Dynamic Thermal Management of Air Cooled Data Centers

Cullen E. Bash, Chandrakant D. Patel, Ratnesh K. Sharma
Hewlett-Packard Laboratories
Palo Alto, CA
Phone 650.236.2748, Fax 650.857.7029, cullen.bash@hp.com

ABSTRACT

Increases in server power dissipation have placed significant pressure on traditional data center thermal management systems. Traditional systems utilize Computer Room Air Conditioning (CRAC) units to pressurize a raised floor plenum with cool air that is passed to equipment racks via ventilation tiles distributed throughout the raised floor. Temperature is typically controlled at the hot air return of the CRAC units away from the equipment racks. Due primarily to a lack of distributed environmental sensing, these CRAC systems are often operated conservatively resulting in reduced computational density and added operational expense.

This paper introduces a data center environmental control system that utilizes a distributed sensor network to manipulate conventional CRAC units within an air-cooled environment. The sensor network is attached to standard racks and provides a direct measurement of the environment in close proximity to the computational resources. A calibration routine is used to characterize the response of each sensor in the network to individual CRAC actuators. A cascaded control algorithm is used to evaluate the data from the sensor network and manipulate supply air temperature and flow rate from individual CRACs to ensure thermal management while reducing operational expense. The combined controller and sensor network has been deployed in a production data center environment. Results from the algorithm will be presented that demonstrate the performance of the system and evaluate the energy savings compared with conventional data center environmental control architecture.

Keywords: Data center cooling, smart cooling, dynamic thermal control, thermal management, dynamic thermal management

NOMENCLATURE

- A — Area (m^2 or ft^2)
- C — Controller output
- G — Flow control gain ratio ($m^3/s{\cdot}C$)
- G_c — Compensator
- H — Control sensor output
- IT — Information Technology cost ($)
- J — CRAC Capacity Utilization Factor (Rated heat extraction capacity/Actual heat extracted)
- K — Amortization and Maintenance Factor
- L — Cooling Load Factor (Power required by cooling resources/Power dissipated by compute hardware)
- M — Number of personnel
- P — Power (W)
- R — Number of Racks
- S — Salary ($/person)
- T — Temperature (°C)
- TCI — Thermal Correlation Index
- U — Cost of Power ($ per KWh)
- $\dot{V}$ — Air Flow rate (m^3/s or CFM)
- σ — Cost of Software Licenses ($)

Subscripts

- crac — CRAC actuator
- critical — Critical data center resource e.g. space
- dep — Depreciation
- i — Rack index
- j — Actuator index
- lower — Lower bound
- ref — Reference
- return — CRAC return
- sup — Supply
- upper — Upper bound

INTRODUCTION

Data center thermal management challenges have been steadily increasing over the past few years due to rack level power density increases resulting from system level compaction [1][2]. These challenges have been compounded by antiquated environmental control strategies designed for low power density installations. The state of the art in data center thermal management consists of a single sensory feedback signal, which acts as a global indication of the heat being dissipated in the room and controls the temperature of the computer room air conditioner (CRAC) supply air as shown in Figure 1a. Other air-cooled configurations also exist as shown in Figure 1b and detailed elsewhere [3]. Typically the CRAC fan speed is fixed throughout operation. This mode of operation allows no local flexibility in how the cooling is delivered to the computers and there is no local state feedback information from different areas of the data center [4]. This style of operation requires less hardware, but is inefficient and not readily adaptable to changes in the environment. Without this flexibility, there is little potential to optimize the operation of the data center. Previous work in this area includes using CFD (computational fluid dynamic) models to create intricate mathematical models of the data center dynamics to find optimal layouts of data center in terms of energy efficiency [5][6][7].

Research has also begun on the dynamic optimization of the data center thermal environment. Bash et. al. have discussed the need for more sensing in the environment placed in closer proximity to critical equipment [4]. Patel et. al. have provided a vision for the use of such a sensor network coupled with standard CRAC components to enable a more energy-efficient control architecture [8]. Boucher et. al. have demonstrated the potential for utilizing standard data center environmental control components to provide dynamic thermal management [9]. Additional work has begun on expedient optimization methods for data center thermal management using model-based approaches. Optimization techniques based on the second law of thermodynamics have been introduced by Shah et. at. [10][11]. Rolander and Rambo have investigated techniques for rapidly assessing changes to the power and cooling parameters of an electronics rack [12][13]. Additionally, Moore et. al. have developed policies that optimize the power consumption of racked equipment based on local sensing of temperature [14].

