sProceedings of the ASME 2017 International Technical Conference and Exhibition on
Packaging and Integration of Electronic and Photonic Microsystems
InterPACK2017
August 29-September 1, 2017, San Francisco, CA, USA

InterPACKICNMM2017-74105

COMPARISON AND EVALUATION OF DIFFERENT MONITORING METHODS IN A DATA CENTER ENVIRONMENT

M. Tradat, S. Khalili, B. Sammakia
State University of New York at Binghamton- SUNY
Binghamton, NY, USA

M. Ibrahim, Th. Peddle, A. Calder, B. Dawson
Panduit Corporation, Tinley Park, IL, USA

M. Seymour, K. Nemati
Future Facilities Ltd.
London, UK

H. Alissa
Microsoft
Redmond, WA, USA

ABSTRACT

The operation of today's data centers increasingly relies on environmental data collection and analysis to operate the cooling infrastructure as efficiently as possible and to maintain the reliability of IT equipment. This in turn emphasizes the importance of the quality of the data collected and their relevance to the overall operation of the data center. This study presents an experimentally based analysis and comparison between two different approaches for environmental data collection; one using a discrete sensor network, and another using available data from installed IT equipment through their Intelligent Platform Management Interface (IPMI). The comparison considers the quality and relevance of the data collected and investigates their effect on key performance and operational metrics. The results have shown the large variation of server inlet temperatures provided by the IPMI interface. On the other hand, the discrete sensor measurements showed much more reliable results where the server inlet temperatures had minimal variation inside the cold aisle. These results highlight the potential difficulty in using IPMI inlet temperature data to evaluate the thermal environment inside the contained cold aisle.

The study also focuses on how industry common methods for cooling efficiency management and control can be affected by the data collection approach. Results have shown that using preheated IPMI inlet temperature data can lead to unnecessarily lower cooling set points, which in turn minimizes the potential cooling energy savings. It was shown in one case that using discrete sensor data for control provides 20% more energy savings than using IPMI inlet temperature data.

Keywords: Experimental Study, Data Centers, Raised Floor, Containment, IPMI, Supply Air Temperature, Cold Aisle.

NOMENCLATURE

BMC	Baseboard Management Controller
CAC	Cold Aisle Containment
CPU	Central Processing Unit
CRAC	Computer Room Air Conditioning
CRAH	Computer Room Air Handler
DC	Data Center
DCM	Data Center Manager
FMS	Facility Management System
ICT	Information and Communication Technology
IPMI	Intelligent Platform Management Interface
N	Fan Rotational Speed
P	Fan Power
PDU	Power Distribution Unit
PUE	Power Usage Effectiveness
RAT	CRAH Return Air Temperature
RPM	Revolution Per Minute
RTT	Ride Through Time
SAT	CRAH Supply Air Temperature
SNMP	Simple Network Management Protocol
tPUE	total Power Usage Effectiveness

INTRODUCTION

Data center cooling energy efficiency is critical to the successful operation of modern large data centers. In 2014, data centers in the U.S. consumed an estimated 70 billion kWh, representing about 1.8% of total U.S. electricity consumption [1]. Given that the cooling infrastructure can average 40% of the total data center energy consumption [2, 3], then the data center cooling energy consumed in 2014 can be approximated at 28 billion kWh. These numbers indicate that improving airflow management in order to raise the efficiency of cooling in data centers can significantly affect operating costs and allow for increased IT capacity, thereby extending the life of the data center. Some of the methods used to improve airflow include, but are not limited to, containment, IT equipment configuration changes, bypass air management (e.g. cable penetrations), recirculation management (e.g. blanking panels). These lead to opportunities to raise cooling energy efficiency by air or waterside economization, variable frequency drives, and increased IT equipment inlet temperatures, etc.

Air containment in its many forms simply provides a physical separation between the supplied cool air and the cabinet hot exhaust air, optimizing airflow distribution in the data center room by preventing mixing of cold and hot air streams. Over the past decade, containment has greatly matured from a strategy considered by data center managers to manage higher thermal loads [4-8], to one of the most widely used strategies for improving the thermal management of data centers. Per the 2014 Uptime Institute Survey [9], 80% of the 1,000 data center operators and IT practitioners surveyed indicated their use of cold or hot aisle containment to improve data center efficiency. This wide use of containment has also driven numerous research efforts in understanding various aspects of using containment [10-16]. Shrivastava et al. [10] presented a comparison between different types of containment systems from the perspective of the cooling energy cost and performance. In addition, guidelines for choosing suitable containment arrangements were introduced. Patterson et al. [11] investigated entrained warm air effect into cold aisle containment (CAC). Their results showed that the recirculation significantly affect the inlet of the lowest servers. Shrivastava and Ibrahim [12] showed the positive impact of cold aisle containment systems on the Ride Through Time (RTT) during failure. They showed that the CAC system increases the RTT five times more. Alissa et al. [13] provided quantitative and qualitative measurements for data centers transient performance during cooling failure in open and contained environments. Their results showed very different responses of the IPMI data, fan RPM, CPU temperature and internal server temperature sensors during failure for CAC compared to the case of open aisle. In addition, they concluded an overestimate of RTT by 70% based on the external inlet air temperature and these temperature fields do not necessarily reflect the IT equipment thermal performance. Makwana et al. [14] investigated the importance of containment sealing. They stated that sealing the containment maximizes the benefits of CAC. Sundaralingam et al. [15] used a multi-dimensional array of sensors for flow management of air inside the CAC system. They introduced and explained that the selection criteria of the CAC based on rack inlet temperature only may not be a best practice. In addition, the authors recommend an over-provisioning for fully sealed contained aisles. Muralidharan et al. [16] investigated the impact of the CAC on thermal performance of data centers. The authors quantified the thermal impact of the CAC by comparing it with different open arrangements (open hot aisle/cold aisle). The study considered different cabinet heat loads at two different CRAH (Computer Room Air Handler) unit RAT (Return Air Temperature) set points. Their results showed a 22% energy saving using the CAC systems over conventional open hot aisle/cold aisle.

