2011 American Control Conference
on O'Farrell Street, San Francisco, CA, USA
June 29 - July 01, 2011

A Holistic and Optimal Approach for Data Center Cooling Management

Rongliang Zhou, Zhikui Wang, Cullen E. Bash, Alan McReynolds, Christopher
Hoover, Rocky Shih, Niru Kumari, and Ratnesh K. Sharma

Abstract— **Efficient and reliable operation of today's data centers, which host IT equipment with ever-increasing power density, relies heavily on the cooling system to meet the thermal management needs of the IT equipment with minimal environmental footprint. The dynamic IT workload, together with the spatial variance of cooling efficiencies, creates both temporal and spatial non-uniformities within the data centers. Most data centers use zonal cooling actuators, such as computer room air conditioners (CRAC), to alleviate the local "hot spots". Without proper localized cooling actuation mechanisms, the cooling capacity is usually over-provisioned that leads to waste of energy. To address this problem, we introduce adaptive vent tiles (AVT) for local cooling adjustment, and develop a holistic multivariable model based on the mass and energy balance principles to capture the effects of both zonal and local cooling actuation on the inlet temperatures of the racks that host the IT equipment. A model predictive controller is then proposed to minimize the total cooling power while meeting the thermal requirements of the racks. The zonal and local cooling actuation is coordinated in such a unified framework for the provisioning, transport and distribution of the cooling resources in the data centers. The proposed holistic cooling approach is validated in a production data center. Experimental results indicate that up to 36% of CRAC units blower power can be saved, compared with the state of the art control solution.**

I. INTRODUCTION

Due to the ever-increasing power density of the IT equipment, today's data centers consume tremendous amount of power. The Environmental Protection Agency (EPA), for example, reported that 60 billion kWh power was consumed by data centers in 2006, accounting for 1.5% of the total electricity usage of the United States [1]. According to [2], [3], about 1/3 to 1/2 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 (TCO) and environmental footprint of data centers.

Figure 1 shows a typical raised-floor air-cooled data center considered in this paper, where rows of IT equipment racks are separated by alternating hot and cold aisles. Air is the main media for heat transport in this open environment. The thermal requirements of IT equipment are usually specified in terms of the inlet air temperatures of the equipment [4]. 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. Neither the cold aisles nor the hot aisles are contained and hence air streams are free to mix. Most of the hot air in the hot aisles returns to the CRAC units, but a small portion might escape into the cold aisles from the top or the sides of 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 re-circulated hot air [5]. The cool air flow through any vent tile, similarly, could come from a number of CRAC units nearby. The complicated air flow distribution in the open environment makes cooling control of such data centers challenging. Moreover, cooling demands of the IT equipment are both temporally and spatially non-uniform due to at least two reasons. First, the IT workload is usually randomly distributed in the data center, and the workload itself is time-varying. Second, the cooling efficiency from the CRAC units to the IT equipment is heavily dependent on the physical layout of the IT and cooling infrastructure.

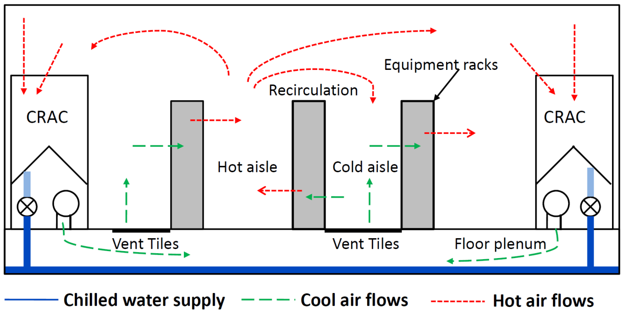

Fig. 1. Typical raised floor data center

In order to handle the complicated dynamics and the non-uniformity of cooling demands, 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 [6], ensures timely response to thermal anomalies of the IT equipment. Localized cooling actuators, on the other hand, help deliver cooling resources to any target area within the data centers. The CRAC units, used in most data centers today, have zonal cooling effects. A few number of server racks near a CRAC unit are significantly affected by the CRAC. Tailoring the cooling resources provisioning through CRAC units to meet the exact needs of the individual servers or racks is virtually impossible. With only the zonal controllers such as CRAC units, the cooling resources can be significantly over-provisioned throughout the data center. Deployment of adaptive vent tiles (AVT) [7] with adjustable openings in the cold aisles provides localized on demand cooling and the opportunities to improve the overall cooling

Sustainable Ecosystems Research Group, HP Labs, Hewlett-Packard Company, 1501 Page Mill Road, Palo Alto, CA 94304-1126. `{firstname.lastname}@hp.com`

978-1-4577-0079-8/11/$26.00 ©2011 AACC

efficiency. The challenge of cooling control brought by the introduction of AVTs, however, is how to orchestrate the zonal (CRAC units) and local (AVTs) actuation for stable and optimal cooling.