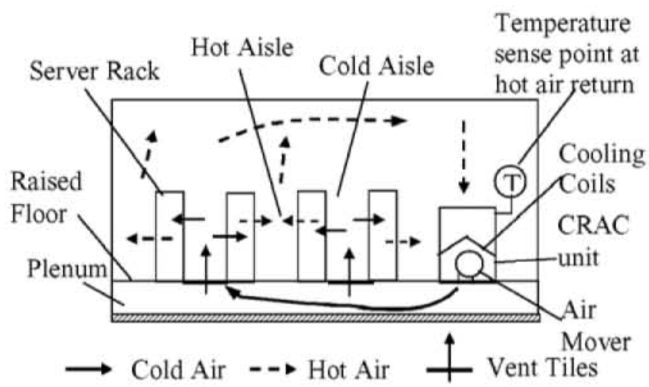

Figure 1a : Typical Raised-Floor Data Center

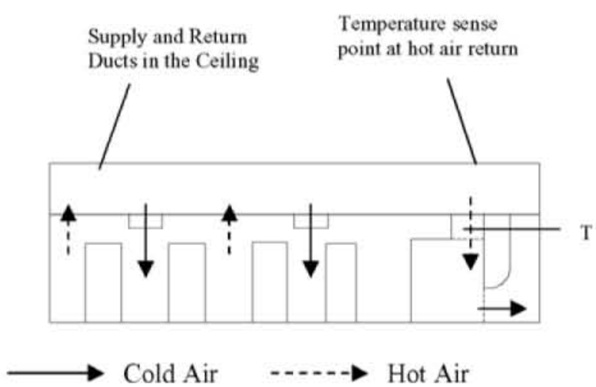

Figure 1b. Overhead Air Distribution

The present work introduces a data center environmental control system that utilizes a distributed sensor network to manipulate the distribution of cooling resources in an air-cooled environment. The control system has been deployed in an operational data center that utilizes a raised-floor infrastructure but is compatible with virtually any air-cooled environment. Experimental data will be discussed that compares the performance (dynamic and thermodynamic) of the controller with conventional control architectures.

CONTROLLER ARCHITECTURE

Conventional Data Center Control

Conventional feedback control systems typically require a plant function that describes the underlying behavior of the system to function satisfactorily. A compensator, such as a Proportional-Integral-Differential (PID), fuzzy logic, or other compensator, is generally combined with the plant to manipulate actuators in response to an error signal. The error signal is the difference between sensory data and the desired operating point of the system. The purpose of the compensator is to provide a stable, accurate response to the error signal [15].

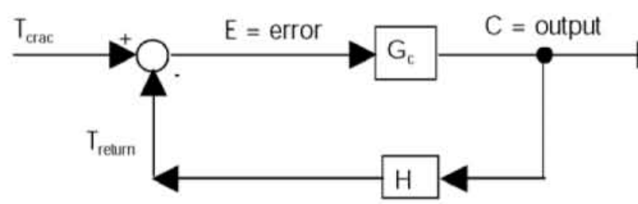

Figure 2 : Typical raised floor data center configuration

Figure 2 shows a block diagram of the feedback control system for a CRAC deployed in a data center environment operating as a self-contained unit. The CRAC return temperature (T_{return}) is compared to the set point (T_{crac}) and an error signal is subsequently fed into a compensator (G_c). The controller sends an output signal to the actuator (valve or compressor etc.) to adjust the temperature of the air supplied by the unit. The magnitude of the output signal is a function of the input error signal, history and system dynamics, and configuration of the compensator. Typically, air flow rate is not varied but in situations where a VFD is integrated into the actuator, the VFD output will vary linearly with the output of the compensator. Not included in Fig. 2 is a plant function that describes the response of the return sensor (H) to the actuation. The plant function is left out because the physics governing the sensor/actuator relationship are complex, involving non-linear fluid dynamics and heat transfer. Moreover, this relationship is affected by the operation of other actuators deployed in the environment and can be best described using the Reynolds-Averaged Navier-Stokes equations which are time-consuming to solve and not practical for use in a real-time controller [12].

Commissioning and Plant Function Discovery

Figure 3 shows a plan view of a typical raised-floor data center. Extending from each CRAC are bubbles indicating the extent of influence each CRAC has over equipment placed in the room. The computer equipment is represented by the rectangular rows in the figure. The shape of the regions defined by the bubbles are governed by the plant function of the system and are primarily influenced by data center geometry, layout, and CRAC flow rate with secondary dependencies on rack-level flow rate. The regions can be isolated, overlapping or discontinuous as shown in Fig. 3 and do not always follow intuition.

The commissioning of the control system is the process by which these regions of influence are discovered. It is conducted as follows:

1. Each actuator's supply air temperature is sequentially perturbed;
2. The subsequent change in temperature at each rack inlet sensor is recorded;
3. The Thermal Correlation Index is calculated for each sensor/actuator combination according to Eq. 1;
4. Regions of influence, or Families, are defined based on this response.