The true benefit of containment lies in the separation of cold and hot air streams, which provides the opportunity to closely match cooling airflow to IT equipment airflow, thereby promoting a uniform cabinet inlet temperature profile. It also allows the increase of the cold air supply temperature, while maintaining inlet temperatures acceptable for the deployed IT equipment, which results in cooling energy savings and increased cooling equipment efficiency [17]. Therefore, to truly gain the benefits of containment, an effective monitoring system must be used to accurately measure IT equipment inlet temperatures, as well as IT equipment airflow needs.

Legacy data centers use CRAH or CRAC (Computer Room Air conditioning) return air temperature sensors as the primary control-variable to adjust the air temperature and air volume supplied to the IT equipment. This control approach significantly limits energy efficiency because it does not control the air temperatures available for IT equipment cooling and nor does it verify the temperatures actually delivered. Server manufacturers have agreed that their main operational parameter is the air temperature provided at the inlet of the server itself, not the proxy temperature returning to the cooling device. Therefore, a reasonable hypothesis is that a much higher degree of monitoring and efficient control would be achieved by using front-panel, inlet temperature sensor data. The majority of servers have a host of platform information available from their Information and Communications Technology (ICT) management network. Server front-panel, i.e., server inlet, air temperature is monitored and available over the network through each server's management connection that supports Simple Network Management Protocol (SNMP) or Intelligent Platform Management Interface (IPMI). IPMI and now-standard hardware called a Baseboard Management Controller (BMC) - allow remote administrators to monitor the health of servers, deploy (or remove) software, manage hardware peripherals like the keyboard and mouse, reboot the system and update software on it [18].

More advanced data centers base their control policy on distributed discrete sensors mounted at the inlet of the IT equipment (or at least the front door of the rack). Finally, some suggested to use the front-panel temperature sensor data from the ICT network as input information to the facility management system (FMS) for control of the cooling system in a data center. Patterson et al. [19] investigated the appropriate control strategy in a contained environment, where they tested three control designs based on temperature, pressure and velocity. The study showed that the best level of control to supply sufficient airflow to the IT equipment in a contained cold aisle is using pressure based control. Nishi et al. [20] addressed the cooling inefficiency resulting from airflow mismatch between the cooling requirements of the IT equipment vs the supply air conditions from the facility-cooling infrastructure. They proposed and outlined a method to estimate the real time volumetric airflow based on fan's RPM. Then, the estimated volumetric airflow and IT exhaust temperature were used as input parameters to the Intel Data Center Manager (DCM) by using the IPMI commands. Alissa et al. [21] showed that the server's IPMI average fans' speed, and discrete pressure reading from containment can be used to generate a flow curve model. This model collapses the server impedance and effective total fan curve into one. The flow curve can be used in real time airflow prediction that is inclusive of all operational CAC pressure differential values. Tradat et al. [22] showed that the difference between the discrete and IPMI inlet temperature of the IT equipment increased as SAT increased. This was due to the negative pressure differential build up inside the containment. Furthermore, the authors identified a value of the supply air temperature at which IT equipment fans start to ramp up. To the authors' knowledge using IPMI data as a monitoring and control strategy for data center cooling system is scarce in the literature.

This study presents an experimentally based analysis and comparison of environmental data collection using two different approaches; one using a discrete sensor network, and another using available data from installed IT equipment (IPMI data). The comparison looks closely into the effect of using both approaches for data center cooling control.

DATA CENTER LABORATORY

The ES2-Binghamton University Data Center Laboratory was used for all the testing conducted for this study. The lab is a 2,315 ft^2 (215 m^2) space with a 3 ft. (0.91 m) raised access floor. It is equipped with a down-flow chilled water based cooling unit, which is rated at 32 tons (114 kW) of cooling capacity and 16,500 CFM of airflow capacity. The unit is equipped with a variable frequency drive on its blower motor so that airflow can be modulated. IT equipment cabinets are placed in the laboratory in a traditional alternating hot aisle/cold aisle arrangement. Aisle C, of primary interest in this study, is a contained cold aisle, with end-of-aisle doors and a horizontal barrier across the aisle at the cabinet tops. A map of the aisle locations and tile numbering system is shown in Figure 1. [23].