In most cases, the internal CRAC control can regulate the chilled water valve opening to track the given reference of Supplied Air Temperature (SAT). The flow rate of the supplied air can also be tuned continuously if a Variable Frequency Drive (VFD) is available for each CRAC unit to vary the speed of its blowers. In the previous decoupling based design [8], [9], one AVT controller adjusts the air distribution within each zone to meet the rack inlet temperature thresholds while minimizing the total air flow demand. The pressure in the underfloor plenum is maintained at certain reference value through CRAC units VFD control to satisfy the cooling needs of the zones. It is difficult to extend this design to include the tuning of SATs of the CRAC units, which have significant effects on both the rack inlet temperatures and the cooling efficiency. It is also challenging to handle the strong interactions between the two separate control loops. More importantly, the reference pressure, which determines the cooling resources provisioned and the blower power, is difficult to optimize.

In this paper, a holistic modeling, control and optimization framework is developed to integrate the provisioning, transport, and distribution of the cooling resources. The zonal and local cooling actuation is coordinated, in a unified framework with Model Predictive Controller (MPC), to minimize the total flow and thermodynamic work done by the cooling system. While MPC has been applied to thermal management in various contexts before [10], [11], the work to be presented is different in its scale since it is targeting the thermal regulation of hundreds of racks in large scale data centers with tens of CRAC units. Moreover, the commonly seen temperature tracking problem is now replaced by an energy minimization problem subject to temperature constraints.

The other sections of this paper are organized as follows. Section II derives the holistic models using the energy and mass balance principles. Based on the models developed, integrated cooling control and optimization using MPC is presented in Section III. Section IV discusses the controller implementation and experimental results. The paper is summarized in Section V with discussion on the future work.

II. HOLISTIC 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.

A. *Cool and Recirculated Hot Air 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. The recirculation of hot air in the cold aisle generates entropy and lowers the cooling efficiency of the data centers [5].

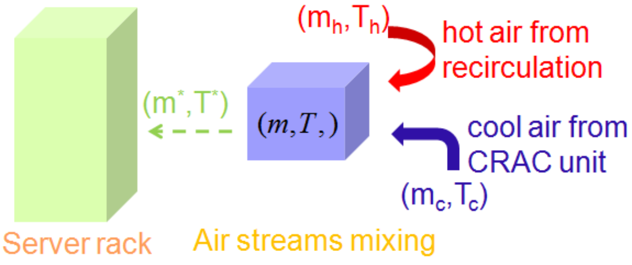

Fig. 2. Air mixing at the rack inlet

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 Figure 2. Cool and recirculated hot air flows with mass and temperature (m_c, T_c) and (m_r, T_r) 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,

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

and from energy balance principle,

$$m^* h^* = mh + m_c h_c + m_r h_r. \tag{2}$$

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 to be constant.

Combining equations (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:

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

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 $m_r(T_r - T)$, 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.

B. *Cool Air Flow From CRAC Unit 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, and vent tile openings to regulate the rack inlet temperatures.