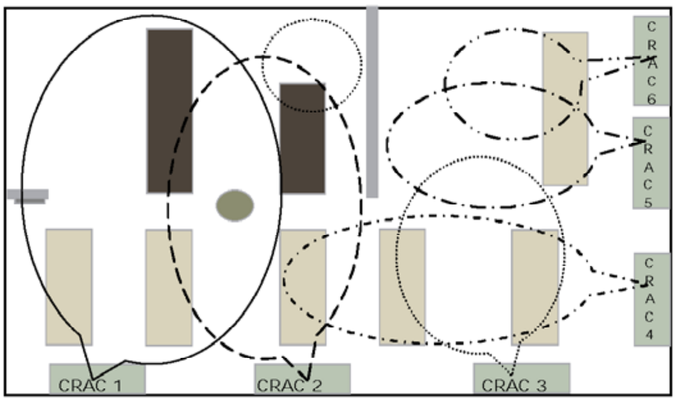

Figure 3 : CRAC Regions of Influence

Equation 1 is the Thermal Correlation Index (TCI). It quantifies the response at the i^{th} rack inlet sensor to a step change in the supply temperature of the j^{th} CRAC. TCI is a static metric based on the physical configuration of the data center. Since it does not contain dynamic information, it can be thought of as the steady-state gain to a step change in input. Additionally, it contains steady-state plant function information and can therefore be used in a control system. The regions of influence that are defined by the TCI metric are stable with time but are functions of data center geometry and infrastructure (e.g. vent tile arrangement) as well as CRAC flow rate uniformity. It is generally necessary to repeat the commissioning process after significant changes or modifications to the data center occur.

$$TCI_{i,j} = \frac{\Delta T_i}{\Delta T_{crac,j}} \qquad (1)$$

The commissioning process can be performed numerically via CFD or in-situ in the data center. In-situ measurements are more accurate while numerical simulations can be done prior to deployment of the controller or prior to green-field construction.

Control Algorithm Structure – Dynamic Smart Cooling

Figure 4 is a flow diagram of the control system architecture for the Dynamic Smart Cooling (DSC) controller. Rack inlet temperature and commissioning information is sent to a rule-based filter that employs thermal management and actuator operational criteria to choose a single sensor from each actuator's family to be used as the control sensor for that actuator. Thermal management criteria consist of the desired rack inlet temperature, which need not be uniform throughout the data center, along with margin of safety thresholds. Actuator operational criteria are primarily employed when multiple actuator regions of influence overlap and consider TCI values in the overlapping regions along with respective actuator operating points to determine the most efficient actuator for the particular sensor and eliminate actuators from interfering with one another.

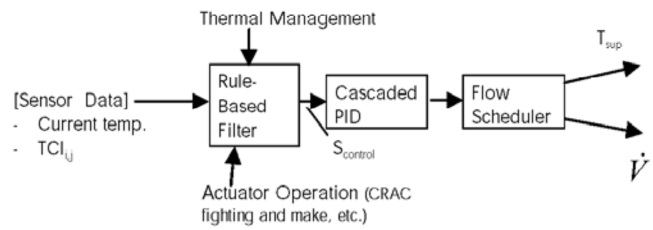

Figure 4 : DSC Controller Flow Diagram for Each Actuator

After filtering the sensor information, an error signal for the actuator ($S_{control}$) is generated and sent to a PID compensator from which the new supply air temperature set point is generated for each actuator. The supply air temperature is fed to a flow scheduler which determines the new air flow rate for each actuator as a function of supply air temperature. The specific functionality is dependent upon actuator type and infrastructure. These new set point signals are transferred to the actuators via the communications infrastructure described further in the next section. In this manner, the initial Multi-Input Multi-Output control system has been simplified to a series of Single-Input Dual-Output controllers. The resulting controller is general and is designed to function with any open, shared environment in which multiple actuators are employed to manage overlapping regions of influence. As such, it is designed to work with current under floor cooling infrastructures along with the increasing varieties of newer configurations meant for higher density deployments like overhead, in-row, and rack-mounted cooling.

EXPERIMENTAL PROCESS

Experiments were conducted in a production data center in Palo Alto, California to gauge the performance of the controller.

Data Center Layout

Figure 5 is a plan view of the experimental facility. The data center is air cooled with a 24 inch raised floor plenum to distribute cool air, power and networking. Six 105 kW (30 ton) CRAC units are arranged around the perimeter of the room. CRACs 1-4 in Fig. 5 are chilled water cooled (Liebert model FH600C-AAEI) while CRACs 5 and 6 are capable of utilizing facility chilled water or operating in a self-contained manner with an internal vapor-compression refrigeration loop driven by two dual-staged compressors (Liebert model DE412W-AAEI). Each CRAC contains two internal temperature control sensors, one is placed on the return side of the CRAC unit as shown in Fig. 1 and the second is placed on the supply side. Only one sensor is active according to the control architecture employed during the experiment. Furthermore, blowers within each CRAC are operated via variable frequency drives (ABB model ACH400) connected to the communications infrastructure as described below.