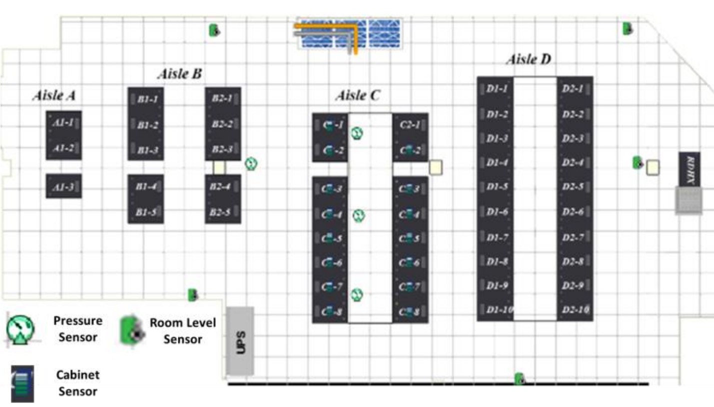

(a)

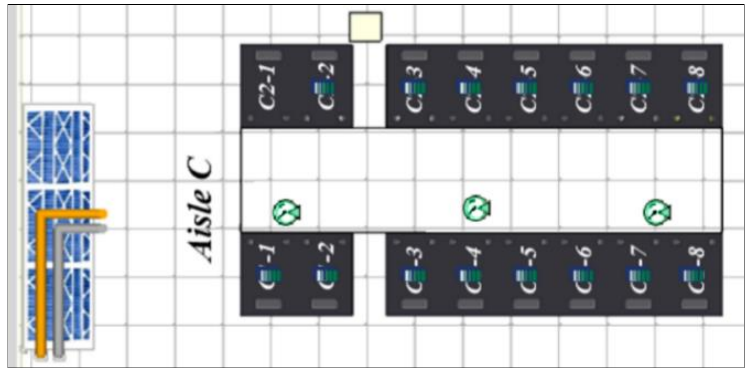

(b)

(c)

Figure 1. (a) Data center laboratory layout including **t**emperature and pressure sensors locations as shown in SynapSoft™. (b) Aisle C tiles and rack matrix. (c) Discrete sensors locations per rack [red circles].

Aisle C is comprised of two rows, with 8 cabinets per row. A total of 242 IT servers are deployed in the 16 cabinets of Aisle C, and all empty RU slots were blanked off. Server types and quantities used in Aisle C are listed in Table 1.

Table 1. Aisle C IT Inventory. [24].

IT Make	Number	Unit Active power[W]
Dell™PowerEdge™2950	128	386
Dell™PowerEdge™R520	64	165
Dell™PowerEdge™C2100	14	281
HP ProLiant DL385 G2	36	330

MEASUREMENTS METHODOLOGY

Volumetric airflow supplied to Aisle C via perforated floor tiles was measured with a Flow Hood, which is a back pressure compensated airflow measurement device, and found the aisle's airflow to be 9530 CFM at SAT of 78 °F. Differential pressure between contained Aisle C and the laboratory air space was measured using a multimeter (ADM-850L). Air temperatures were gathered using multiple sensor types. Discrete SynapSense^{TM} temperature sensors, distributed as shown in Figure 1, marked by red circles; were used to record air temperature in Aisle C. They were located at the inlets to the IT servers with a measured uncertainty of ± 0.5 °F depending on the temperature range. The inlet air temperatures were also reported by IT equipment using IPMI data. The IPMI data also include the server's fan speed, CPU temperature and active power.

EXPERIMENTAL SETUP

During all tests, the IT servers were exercised at 75% CPU utilization through the Linux operating system, which resulting in a total IT power consumption of 57 kW. The first set of test cases were performed with cooling unit supply air temperature set points varying from 64 °F to 78 °F in 2 °F increments. Cooling unit blower speed was adjusted via the variable frequency drive to maintain a neutral or very slightly positive pressure differential in Aisle C compared to the rest of the laboratory space $(0 - 0.002'' \ H_2O)$. This measurement was treated as an indicator of balanced volumetric airflow. That is, the amount of air being supplied into Aisle C by the cooling unit was equal to, or slightly higher than, the amount of air being drawn through the servers in Aisle C. Each set of test conditions was allowed to operate for an extended time period to ensure that a steady state condition had been achieved.

In another set of test cases, the cooling unit supply air temperature was set at 64 °F, 72 °F, 74 °F, 76 °F, and 78 °F with the differential pressure of Aisle C maintained at $-0.01'' \ H_2O$. The negative pressure is treated as an indicator of under-provisioning Aisle C, where there is less volumetric air supplied by the cooling unit compared to the airflow drawn by the servers.

A final test case was conducted at a cooling unit supply air temperature of 64 °F and a differential pressure of $0.01'' \ H_2O$ in Aisle C. The positive pressure is treated as an indicator of over-provisioning Aisle C, where there is more airflow supplied by the cooling unit compared to the airflow drawn by the servers.