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 for normal CRAC units operation, the total cool air flow $\dot{m}_{CRAC}$ delivered by the blower of each CRAC unit can be determined by the fan law:

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

in which VFD stands for the speed of the blower in the percentage of the maximum VFD setting. The coefficient k_{CRAC} 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. Each adaptive vent tile can be treated as an adjustable valve. In order to determine the total cool air flowing through a vent tile, we define the normalized tile opening $\underset{\sim}{U}_{tile}$ as:

$$\underset{\sim}{U}_{tile} = U_{tile} / \sum_{N_{tile}} U_{tile},$$

in which U_{tile} is the vent tile mechanical opening ranging from 0 to 100 percent and N_{tile} is the number of tiles in a cold aisle. For a cold aisle cooled exclusively by a single CRAC unit, if it is assumed that the cool air mass flow rate $\dot{m}_{tile}$ through each individual tile is proportional to its normalized opening $\underset{\sim}{U}_{tile}$, then we have:

$$\dot{m}_{tile} = \dot{m}_{CRAC} \cdot \underset{\sim}{U}_{tile}. \qquad (5)$$

It can be shown that the cool air flow distribution under this assumption is consistent with the mass balance principle, since

$$\dot{m}_{CRAC} = \sum_{N_{tile}} \dot{m}_{tile}. \qquad (6)$$

The cool air flows leaving the vent tiles are free to mix above the floor. As a result, the cool air flow $\dot{m}_c$ that reaches a rack inlet might come from several vent tiles in its vicinity:

$$\dot{m}_c = \sum_{N_{tile\_v}} b_{tile} \cdot \dot{m}_{tile}, \qquad (7)$$

in which N_{tile_v} stands for the number of vent tiles nearby that have significant influence over the cool air flowing into the rack, and the contribution of each vent tile is quantified by b_{tile}.

Equations (4), (5) and (7) together give

$$\dot{m}_c = k_{CRAC} \cdot VFD \sum_{N_{tile\_v}} b_{tile} \cdot \underset{\sim}{U}_{tile}, \qquad (8)$$

which describes how the rack inlet cool air flow is affected by one specific CRAC unit and the vent tiles near the rack. For multiple CRAC units deployment, we can sum up the total cool air flows into a rack inlet from all the CRAC units as following,

$$\dot{m}_c = \{ \sum_{N_{CRAC}} k_{CRAC} \cdot VFD \} \cdot \{ \sum_{N_{tile\_v}} b_{tile} \cdot \underset{\sim}{U}_{tile} \}, \quad (9)$$

in which N_{CRAC} is the number of CRAC units that significantly affect the cool air flow reaching the rack.

C. Control Oriented Rack Inlet Temperature Model

In this section, we extend the models to capture the effects of SAT tuning on the rack inlet temperatures.

Substitute equation (9) into equation (3) and replace T^* and T in equation (3) with rack inlet temperatures at time steps $k+1$ and k respectively, we have the following discrete model for rack inlet temperatures:

$$\begin{aligned} T(k+1) \quad = \quad & T(k) + \{ \sum_{i=1}^{N_{CRAC}} g_i \cdot [SAT_i(k) - T(k)] \cdot \\ & VFD_i(k) \} \cdot \{ \sum_{j=1}^{N_{tile}} b_j \cdot \underset{\sim}{U}_j(k) \} + C, \quad (10) \end{aligned}$$

in which g_i quantifies the influence of CRAC unit i and C represents the rack inlet temperature increase brought by recirculation. Notice that in equation (10), $b_j \neq 0$ $(1 \leq j \leq N_{tile})$ only when vent tile j is in the vicinity of the rack.

The model represented by equation (10) describes the dynamics of a single rack inlet temperature. The model parameters (g_i, b_j, C) with $1 \leq i \leq N_{CRAC}$ and $1 \leq j \leq N_{tile}$ need to be identified for each rack inlet temperature of interest. For multiple rack inlet temperatures, the vector form of equation (10) is:

$$\overline{T}(k+1) = \overline{T}(k) + \overline{F} \cdot B \cdot \overline{\underset{\sim}{U}}(k) + \overline{C}, \qquad (11)$$

in which

$$\begin{aligned} \overline{T} &= [T_1, T_2, \cdots, T_{N_{rack}}]^T, \\ \overline{F} &= diag(F_1, F_2, \cdots, F_{N_{rack}}), \\ F_i &= \sum_{j=1}^{N_{CRAC}} g_{ij} [SAT_j(k) - T_i(k)] VFD_j(k), \\ & 1 \leq i \leq N_{rack}, 1 \leq j \leq N_{CRAC} \\ B &= [b_{ij}], 1 \leq i \leq N_{rack}, 1 \leq j \leq N_{tile}, \\ \overline{\underset{\sim}{U}} &= [\underset{\sim}{U}_1, \underset{\sim}{U}_2, \cdots, \underset{\sim}{U}_{N_{tile}}]^T, \\ \overline{C} &= [C_1, C_2, \cdots, C_{N_{rack}}]^T, \end{aligned}$$

and "·" denotes matrix multiplication. The holistic and multivariable model above integrates the zonal cooling actuation of CRAC units and local cooling actuation of vent tiles, and lays the foundation for the control and optimization work to be discussed next.

