Due to the ongoing IT equipment upgrade, some rows are not filled to their full capacities and temperature sensors for some racks might not function properly. In this paper, we only consider the total of 220 rack inlet temperature sensors. The five inlet temperature sensors for each rack are named $T1$ through $T5$ from the rack bottom up to the top. Temperature G9.T5, for example, refers to the inlet temperature at the top of rack G9.

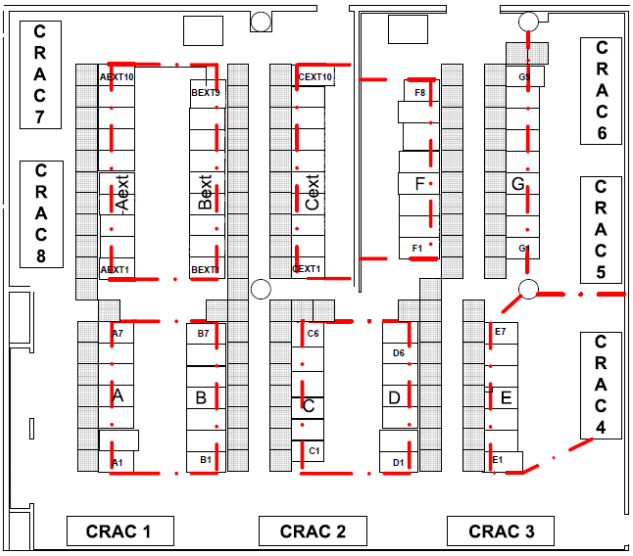

FIGURE 5. LAYOUT OF THE EXPERIMENTAL DATA CENTER

4.2 Zone Partition for Decentralized Control

For decentralized cooling control design, the entire data center needs to be partitioned into a number of zones. Each zone contains both cooling actuator (CRAC units) and a number of rack inlet temperatures to be maintained at or below the specified thresholds. Targeting a completely decentralized design, we have chosen to allocate only one CRAC unit to each zone. The experimental data center is thus divided into 8 zones, with the i^{th} CRAC unit in the i^{th} zone.

In order to limit the interactions between individual zones and enhance stability, it is essential to group the inlet temperatures with the CRAC units that most effectively influence them. The influence of a CRAC unit on a specific rack inlet temperature can be captured by the thermal correlation index (TCI) value. As defined in Eqn.10, $TCI_{i,j}$ quantifies the response of the i^{th} rack inlet temperature to a step change in the SAT of the j^{th} CRAC unit [8].

$$TCI_{i,j} = \frac{\Delta T_i}{\Delta SAT_{CRAC,j}}. \tag{10}$$

All the TCI values $[TCI_{i,j}](1 \le i \le 220, 1 \le j \le 8)$ constitute a static gain array from the system inputs to the outputs.

For rack inlet temperature T_i, define

$$TCI_{i,max} = max(TCI_{i,j}), \qquad 1 \le j \le 8,$$

then T_i is assigned to the j^{th} zone containing CRAC unit j if

$$TCI_{i,j} \ge \lambda \cdot TCI_{i,max},$$

in which λ is a threshold to adjust the extent of overlapping of rack inlet temperatures between individual zones. Choosing $\lambda = 1$, for example, results in disjoint partitions of all the rack inlet temperatures among the 8 zones.

Table 1 lists the number of rack inlet temperatures of each zone when choosing $\lambda = 0.5$ for the partition. Because of the overlapping between neighboring zones, the number of rack inlet temperature sensors of the 8 zones add up to 338, more than the aforementioned total of 220 rack inlet temperature senors deployed in the data center. The decentralized control experimental results to be presented later are based on this partition.