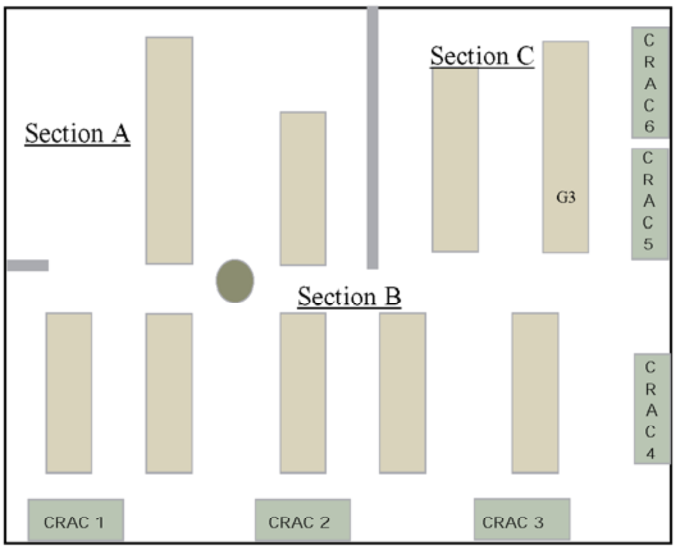

Figure 5 : Experimental Test Bed

The data center is configured into three distinct regions labeled Section A, B and C in the figure. Section A houses equipment used to support customer engagements. Section B houses equipment to support the IT needs of the business unit, and Section C houses equipment used to support the research functions of the lab. This latter section can be isolated from the rest of the room via flow dampers under the raised floor at the borderline of the sections and a curtain above the floor between Sections B and C. A solid wall above the floor separates Sections A and C. In this manner, isolated environmental control experiments can be conducted within Section C with minimal impact on the larger room.

The data center contains approximately 230 m^2 (2500 ft^2) of raised floor space and houses 76 racks of equipment. Total power dissipation of the computational equipment varies from 100 kW – 500 kW with a maximum distributed power density of 1800 W/m^2 (200 W/ft^2). Local power densities can exceed 3100 W/m^2 (350 W/ft^2). Over 200 temperature sensors are attached to racks around the room according to ASHRAE guidelines and are described in more detail below [16]. Power into the data center is monitored at power distribution units in the room. Additionally, power consumed at the rack and by the CRAC actuators are monitored by individual power meters (Power Distributuion Inc. model ION 7330).

Data Collection Architecture

Figure 6 shows a schematic of the data collection architecture. Data is collected from CRAC units, VFD units, PDUs and temperature sensors deployed in the data center. CRAC unit and VFD unit data is collected through the OPC server, as shown in Figure 6. Also shown in Fig.6 are the sensor networks mounted on racks, which provide temperature data over the IP network. The sensor network consists of digital temperature sensors connected to a base station in different

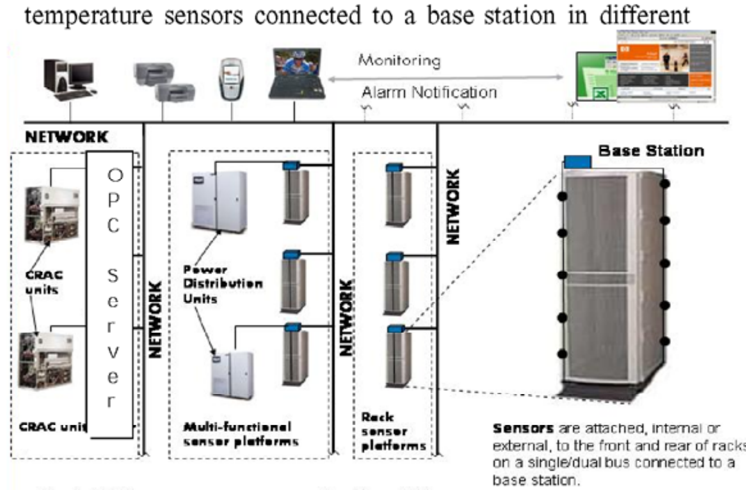

Figure 6: Schematic of the data collection architecture

topologies. The base station is connected to the data center IP network. The base station collects temperature data from the sensors and sends the data packets to monitoring clients or alarm notification devices. When temperatures exceed a preset limit, secure alarm notification is provided to pagers, cell phones and PDAs. Monitoring clients include a web-based authenticated intranet tool with database and visualization functions. This plug and play network is self-configuring requiring minimal or no human intervention upon deployment. Figure 6 is a schematic of the sensor network architecture.

The digital output from each sensor is accurate to ± 0.5°C in the range of interest. Since the sensor is primarily a transistor with compensation for leakage, no calibration is needed.

Communications System

The communication fabric supporting the smart cooling control and monitoring of the data center is primarily event – driven and IP-based; with field devices connected on a floor level network with appropriate IP interfaces. In case of non-IP devices, the proprietary field devices are accessed using standard, interoperable interfaces. Event driven interfaces are provided to devices that are incapable of doing so. For legacy devices with proprietary buses, appropriate drivers and bus converters are used to convert data packets into Ethernet frames. Device permitting, all communications are two-way and have built-in redundancies.