III. INTEGRATED CONTROL AND OPTIMIZATION

Data center cooling is in essence 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 threshold.

A. Control System Structure

Figure 3 shows the proposed control system structure, in which the three cooling knobs available to the controller are the CRAC unit SAT, CRAC unit VFD, and vent tile openings. 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 zonal and local cooling settings in response to the dynamic IT workload. The cooling resources provisioning, transport, and distribution are coordinated since they are considered simultaneously in the same framework to minimize the cooling power.

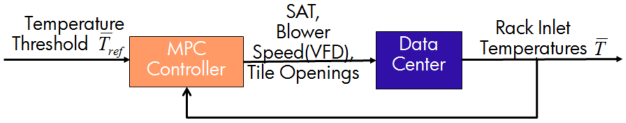

Fig. 3. Control system structure

B. MPC Problem Formulation

Following is the optimization problem in the vector form:

$$
\begin{aligned}
J(\overline{VFD},\overline{SAT},\overline{U}) = & \sum_{i=1}^{hu-1}\{\{\sum_{l=1}^{N_{CRAC}} VFD_l^3(k+i)\}_{R_{VFD}} \\
& +\{\sum_{l=1}^{N_{CRAC}}(SAT_{max}-SAT_l(k+i))\}_{R_{SAT}} \\
& +||\Delta\overline{VFD}(k+i)||^2_{W_{VFD}}+||\Delta\overline{SAT}(k+i)||^2_{W_{SAT}} \\
& +||\Delta\overline{U}(k+i)||^2_{W_{tile}}\}
\end{aligned}
$$

subject to:

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

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

In the constrained optimization above,

$$
\begin{aligned}
\overline{SAT} & = [SAT_1, SAT_2, \cdots, SAT_{N_{CRAC}}]^T, \\
\overline{U} & = [U_1, U_2, \cdots, U_{N_{tile}}]^T, \\
\overline{VFD} & = [VFD_1, VFD_2, \cdots, VFD_{N_{CRAC}}]^T, \\
\Delta\overline{VFD}(k+i) & = \overline{VFD}(k+i)-\overline{VFD}(k+i-1), \\
\Delta\overline{SAT}(k+i) & = \overline{SAT}(k+i)-\overline{SAT}(k+i-1), \\
\Delta\overline{U}(k+i) & = \overline{U}(k+i)-\overline{U}(k+i-1), \\
SAT_{max} & = max(\overline{SAT}_{max}),
\end{aligned}
$$

and $\overline{VFD}$, $\overline{SAT}$, and $\overline{U}$ remain constant from time step $hu-1$ to time step $hp-1$. The hu and hp are the control horizon and prediction horizon respectively, with $hu \le hp$.

The objective function J penalizes both the total cooling power consumption and the rate of change of cooling actuation. The CRAC units 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. The rate of change of control actions is penalized for the purpose of system stability.

Among the optimization constraints, $\overline{T}_{ref}$ is the rack inlet temperature threshold. Cooling control inputs including the blower speeds $\overline{VFD}$, supply air temperatures $\overline{SAT}$, and vent tile openings $\overline{U}$ are constrained by their respective physical or specification 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 negligible power savings.

IV. EXPERIMENTAL SETUP AND RESULTS

The proposed holistic modeling and control approach was implemented and evaluated through experiments in the experimental area of a production data center. We present part of the experimental results in this section.

A. TestBed

Figure 4 shows the data center with 10 rows of racks and 6 CRAC units. The experimental area is confined to the upper right section with AVTs fully populated in the cold aisle. It is isolated from the rest of the data center by walls, a removable curtain above the floor and dampers beneath the floor. The experimental area hosts two rows of IT equipment racks, row F and row G with 8 and 9 racks respectively. Each of the 17 racks is fully instrumented with 5 temperature sensors in the front and another 5 on the back, but only the 17 temperatures located at the top front of the racks are chosen to be regulated. The 20 AVTs, labeled VF and VG in the figure, line the two rows of the racks in the cold aisle. Two chilled water cooled CRAC units along the right wall, CRAC #5 and CRAC #6, cool the experimental area.