TABLE 1. NUMBER OF RACK INLET TEMPERATURE SENSORS FOR EACH ZONE IN A PARTITION WITH $\lambda = 0.5$

Zone	1	2	3	4	5	6	7	8
Number of inlet temperature sensors	8	19	26	69	72	59	35	50

4.3 Implementation of Decentralized Controller

Following the model structure of Eqn.(8), a multiple-input-multiple-output (MIMO) model is obtained through system identification experiments for each of the 8 partitioned zones. The inputs of each MIMO model are VFD and SAT of the CRAC unit in the zone, and the outputs are the rack inlet temperatures within the zone.

For initial development, the decentralized MPC controllers are implemented in MATLAB running on a Windows server. Each of the 8 MPC controllers monitors the thermal status within its zone and adjusts the CRAC unit VFD and SAT accordingly. Although currently running on the same machine, the decentralized MPC controllers can be readily implemented on different servers. The Matlab optimization toolbox [12] function "fmincon" is used to solve the constrained optimization problem. In the objective function J, R_{VFD} and R_{SAT} are chosen to reflect the actual blower power and chiller power consumption, and the weights W_{VFD} and W_{SAT} are chosen carefully to ensure satisfactory transient performance. The control interval is set to 25 seconds to allow for sufficient time for computation. The control horizon hu and prediction hp are 1 and 5, respectively.

Constraint relaxation method [13] is applied to address the feasibility problem due to the temperature constraints. In the case of temperature threshold violations, the temperature constraints in the prediction horizon are relaxed to approach the thresholds

asymptotically. It is observed in the experiments that the temperature constraints relaxation help avoid aggressive control actuation as well.

4.4 Experimental Results–Step Change of Rack Inlet Reference Temperatures

Extensive experiments were done to validate the stability of the MPC controllers. In this and the next sections, we show two examples where the closed-loop system was disturbed by step changes of the reference temperature values or by step changes of the tile openings.

Figure 6 shows the controllers outputs, the SAT and VFD trajectories of the CRAC units, when the rack inlet temperature reference values $\overline{T}_{ref}$ of row F were reduced by $2^{\circ}C$, and then changed back to what they were after two hours. The outputs of the MPC controllers in zones 4, 5 and 6 are included because they cover 9, 36 and 25 temperature sensors in row F, respectively. No temperatures in row F belonged to zones 1, 2, 3, 7, and 8, and the related CRAC unit settings were not affected significantly by the T_{ref} changes in row F. The output trajectories of CRAC unit 3 are included in the figures as an example.

The data center had reached a steady state before the step changes were applied. The maximum rack inlet temperature violation $T_{vio\_max} = max(\overline{T} - \overline{T}_{ref})$ for those in row F was -0.9$^{\circ}C$, occurring at rack inlet temperature F8.T1 (not shown in the figures). Note that the negative violation value implied that all the temperatures were below their references (or thresholds). The first step change in T_{ref} 20 minutes after the experiment started instantly increased T_{vio_max} from -0.9$^{\circ}C$ to 1.1$^{\circ}C$ for both zones 5 and 6 since the temperature F8.T1 became the dominant sensor, e.g., the one with the largest violation, in both zones. The MPC controllers in these two zones responded to the T_{vio_max} excursion by reducing the SATs and increasing the blower speeds to increase the cooling resource provisioning as shown in Fig.6(a) and 6(b), and converged to new steady states after a few steps. The temperature F8.T1 was not part of zone 4, however. Upon the step change of the reference, the T_{vio_max} (of those temperatures in Row F but belong to Zone 4) were increased from -0.3$^{\circ}C$ to 0.2$^{\circ}C$ only since these temperatures were at least 1.8$^{\circ}C$ below their thresholds before the reduction of T_{ref}. The minor temperature violation of zone 4 led to slight decrease of its SAT and increase of VFD.

Similar explanations apply when two hours after the initial step change the T_{ref} values for row F were changed back to what they were, which caused the SATs of CRAC units 5 and 6 to increase and the VFD to decrease.