RESULTS AND DISCUSSION

IPMI DATA ANALYSIS

This section focuses only on the server level data as provided by the IPMI protocol to better understand the behavior of the three classes of servers installed. Figure 2 (a) shows the behavior of the inlet temperature of all servers installed in Aisle C as provided by the IPMI protocol with a varying SAT set point. As expected, as the SAT increases, the IPMI inlet temperatures of all servers increase regardless of server type or location. The figure also shows that for all SAT set points, the Dell PowerEdge 2950 -112 server at the top of rack C1-7 reports the highest IPMI inlet temperature. In taking a closer look at the Dell PowerEdge 2950 -112 server in Figure 2 (b), it can be noted that the server inlet temperature is consistently about 10 °F higher than the supply temperature and the discrete sensors are 2 °F higher than the supply air temperature. Figure 2 (b) also shows the inlet temperature of the Dell PowerEdge 2950 -50 server at the bottom of rack C1-4, which reported the lowest IPMI inlet temperature for all SAT set points. In this case, however, the figure shows that the Dell PowerEdge 2950 -50 server is consistently about 2 °F higher than the SAT set point. These results are counterintuitive given that all servers are in a contained cold aisle, with the containment maintained at neutral pressure to make sure sufficient airflow is supplied to all servers.

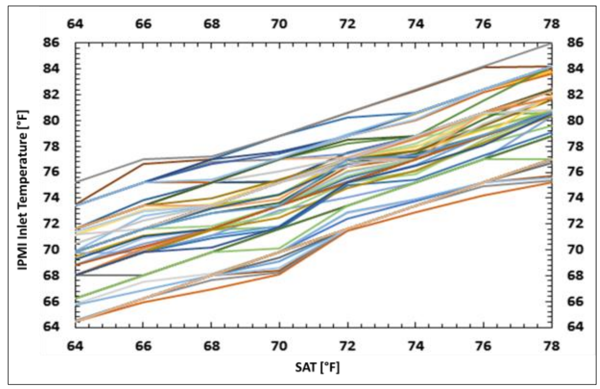

(a)

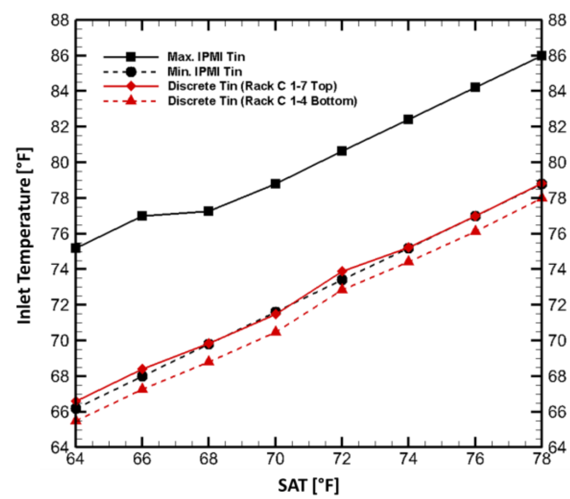

(b)

Figure 2. (a) IPMI inlet temperature for all servers in Aisle C. (b) IPMI and discrete inlet temperature vs SAT of both Dell PowerEdge 2950 -50 and 112 servers.

In further analyzing the data provided by the IPMI protocol, given the critical nature of the CPU we first focus on the server CPU temperature and fan speed. Figure 3 (a) shows a box plot of the average CPU temperature of all servers installed in Aisle C for varying SAT set points at neutral pressure. It can be noted that the CPU temperature of all servers never reached or exceeded the CPU temperature threshold of 190 °F [25].

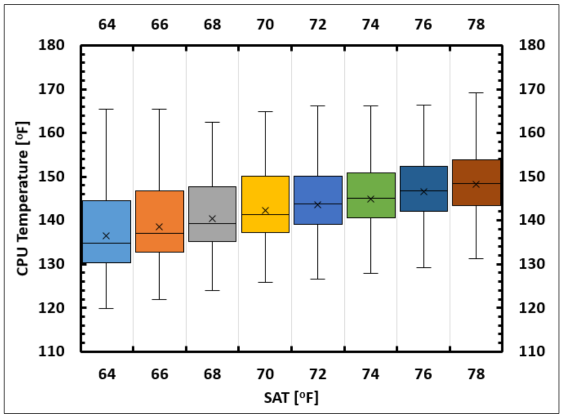

(a)

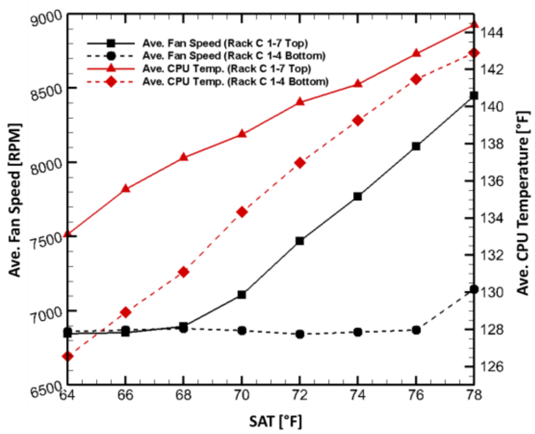

(b)

Figure 3. (a) CPU temperature variation. (b) Ave. fan speed and CPU temperature vs SAT.