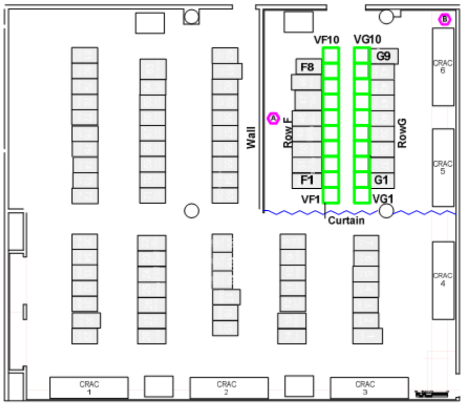

Fig. 4. Layout of the experimental data center

B. MPC Controller Implementation

Following the holistic model structure as in equation (11), a multiple-input-multiple-output (MIMO) model was identified through system identification experiments. The model takes the 17 temperatures as the outputs, and the SAT and VFD of CRAC #5 and CRAC #6 as well as the openings of the 20 AVTs as the inputs.

The MPC controller was implemented in Matlab running on a Linux server. The Matlab optimization toolbox [12] function “fmincon” was used to solve the constrained optimization problem. In the objective function J, R_{VFD} and R_{SAT} were chosen to reflect the actual blower power and chiller power consumption, and the weights W_{VFD}, W_{SAT}, and W_{tile} were chosen carefully to ensure satisfactory transient performance. The control interval was set to allow for sufficient time for computation.

Constraint relaxation method [13] was 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 were relaxed to approach the thresholds asymptotically. It was observed in the experiments that the temperature constraints relaxation helped avoid aggressive actuation as well.

C. Integrated Control of VFD and AVT

The holistic modeling and control approach was first applied to the integrated control of VFD and AVT. In this case, SAT was a fixed given value while CRAC units blower speeds and AVT openings were dynamically tuned by the MPC controller. During the experiment, the rack inlet temperature threshold was set to be 27 $°C$ uniformly, SAT was maintained at 17.8 $°C$ (by the internal PI controllers of the CRAC units), and the two CRAC units were configured to operate with the same blower speeds. The control horizon hu and prediction horizon hp were both set to be 3, and the control interval was 30 seconds.

Figure 5(a) shows the trajectories of all the 17 rack inlet temperatures during the five-hour experiment, with the temperature threshold denoted by the dotted straight black line. Despite of the varying IT workload that the experimental area might experience, all the rack inlet temperatures were kept below the threshold most of the time. The rare temperature violations, once appeared, were suppressed quickly. Figure 5(b) shows that the average opening of all the 20 vent tiles remained above 60% during the experiment, about 20% higher than that achieved by the approach presented in [9]. The much higher tile openings implied lower air flow resistance and hence less flow work. It was also observed that vent tiles at the ends of the two rows had noticeably higher average openings than those in the middle of the rows. This can be attributed to the more significant recirculation from the sides of the racks at these locations, and the consequent AVT adjustment to maintain the inlet temperature thresholds.

Figure 6 depicts the VFD setting and the total blower power of the two CRAC units. During the five-hour experiment, VFD setting varied by less than 10% after the initial transient period. It was also found that the average total power of the two blowers was 5.1kW over the entire experiment, 36% less than the approach in [9] where the CRAC units VFD and AVT were tuned by two separate control loops.

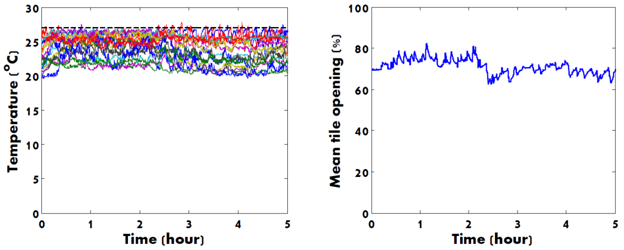

(a) Rack inlet temperatures (b) Average opening of all AVTs

Fig. 5. Rack inlet temperatures and average opening of all AVTs

The significant blower power savings come from the coordinated cooling resources provisioning and distribution. The zonal actuation (CRAC units) provides just sufficient amount of cool air in response to the dynamic IT workload, which is routed through the local actuation (AVTs) to the racks according to their individual cooling needs. In comparison, the cooling resources provisioned in [9] is highly dependent on the pressure reference that the CRAC units VFDs are tuned to maintain. The optimal pressure reference that guarantees adequate cooling resources provisioning is difficult to find, and a conservative high reference value unavoidably leads to higher blower speeds and hence larger blower power consumption.