4.5 Experimental Results–Step Change of Vent Tile Opening

In order to evaluate the ability of the decentralized controller to reject external disturbances, the opening of the vent tiles in front of the racks were manually changed. In the experiment to be discussed in this section, the opening of the vent tile in front of rack G9 was tuned. This rack is located at the end of the row, where serious recirculation usually happens when some exhaust air from the hot aisle can escape the plastic strips containment. Changes of the vent tile opening affected the recirculation and hence the inlet temperatures significantly.

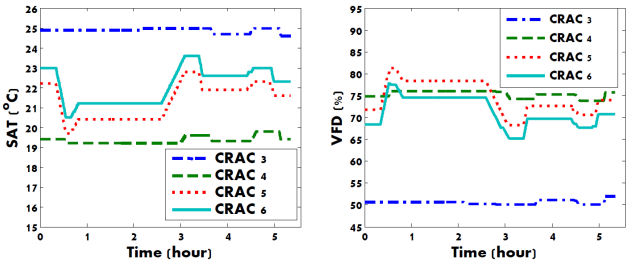

(a) SAT (b) VFD

FIGURE 6. RESPONSE TO RACK INLET TEMPERATURE REFERENCES STEP CHANGE – CRAC UNITS SAT AND VFD

Figure 7 and Figure 8 show one set of experimental results when the opening of the vent tile in front of G9 was increased from 50% to 100% at the time of one hour. With the larger tile opening, more cool air was directed to rack G9, which suppressed the recirculation and lowered the rack inlet temperature. Figure 8(a) shows that rack inlet temperature G9.T5, which was subject to more recirculation than the other inlet temperatures in the rack, drops more than 3$^{\circ}C$ because of the completely opened tile in front of the rack.

Prior to this step change in tile opening, the cooling system had reached a steady state as shown in Fig.7. The temperature G9.T5 had the maximum temperature violation of -0.2$^{\circ}C$ for both zone 5 and 6. The quick decrease in the temperature of G9.T5 drove T_{vio_max} for both zones 5 and 6 from -0.2$^{\circ}C$ to -1$^{\circ}C$. The controllers for the zones 5 and 6 responded with increased SAT and lowered VFD as shown in Fig.7 to save cooling power. An important fact worth mentioning is that as rack inlet temperature G9.T5 decreased to around 27$^{\circ}C$, 1$^{\circ}C$ below its reference temperature, it was no longer the rack inlet temperature with the highest temperature violation for zones 5 and 6. The rack inlet temperature G3.T3 took over the place of G9.T5 as the one with the maximum violation for both zones 5 and 6 during the transient process after the tile opening change. As the cooling system reached a new steady state, rack inlet temperature with the maximum violation for both zone 5 and 6 alternated between G3.T3 and F5.T5, as shown in Fig.8(b). The reference temperature for both G3.T3 and F5.T5 was 25$^{\circ}C$.

Since the temperatures of rack G9 belonged to only zones 5 and 6 but not others, the tile opening change in front of rack G9 didn't have significant influence beyond these two zones. As a result, the SAT and VFD of CRAC units 3 and 4 remained almost

constant during the experiment.