In looking closely at how the server fan speed behaves with CPU temperature, Figure 3 (b) shows the behavior of the average CPU temperature vs the average fan speed of the Dell PowerEdge 2950's-112 and 50. It can be noted that the Dell PowerEdge 2950-112 fan speed begins to increase at a CPU temperature of about 138 °F, while the Dell PowerEdge 2950-50 server fan speed begins to increase at a CPU temperature of about 142 °F. Since the fan speed is not uniquely controlled by CPU temperature, the next step was to analyze the behavior of the Dell PowerEdge 2950 server fan speed vs the server inlet temperature. Figure 4 shows the variation in the Dell PowerEdge 2950 IPMI inlet temperature vs fan speed for 90 of the 92 Dell PowerEdge 2950 servers installed in Aisle C. Two of the servers were giving erroneous results.

The figure clearly shows the relationship between the two variables, where the Dell PowerEdge 2950 server fan speed always starts to increase at an IPMI inlet temperature of 77 °F and beyond. While this was particularly true for the Dell PowerEdge 2950 servers, it was not the case for the other two classes of servers installed in Aisle C.

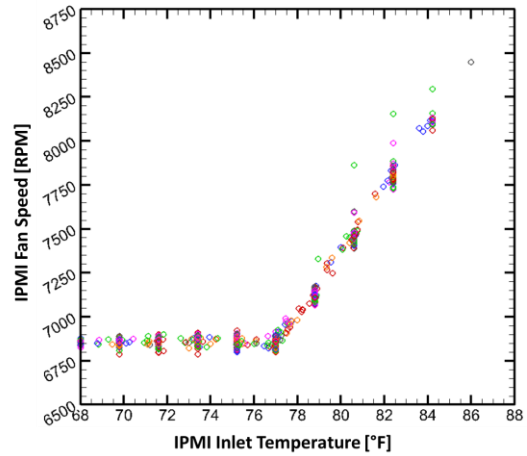

Figure 4. IPMI fan speed vs IPMI inlet temperature for the D2950 servers.

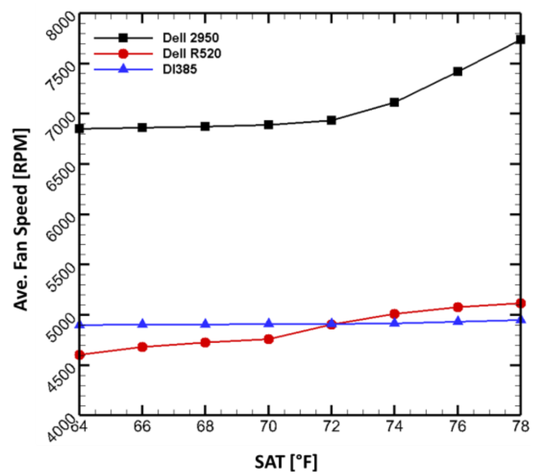

Figure 5. Average fan speed vs SAT for the three server classes.

Figure 5 shows the variation of average fan speed for the three classes of servers in Aisle C, at different SAT set points. The HP ProLiant DL385 servers maintained the same average fan speed for all SAT set points tested. The Dell PowerEdge R520 servers' fans started ramping up slightly after a SAT of 70 °F by about 500 RPM (approximately 4% of increase in fan speed). Finally, the Dell PowerEdge 2950 servers' fans speed started ramping up beyond a SAT of 72 °F by about 1000 RPM (approximately 8% of increase in fan speed). These results clearly indicate how different servers may behave despite having the same temperature and pressure environment. It is therefore critical for a data center operator to understand the behavior of their installed IT equipment before initiating a change in the data center control set points as the IT control algorithms may not be able to respond sufficiently to compensate. For example if the cooling system is not supplying sufficient cold air, then increasing the servers fan speed may not be sufficient to resolve the cooling of the server. The cooling energy savings that are realized from raising temperature set points or lowering cooling unit fan speeds may come at the cost of raising the IT equipment power consumption, mainly due to increase in server fan speed. The following sections discuss this concept in further detail.

IPMI VS DISCRETE DATA

This section focuses on comparing the two measurement methods used in this study and analyzing the validity of using either method for monitoring data center thermal performance. Figure 6 (a) displays the data distribution for both the IPMI inlet temperature measurements and the discrete sensor inlet temperature measurements at neutral containment pressure and for various SAT set points. The figure clearly shows the consistently large variation in the IPMI inlet temperature measurements for all SAT set points, reaching a delta as high as 11 °F between the minimum and maximum IPMI server inlet temperatures. On the other hand, the discrete sensor measurements are consistently close in value across Aisle C with a maximum delta of 2 °F between the minimum and maximum inlet temperature sensor measurements. Given that the infrastructure controls are designed to provide air at uniform temperature and at sufficient volume flow rate we expect a uniform inlet temperature. It would be dangerous to use the IPMI sensors for cooling infrastructure control without greater understanding of the IPMI measurements as they report significantly different temperatures.

Figure 6 (b) displays a statistical comparison between the two monitoring methods, showing both the standard deviation and mean temperature values of all IPMI inlet temperatures and discrete sensor measurements at various SAT set points. The IPMI inlet temperature standard deviation and mean values are consistently higher than the discrete sensor measurements. Both Figures 6 (a) and 6 (b) illustrate the difficulty of using the server IPMI inlet temperature measurements in understanding the true thermal environment around the servers in Aisle C. For data center operators looking to optimize on their cooling control for Aisle C, it is difficult to establish what IPMI inlet temperature reading to use to control the cooling infrastructure.