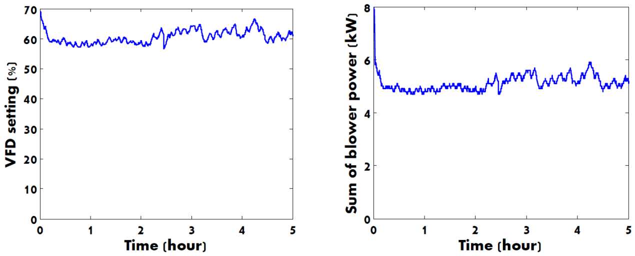

(a) VFD setting (b) Sum of blower power

Fig. 6. Speed and power of CRAC unit blowers

D. Integrated Control of SAT, VFD, and AVT

Power savings of the chiller plant from reduction of the thermodynamic work could be exploited when SAT regulation is integrated with the control of VFD and AVTs. Our preliminary investigation assumed that the chiller power consumption increases linearly with reduced chilled water supply temperature and hence reduced SAT of the CRAC units. In the objective function J of the MPC controller, experimental data in previous work of chiller operation optimization [14] were used to set up appropriate R_{SAT}.

The experiment for integrated control of SAT, VFD, and AVT lasted for 11 hours. The damper along the left wall was closed, while the curtain and dampers along the curtain were open, thus allowing air flow disturbance from the neighboring area both above and beneath the floor. The control horizon hu was set to 2, and the prediction horizon hp was set to 8 for a more conservative controller in consideration of the disturbance. The control interval was 30 seconds. The uniform rack inlet temperature threshold was set to $25°C$.

Figure 7(a) shows the trajectories of all the 17 rack inlet temperatures throughout the experiment, in which the temperature threshold is denoted by the dotted straight black line. The measured maximum rack inlet temperature remained below the specified temperature threshold, with occasional temperature violations by at most $0.4°C$. The average opening of all the 20 vent tiles was maintained steadily around 71% during the experiment as shown in Figure 7(b), and it was also found that the opening variation of most tiles was within 20%.

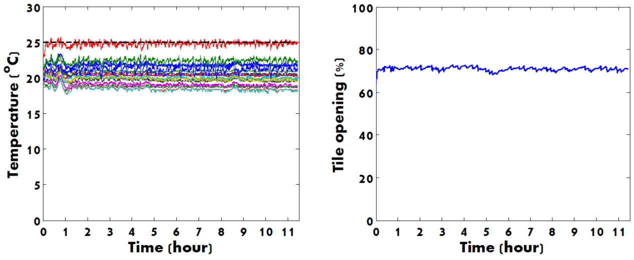

(a) Rack inlet temperatures (b) Average opening of all AVTs

Fig. 7. Rack inlet temperatures and average opening of all AVTs

Figure 8(a) and Figure 8(b) show the VFD and SAT set-points of the CRAC units, respectively. Despite the varying IT workload and air flow disturbances from the neighboring area, CRAC units VFD and SAT varied within a reasonable range ($68\% \sim 78\%$ and $16.9°C \sim 18.3°C$, respectively). During the experiment, the actual power consumption of the blowers in the two CRAC units was temporarily unavailable. It was found, however, that the nominal "total cooling power", defined by the objective function J of the MPC controller, did decrease from $13.8kW$ in the first hour when the controller was enabled to $13.2kW$ on average in the next 10 hours of the experiment, indicating the proper functioning of the optimal controller. The limited energy saving observed might be due to the fact that the experiment was started from a relatively energy efficient state with moderate VFD setting of 77.5% and high SAT of $18.3°C$. More control experiments are needed to determine if the optimal controller can drive the cooling system from an initial energy inefficient state to a highly energy efficient state.