RESULTS/DISCUSSION

Two series of experiments were performed to gauge both the thermodynamic and dynamic performance of the control system and compare it to a conventional scheme. Unless otherwise noted, CRACs configured in "conventional" mode utilize their internal sensor located in their return air streams for control and are operated at a constant air flow rate of 95% maximum throughout the experiment. An internal PID controller on each unit is used to regulate the unit's operation. Under conventional control, no information from the distributed sensor network is used. CRACs operated in "dynamic smart cooling" (DSC) mode utilize the internal sensor located in their supply stream to adjust operation according to the set points they receive from the controller as described above. Airflow rates are also adjusted in real-time by the controller according to Eq. 2. In DSC mode, set points for all rack inlet sensors are independently adjustable and were set to a uniform 25 C unless otherwise noted. For all experiments, the CRACs are configured to operate in chilled water mode unless otherwise noted.

Energy Consumption

The energy consumption of the environmental control system operating in conventional mode was compared with operation in DSC mode. All sections of the data center (i.e. all six CRACs) as shown in Fig. 5 were used in this experiment. Two DSC modes were evaluated to gauge the relative impact on energy consumption. Furthermore, in Uniform mode, the actuator regions of influence (Fig. 3) do not vary with VFD speed.

- Uniform Flow: All CRAC blowers were set to the same flow rate according to the highest blower set point of all the CRACs in the data center.
- Independent Flow: CRAC blower speeds were allowed to vary independently according to the set points generated by the controller.

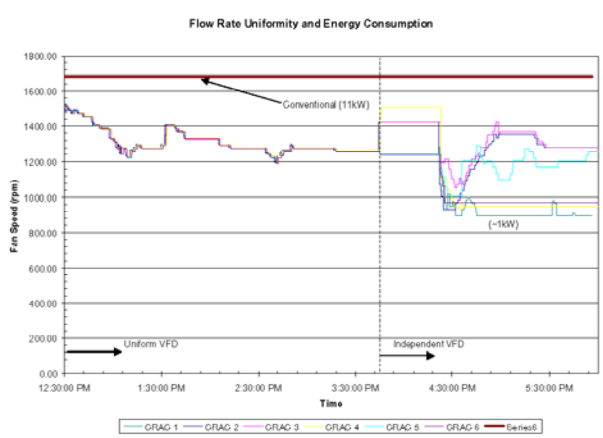

Figure 7 : Blower Speed versus Time

Figure 7 charts the blower speed with time for each of the DSC experimental scenarios. The blower speed in the conventional experiment overlays the DSC data. Steady-state is reached when the blower speed fluctuations cease. Note that the blower speeds for each actuator separate under independent flow control prior to reaching steady state.

Table 1 lists the energy consumption of the environmental control system (CRACs and central chiller) for each of the experiments. Energy consumption of the central chiller used to supply chilled water to the CRACs was not measured but was estimated by varying chiller coefficient of performance (COP) with CRAC supply temperature. The results show a reduction in energy consumption of 49% using uniform flow and 58% using independent flow with most of the savings coming from the reduction in flow work (via reduction in blower speed). These results are applicable to the data center under test only. Changes in data center loading, in particular, will impact energy savings and will be discussed further in the next section. Nevertheless, the savings in energy stem from the fact that, with a distributed sensor network in place, CRAC units can be operated less conservatively than in conventional mode where rack inlet temperatures are largely unknown.

Table 1 : Experimental Test Bed

Experiment	Power (kW)	Savings (%)
Conventional	83.6	-
Uniform Flow	42.9	49
Independent Flow	35.4	58

System Performance

The previous results demonstrate the impact of having more information on the operation of an environmental control system. More information typically translates to less conservative operation. However, with less conservative operation comes reduced operating margins. Therefore, experiments were conducted to compare the dynamic response of the DSC controller with that of a conventional controller. For each of these experiments, only the research area in Fig. 5 was used and was configured to be isolated from the rest of the data center. In this configuration, the research area is N+1 redundant in CRAC actuators. Experiments were configured to mimic data centers with underutilized environmental control equipment (30% utilization) and well-utilized equipment (80% utilization). Part way through each experiment one of the CRACs was intentionally shut down to mimic a failure scenario.

Figure 8 is a chart showing the behavior of a conventional controller in an underutilized facility when CRAC 6 fails. The total power dissipation in the section was 31 kW resulting in a CRAC utilization factor (J) of 3.3. CRAC 5 supply and return temperatures are shown along with the inlet temperature at the hottest rack (Rack G3 in Fig. 5) in Section C. Since the CRAC is responding only to changes in its return temperature, the rack inlet temperature rises 3.4 C where it stabilizes. Figure 9 shows a similar experiment with DSC control. The supply temperature is displayed along with inlet temperature at rack G3. Rack G3 set point temperature was 24 C and the air flow rates of CRACs 5 and 6 were 63% and 55% respectively. Upon failure of CRAC 6, the rack inlet temperature increases by 3 C and, as the CRAC supply temperature drops, is eventually reduced below its value prior to CRAC 6 failure, preserving thermal management. In this configuration and prior to the failure, CRACs 5 and 6 consume approximately 60% less energy than in the conventional approach yet the system is able to respond faster to the failure due the use of the locally distributed sensor network. Additionally, since control is at the rack inlet level rather than the CRAC return, the rack inlet temperature is reduced to the set point condition rather than stabilizing at the higher temperature.