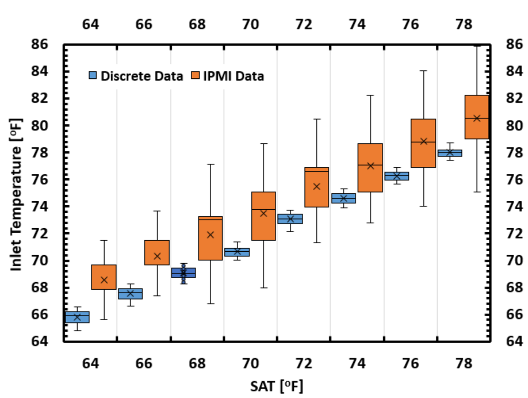

(a)

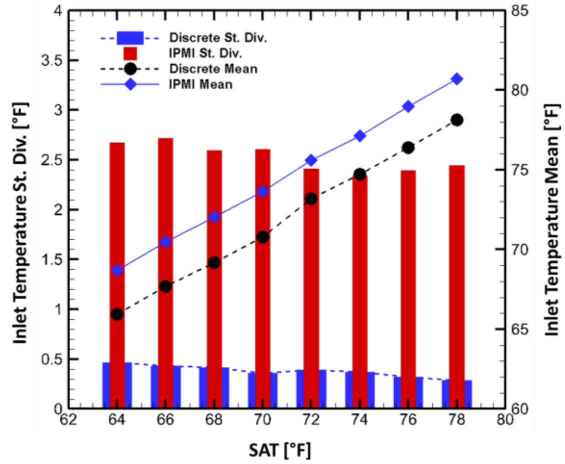

(b)

Figure 6. (a) IPMI vs discrete inlet temperature variation. (b) Standard deviation and mean values of IPMI and discrete inlet temperature.

COOLING CONTROL

As illustrated in the previous section, the IPMI inlet temperature data was highly variable at all SAT set points. To understand the impact of using this data for cooling control, it is essential to look at the cost of cooling Aisle C at the various SAT set points. Figure 7 (a) displays the annual costs of operating the chiller and the CRAH blowers for varying SAT set points at an energy cost of 10 ¢/kWh. The CRAH blower power was measured while the chiller power (W_{net}) was calculated using equation (1). [13].

$$W_{net} = \frac{Q_{removed}}{COP_R} \qquad (1)$$

The figure also shows the total annual cooling cost which is calculated as the sum of the chiller operating cost and the blowers operating cost. As expected, as the SAT set point increases, the chiller operating costs decrease. However, the CRAH blower operating costs increase slightly beyond a SAT of 72 °F. This occurred due to the increase in server fan speeds as displayed in Figure 5, which caused the CRAH blower speed to increase to maintain the neutral pressure set point inside Aisle C. Despite the increase in blower operating cost, the total annual cooling cost for Aisle C always decreased with increasing SAT set points. Looking back at Figure 6 (a), if the data center operator takes the conservative approach and controls the CRAH SAT set point based on a target maximum server inlet temperature of 75 °F. The IPMI inlet temperature measurements would require a SAT set point of 64 °F, while the discrete sensor measurements would require a SAT set point of 72 °F. This difference in SAT set point results in a difference in total cooling costs of $5,745, which is approximately a 20% reduction in cooling cost.

A similar analysis can be conducted for higher server inlet temperature set points; however, the increase in server fan speed must be accounted for to quantify the true savings from set point adjustments [26]. For example, to maintain a target server inlet temperature of 80 °F, the IPMI inlet temperature measurements would require a SAT set point of 72 °F, while the discrete sensor measurements would require a SAT set point of 78 °F. The total cooling cost difference between the two set points is $2,030, which is approximately a 9% reduction in cooling cost. However, the servers fan power in Aisle C increased due to higher SAT set points as shown in Figure 7 (b). More specifically, the difference in server fan power annual cost between a SAT set point of 72 °F and 78 °F is $703. Therefore, the true energy savings between a SAT set point of 78 °F vs 72 °F are $1,327.