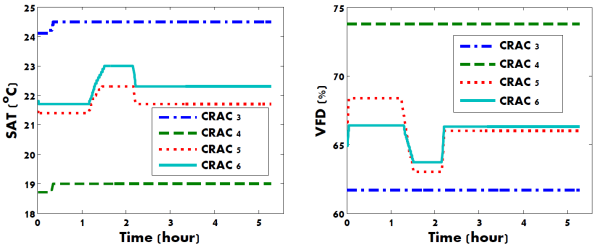

(a) SAT (b) VFD

FIGURE 7. RESPONSE TO TILE OPENING STEP CHANGE – CRAC UNITS SAT AND VFD

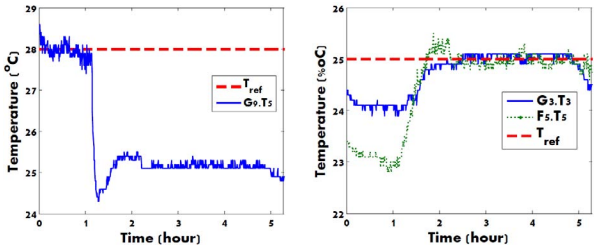

(a) BEFORE TILE STEP CHANGE: G9.T5 (b) AFTER TILE STEP CHANGE: G3.T3 AND F5.T5

FIGURE 8. RESPONSE TO TILE OPENING STEP CHANGE – INLET TEMPERATURE CORRESPONDING TO T_{vio_max}

5 CONCLUSIONS AND FUTURE WORK

In this paper, we derive a dynamic model for data center cooling management and propose a decentralized MPC controller design. The decentralized control system structure lowers the risk of failure suffered by centralized controllers, and can be more easily scaled up to large-scale data centers. The experimental results in a production data center with hot aisle containment validate the ability of the decentralized controller to handle reference change of rack inlet temperatures and to reject vent tile opening disturbance. To extend the work presented, the authors are working to compare the performance of the decentralized controller with that of the centralized controller, and to incorporate the tuning of adaptive vent tiles into the decentralized control framework.

REFERENCES

[1] Greenberg, S., Mills, E., Tschudi, B., Rumsey, P., and Myatt, B., 2006. "Best practices for data centers: Results from benchmarking 22 data centers". In Proceedings 2006 ACEEE Summer Study on Energy Efficiency in Buildings, ACEEE.

[2] Patel, C. D., Bash, C. E., Sharma, R.K., Beitelmam, M., and Friedrich, R.J., 2003. "Smart cooling of data centers". In Proceedings ASME 2003 International Electronic Packaging Technical Conference and Exhibition (InterPACK2003), ASME, pp. 129–137.

[3] ASHRAE, 2005. *Datacom Equipment Power Trends and Cooling Applications*. ASHRAE, Atlanta, GA.

[4] Bash, C.E., Patel, C.D., and Sharma, R.K., 2003. "Efficient thermal management of data centers – immediate and long-term research needs". *HVAC&R Research,* ***9***(2), pp. 137–152.

[5] Sullivan, R. F., 2000. "Alternating cold and hot aisles provides more reliable cooling for server farms". *Uptime Institute*.

[6] Wu, K., 2008. "A comparative study of various high density data center cooling technologies". MS Thesis, Stony Brook University, Stony Brook, NY.

[7] Martin, M., Khattar, M., and Germagian, M., 2007. "High-density heat containment". *ASHRAE Journal,* ***49***(12), pp. 38–43.

[8] Bash, C. E., Patel, C. D., and Sharma, R. K., 2006. "Dynamic thermal management of air cooled data centers". In Proceedings The Tenth Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronics Systems (ITHERM'06), IEEE, pp. 445–452.

[9] Zhou, R., Wang, Z., Bash, C. E., Hoover, C., Shih, R., McReynolds, A., Sharma, R.K., and Kumari, N., 2011. "A holistic and optimal approach for data center cooling management". In Proceedings 2011 American Control Conference (ACC 2011), IEEE, pp. 1346–1351.

[10] Munson, B.R., Young, D.F., and Okiishi, T.H., 2005. *Fundamentals of fluid mechanics*, 5th ed. John Wiley & Sons, Inc.

[11] Scattolini, R., 2009. "Architectures for distributed and hierarchical model predictive control-a review". *Journal of Process Control,* ***19***(5), pp. 723–731.

[12] Coleman, T. F., and Zhang, Y., 2003. *Optimization Toolbox For Use with MATLAB: User's Guide*, 2 ed. The MathWorks, Inc., Natick, MA.

[13] Oliveira, N. M. C., and Biegler, L. T., 1994. "Constraint handing and stability properties of model-predictive control". *AIChE journal,* ***40***(7), pp. 1138–1155.