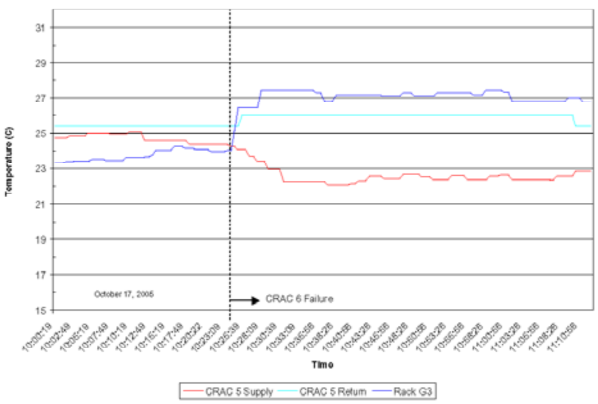

Figure 8 : Underutilized with Conventional Control

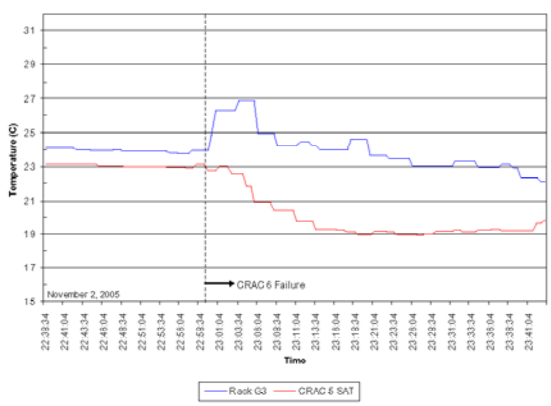

Figure 10 : Underutilized with DSC Control

Figure 10 shows the result of an experiment similar to that of Figure 8 except Section C is well-utilized by maximizing the operation of several racks in the area to increase the load by an additional 50 kW reducing the CRAC utilization factor (J) to 1.25. After the failure of CRAC 6, the inlet temperature at the critical rack increases 5.6 C and stabilizes. CRAC supply air temperature does not change but the return air temperature was kept within set point by the operation of the internal chilled water valve. The results indicate that most of the load in the room prior to the failure of CRAC 6 was already taken by CRAC 5 and that the failure of CRAC 6 primarily impacted the airflow distribution under the plenum, and subsequently, through the ventilation tiles. Figure 11 is the same experimental condition as Fig. 10 with DSC control. The set point of rack G3 is 23 C. Prior to the failure, the blowers of CRACs 5 and 6 were operating at 77% and 80% respectively. At these operating conditions the CRACs consume 28% less energy than in the conventional mode of operation. The savings is lower than in the underutilized case and reflects the fact that increased energy savings are possible when compute equipment power consumption is reduced. The temperature rise at rack G3 subsequent to the failure was 5.0 C, lower than with conventional control but within the margin of error of the experiment.

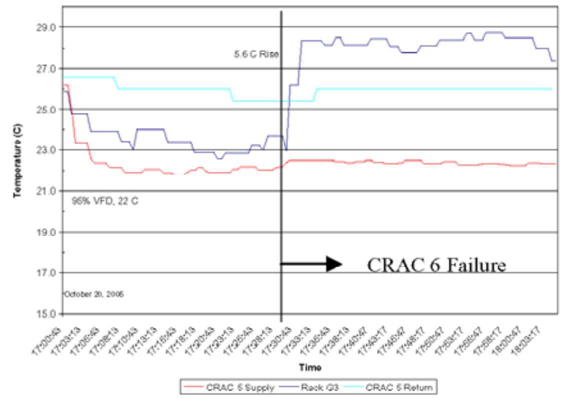

Figure 9 : Well-utilized with Conventional Control

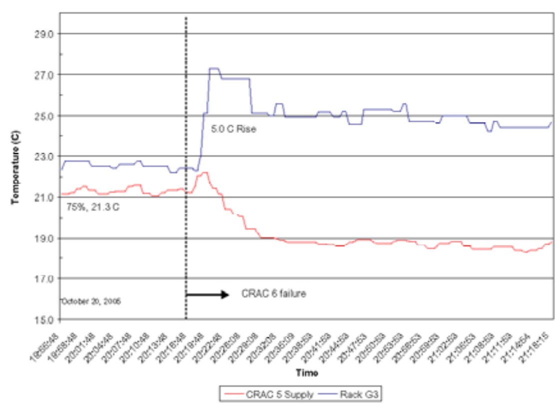

Figure 11 : Well-utilized with DSC Control

Figure 12 shows the results of a complimentary experiment to that of Fig. 11 but with the CRAC blowers held to a constant 95% of maximum throughout the test. This provides a direct comparison to the conventional controller for evaluating the speed of the dynamic response. After the failure of CRAC 6, the inlet of rack G3 rises by 4.0 C which represents a 30% reduction compared to conventional control. This, along with the results from Fig. 11, indicates that Section C can handle additional load with the same operational margin of safety as a conventional controller. Quantification of this additional load can be difficult to ascertain precisely but if we assume that it is evenly distributed among the compute equipment and doesn't impact local recirculation we can estimate that 30% additional load (24 kW) can be handled with DSC for the same response characteristics as a conventional controller.