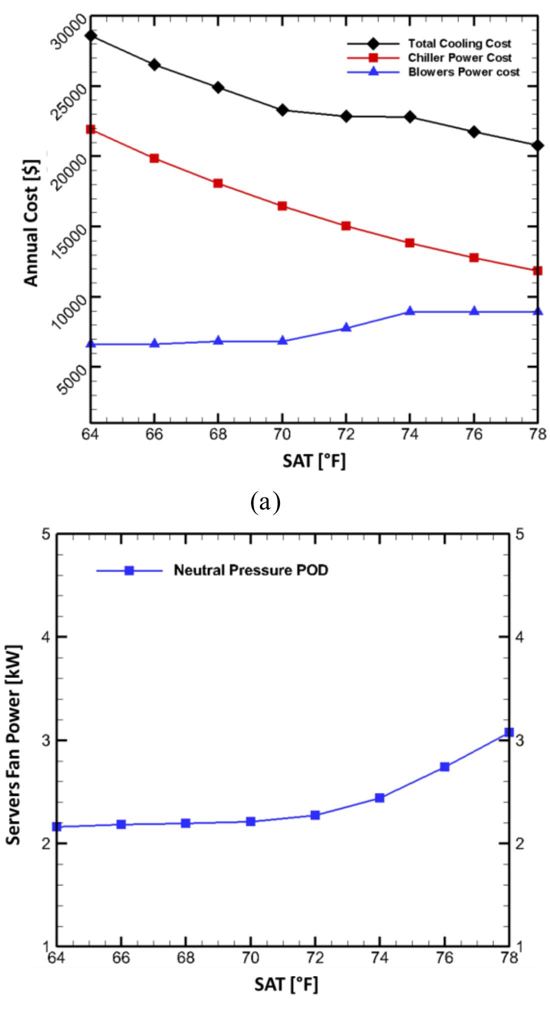

(b)

Figure 7. (a) Total cooling cost. (b) Total Aisle C servers' fan power.

The increase in server fan power (calculated using fan laws from the IPMI reported fan speed as described later) observed in Figure 7 (b) is not only a function of SAT, but also CAC pressure. This is discussed in further details in the following section.

IMPACT OF CAC PROVISIONING

The test data presented so far has focused on maintaining neutral pressure inside the CAC to ensure that sufficient airflow is supplied to the IT equipment. In most cases however, data center facilities deploying CAC will rarely monitor pressure inside the containment. Therefore, most data center operators do not necessarily know whether they are supplying sufficient air to their IT equipment or not. In this section, we consider the case of an under-provisioned CAC at $-0.01''\ H_2O$ of CAC pressure, at SAT set points of 64, 72, 74, 76, and 78 °F.

Figure 8 (a) displays the data distribution for both the IPMI inlet temperature measurements and the discrete sensor inlet temperature measurements at $-0.01''\ H_2O$ of containment pressure and for various SAT set points. Maintaining a lower containment pressure by undersupplying airflow makes matters worse. The variations in temperature measurements get larger which makes it more difficult to interpret those measurements. The figure clearly shows the consistently large variation in the IPMI inlet temperature measurements for all SAT set points, similar to the neutral pressure case, however the delta between the minimum and maximum IPMI server inlet temperatures reaches 21 °F compared to 11 °F for neutral pressure. With the current state of understanding of the IPMI data, it is difficult to recommend a cooling control strategy based on IPMI sensor data for the specific environment in aisle C. Also in Figure 8 (a), the discrete sensor measurements are showing larger variation when compared to the neutral pressure case, where the delta between the minimum and maximum inlet temperature sensor measurements reaches 6 °F, versus 2 °F in the neutral pressure case. This data highlights the possibility of using the variation in discrete sensor data to gauge how undersupplied a containment is, where the temperature variation can be expected to increase as the CAC is further undersupplied, and will plateau as the CAC is oversupplied.

Figure 8 (b) displays the data distribution for both the IPMI inlet measurement and the discrete sensors inlet temperature at -0.01 , 0.0 and $0.01''\ H_2O$ of containment pressure and for SAT value of 64 °F. The discrete sensors show the expected behavior. This is, the temperature variation decreases as the CAC is pressurized. In contrast, the behavior indicated by the IPMI sensors was not as expected. The temperature variation decreases between the under-provisioned and neutrally provisioned cases then increases with over-provisioning. It is also worthy to note that the minimum IPMI inlet temperature in the over-provisioned case was 62 °F while the supply air temperature was 64 °F. This perhaps reflects the lower quality sensors that can be used in mass production.

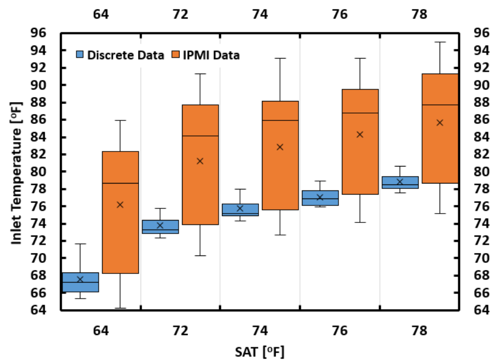

(a)

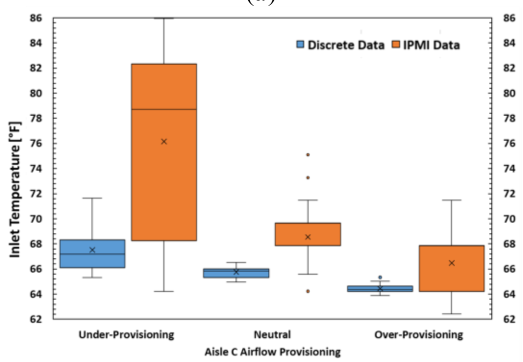

(b)

Figure 8. IPMI vs discrete inlet temperature variation (a) Under-provisioning. (b) Different provisioning cases for SAT value of 64 °F.

Figure 9 displays the effect of under-provisioning the CAC on total server fan power in Aisle C and compares the results to the neutral pressure case. The results clearly show that even at 64 °F of SAT, starving the IT servers of cool air causes their fans to spin up and draw an extra 0.8 kW. As the SAT increases, the total server fan power also increases reaching 5 kW at 78 °F SAT for the under-provisioned case, which accounts for 10% of the total IT equipment power. At neutral pressure and 78 °F SAT, the total server fan power accounts for 6% of the total IT equipment power. These results clearly highlight the effect of an under-provisioned CAC on the IT equipment environment and performance.