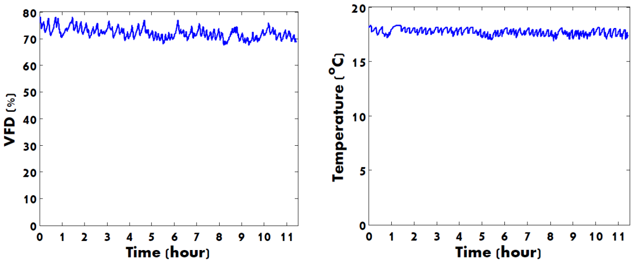

(a) VFD setting (b) SAT

Fig. 8. VFD and SAT

V. CONCLUSIONS AND FUTURE WORK

In this paper, a holistic data center cooling scheme is proposed to optimize the provisioning, transport, and distribution of cooling resources. Based on the integrated model developed, a model predictive controller is designed to coordinate the zonal and local cooling actuation that can minimize the cooling power consumption. It is found through experiments in a production data center that the proposed scheme can reduce the CRAC units blower power up to 36%. To extend the work presented, the authors are working on its scalability in large scale data centers and a decentralized cooling control framework with self-organizing and optimization capabilities.

VI. ACKNOWLEDGEMENTS

The authors would like to thank Dr. He Bai and Dr. Tiejun Zhang for their enlightening discussions on the design and implementation of the model predictive controller.

REFERENCES

[1] U.S. Environmental Protection Agency (EPA). *Report to congress on server and data center energy efficiency*. public law 109-43, August, 2007.

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

[3] Patel, C. D., Bash, C. E., Sharma, R., Beitelmam, M., and Friedrich, R.J. Smart cooling of data centers. In *IPACK03, The Pacific Rim/ASME International Electronic Packaging Technical Conference and Exhibitions*.

[4] ASHRAE. Datacom equipment power trends and cooling applications. Atlanta, GA, 2005.

[5] Cullen E. Bash, Chandrakant D. Patel, Ratnesh K. Sharma. *Efficient thermal management of data centers – Immediate and Long-term research Needs*, volume 9, no. 2. HVAC&R Research, Apr 2003.

[6] Bash, C. E., Patel, C. D., Sharma, R. Dynamic thermal management of air-cooled data centers. In *The 10th Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm '06)*, 2006.

[7] Beitelmal, M. H., Wang, Z., Felix, C., Bash C. E., Hoover, C. McReynolds, A. Local cooling control of data centers with adaptive vent tiles. In *the Proceedings of the ASME/Pacific Rim Technical Conference and Exhibition on Packaging and Integration of Electronic and Photonic Systems, MEMS and NEMS (InterPack 2009)*, 2009.

[8] Wang, Z., McReynolds, A., Felix, C., Bash C. E., Hoover, C. Beitelmal, M. H.,Shih, R. Kratos: Automated management of cooling capacity in data centers with adaptive vent tiles. In *the Proceedings of the ASME 2009 International Mechanical Engineering Congress & Exposition (IMECE2009)*, Lake Buena Vista, Florida, USA, 2009.

[9] Zhikui Wang, Cullen Bash, Christopher Hoover, Alan McReynolds, Carlos Felix, Rocky Shih. Integrated management of cooling resources in air-cooled data centers. In *6th annual IEEE Conference on Automation Science and Engineering*, Toronto, Ontario, Canada, 2010.

[10] A. Bemporad, S. Di Cairano, E. Henriksson, and K.H. Johansson. Hybrid model predictive control based on wireless sensor feedback: An experimental study. *International Journal of Robust and Nonlinear Control*, 20(2):209–225, 2010.

[11] C. Vermillion, J. Sun, K. Butts. Predictive control allocation for a thermal management system based on an inner loop reference modeldesign, analysis, and experimental results. *IEEE Transactions on Control Systems Technology*, pp(99), July 2010.

[12] T. F. Coleman and Y. Zhang. Optimization Toolbox For Use with MATLAB: User's Guide Version 2. 2003.

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

[14] Monem H. Beitelmal and Cullen E. Bash. *Optimizing centralized chiller system performance for dynamic smart cooling*. Technical Report No. HPL-2007-41, Hewlett Packard Laboratories, March 15, 2007.