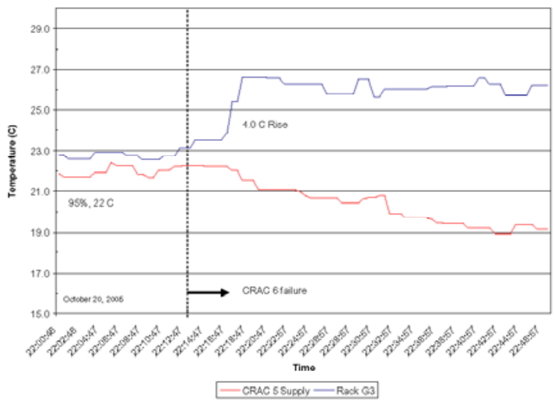

Figure 12 : Well-utilized with DSC Control at 95%

Apart from energy savings, these results indicate that the utility of the actuators has significantly increased with DSC over conventional control. This increased utilization can benefit the operation of the data center in several ways:

1. The environmental control system can operate on thinner margins with equal or better reliability than conventional control due to the improved dynamic response. Thinner margins translate directly to improved operational efficiencies.
2. Improved utilization potential can enable increased computational equipment power density without the need for additional hardware.
3. Higher margins of safety can be realized than with conventional control.

The above results show the dynamic response of the controller due to the loss of an actuator. Data on the controller's response to local disturbances (vent tile blockage, etc.) has not yet been presented but it should be clear that this system will respond in a manner that is more effective than conventional controllers given the proximity of the sensors to the disturbance. Additionally, data on refrigerated DX units has not been presented but, given that these units are typically less efficient then chilled water systems, the savings afforded by the DSC controller are expected to increase over that realized in this work.

Data Center Cost Model

The improved utilization of critical data center space and savings in recurring cost of power can be quantified through application of an appropriate cost model. Patel and Shah introduced a model that examines the burdened cost of power delivery and can be used to quantify the savings [17]. The savings can be used to determine the payback associated with addition of the dynamic smart cooling elements with respect to savings associated with recurring cost of power and utilization of critical space and resources. The model, as shown by Equation 2, accounts for real estate by use of standard appraisal techniques in a given geography to determine cost per unit area ($A_{critical}$). It accounts for capital expense and maintenance of the redundant power and cooling infrastructure such as uninterruptible power supplies, generators, hydronics, chillers, pumps, air handling units, etc by applying "burdening" factors (K_1 and K_2 for power and cooling respectively) to the power consumed ($P_{consumed\ hardware}$) by the compute, network and storage hardware in the data center. It applies the cooling load factor, L_1, to account for the amount of electricity used by the cooling equipment to remove a given watt of power from the compute hardware in the data center. The model uses standard electric grid pricing ($U_{\$grid}$). Furthermore, within the burdening factors K_1 and K_2, the model applies a factor (J) to account for utilization of expensive capital equipment. As an example a data center with 1 MW of power and cooling capacity being used at a compute load of 100 KW would have a J factor of 10. Thus lowering of J results in savings by reducing the burdening factors K_1 and K_2 and better realization of capital and maintenance costs.

The second half of the equation is used to determine software licensing and personnel cost, where M represents the number of personnel servicing a given rack R. S represents the salary, IT_{dep} represents the depreciation of compute equipment and σ_1 represents software licensing cost per rack.

$$Cost_{total} = \left(\frac{\$}{m^2}\right)\left(A_{critical}, m^2\right) + \left(1 + K_1 + L_1 + K_2 L_1\right) U_{\$,grid} P_{\substack{consumed \\ hardware}} + R\left(M_{total} S_{avg} + IT_{dep} + \sigma_1\right) \quad (2)$$

The availability of dynamic control scheme outlined in this paper when applied to this model enables cost savings by:

- Better assessment of actual compute load to rated capacity, and enabling strategies for response to failures, thereby increasing the utilization and reducing the utilization factor, J, by approximately 30% in a well-utilized facility.
- Reduction of the cooling load factor, the multiplier (L_1) used in the model, by half results in direct savings in electricity costs for cooling equipment based on the diurnal heat load profile of the racks. As an example, when racks of computer equipment idle for half a day at 60% of maximum power dissipation, direct savings in recurring cost of cooling results from the use of "dynamic" cooling control system. The load factor for the data center discussed in this report is reduced from approximately 0.8 to 0.4 when the controller is changed from conventional mode to DSC.

- Proper static and dynamic provisioning of capital intensive cooling resources reduces the amortization and maintenance factor (K_1, K_2) by reduction in excessive redundancy for 24/7 operation. As an example, in a given section of the data center shown in Fig 5., the quantified results can enable the use of one less CRAC (17% reduction in usage) with associated amortization and maintenance burden.
- Automation enabled by sensing and control, while not quantified in this paper, can reduce the number of personnel per rack (M) in the cost model.

CONCLUSION AND NEXT STEPS

This paper outlines an architecture and control scheme for dynamic thermal management of air cooled data centers. The control scheme was tested in a "smart" data center [8], and results show 50% reduction energy consumption by cooling resources in addition to improvement in use of critical data center space. In order to enable further quantification, a general cost model defined by Eq. 2, was proposed. Application of the savings to the cost model can reduce the total cost of space and power by 25%. The reductions stems from savings in recurring cost of power used by cooling resources (50% improvement) and better utilization of space.

Future research will focus on the utility of the Thermal Correlation Index, the response characteristics of DX actuators, and incorporation of local actuators into the control architecture.

ACKNOWLEDGMENTS

The authors would like to acknowledge Tim Boucher and Cliff Federspiel, previously at University of California at Berkeley and David Auslander at UC Berkeley for their numerous contributions during the initial stages of this research and in particular their contributions to the commissioning process. We would also like to thank Sara McAllister of UC Berkeley for her assistance with conducting experiments.

REFERENCES

[1] The Uptime Institute, "Heat density trends in Data Processing, Computer systems and Telecommunications Equipment", White Paper issued by *The Uptime Institute*, 2000

[2] C.D. Patel, "A vision of energy aware computing – from chips to data centers", *The International Symposium on Micro-Mechanical Engineering*, ISMME2003-K15, December, 2003

[3] Shrivastava, S., Schmidt, R., Sammakia, B., Iyengar, M., "Comparative Analysis of Different Data Center Airflow Management Configurations", Proceedings of IPACK2005, San Francisco, CA, July 2005

[4] Bash, C.E., Patel, C.D., Sharma, R.K., "Efficient Thermal Management of Data Centers – Immediate and Long-Term Research Needs", *Intl. J. HVAC&R Res.*, Vol.9, No.2, pp., Apr 2003

[5] C.D. Patel, C.E. Bash, C. Belady, L. Stahl, D. Sullivan, "Computational fluid dynamics modeling of high compute density data centers to assure system inlet air specifications", Proceedings of IPACK'01 – The PacificRim/ASME International Electronics Packaging Technical Conference and Exhibition, Kauai, HI, July 2001

[6] C.D. Patel, R.K. Sharma, C.E. Bash, A. Beitelmal, "Thermal Considerations in Cooling Large Scale High Compute Density Data Centers", ITherm 2002 – 8th Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems, San Diego, CA, May 2002

[7] Van Gilder, J., Schmidt, R., "Airflow Uniformity Through Perforated Tiles In a Raised-Floor Data Center", Proceedings of IPACK2005, San Francisco, CA, July 2005

[8] Patel, C.D., Bash, C.E., Sharma, R.K, Beitelmal, A and Friedrich, R.J., "Smart Cooling of Datacenters", Proceedings of IPACK'03 – The PacificRim/ASME International Electronics Packaging Technical Conference and Exhibition, Kauai, HI, IPACK-35059, July 2003

[9] Boucher, T., Auslander, D.M., Bash, C.E., Federspied. C., Patel, C.D., "Viability of Dynamic Cooling Control in a Data Center Environment", *Journal of Electronic Packaging*, June 2006 (Pending Publication)

[10] Shah, A., Bash, C.E., Carey, V., Patel, C.D., "Exergy-Based Optimization Strategies for Multi-Component Data Center Thermal Management: Part I, Analysis", Proceedings of IPACK2005, San Francisco, CA, July 2005

[11] Shah, A., Bash, C.E., Carey, V., Patel, C.D., "Exergy-Based Optimization Strategies for Multi-Component Data Center Thermal Management: Part II, Application and Validation", Proceedings of IPACK2005, San Francisco, CA, July 2005

[12] Rolander, N., Rambo, J., Joshi, Y., Mistree, F., "Robust Design of Air-Cooled Server Cabinets for Thermal Efficiency", Proceedings of IPACK2005, San Francisco, CA, July 2005

[13] Rambo, J., Joshi, Y., "Reduced Order Modeling of Steady Turbulent Flows Using the POD", Proceedings of ASME Summer Heat Transfer Conference, San Francisco, CA, July 2005

[14] Moore, J., Chase, J., Ranganathan, P., Sharma, R., "Making Scheduling "Cool":Temperature-Aware Workload Placement in Data Centers", Usenix 2004

[15] Phillps, C., Harbor, R., *Feedback Control Systems*, 2^{nd} edition, Prentice Hall, New Jersey, 1988.

[16] *Thermal Guidelines for Data Processing Environments*, American Society of Heating, Refrigerating and Air-Conditioning Engineers, Inc, Atlanta, 2003

[17]Patel, C. D, Shah, A., "Cost Model for Planning, Development and Operation of a Data Center", *Hewlett-Packard Laboratories Technical Report,* HPL-2005-107(R.1), June 9, 2005