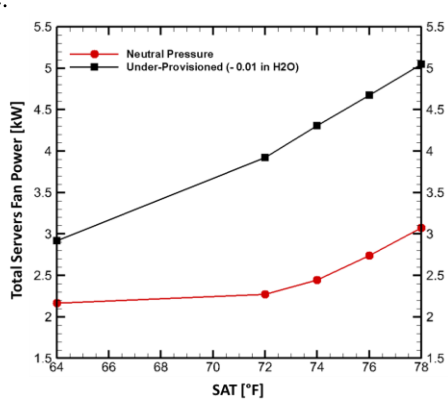

Figure 9. Server's fan power.

In further analyzing the effect of CAC under-provisioning we compared the under-provisioned and the neutral pressure scenarios. Figure 10 (a) displays the temperature gradient of Aisle C as given by SynapSoft™ at $-0.01''\ H_2O$ of containment pressure and for SAT set point of 74 °F. A closer look at the CAC showed that the temperature gradient is non-uniform in the case of under-provisioning. This is due to the recirculation/reverse flow of air from the hot aisle into the cold aisle drawn by negative pressure. In comparison, Figure 10 (b) presents the temperature gradient at neutral or very slightly positive containment pressure at the same SAT value. The positive pressurization minimizes recirculation that results in more uniform temperatures.

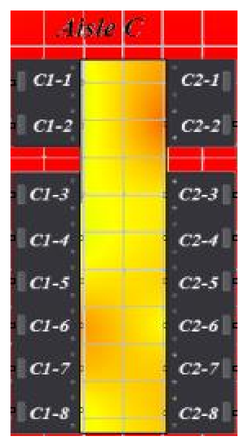

(a)

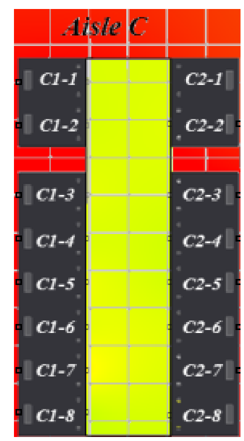

(b)

Temperature [°F]

Figure 10. Aisle C temperature gradient provided by SynapSoft™ Software (a) Under-provisioning. (b) Neutral pressure.

A data center operator that chooses to lower cooling equipment fan speed may gain cooling energy savings, at the cost of raising the IT equipment power consumption. This raises the question of whether common data center metrics used today can capture the effect of set point changes on server fans, and therefore it is discussed in further detail in the following section.

PUE AND tPUE

One of the most common efficiency metrics for DC efficiency is Power Usage Effectiveness (PUE). In simple terms, PUE is the ratio of power delivered to the facility divided by the power delivered to servers, storage, and networking gear. The PUE represented in equation (2), represents a ratio between the total facility power (total IT, chiller, blower and lighting power) and the total IT equipment power.

$$PUE = \frac{IT\ power + Blowers + Chiller + Lighting}{IT\ Equipment\ Power} \quad (2)$$

One of the issues with the PUE metric is the inclusion of server fan power in the total IT equipment power. As shown in the previous section, while raising the cooling equipment set point saves energy at the cooling infrastructure level it may in some cases cause an increase in server fan power, thereby increasing the total IT equipment power. The PUE metric however would not capture this effect and would always show a lower PUE number. Researchers have considered introducing tPUE (Total Power Usage Effectiveness) as a new metric which would capture the effect of non-useful server power changes by replacing the IT Equipment Power in the denominator of equation (2) with Productive IT Equipment Power [27]. This metric is presented below in equation (3).

$$tPUE = \frac{IT\ power + Blowers + Chiller + Lighting}{Productive\ IT\ Equipment\ Power} \quad (3)$$

The issue with the tPUE however is how to define the productive IT equipment power. For the sake of this study, we will assume that the productive IT equipment power is defined as:

$$Productive\ IT\ Equipment\ Power = Total\ IT\ Equipment\ Power - Server\ Fan\ Power \quad (4)$$

Where the server fan power is calculated using server fan curves and the associated fan law:

$$\left(\frac{P_1}{P_2}\right) = \left(\frac{N_1}{N_2}\right)^3 \quad (5)$$

Figure 11 shows how the PUE and tPUE metrics change with different SAT set points while comparing the two cases of neutral and under-provisioned CAC pressure. The first observation to make is that tPUE is consistently higher than PUE because of excluding the total server fan power from the tPUE calculation. The second observation to make is that whether PUE or tPUE are used, both metrics suggest efficiency gains from under-provisioning the CAC, mainly due to the lower cooling fan costs. However, the main observation to highlight is the gradient of the PUE and tPUE plots, especially for the under-provisioned case. It can be noted that the tPUE gradient starts to plateau beyond a SAT of 72 °F, suggesting no gain in efficiency from lowering the SAT while the PUE gradient indicates continues efficiency gains beyond 72 °F SAT. This shows the value of using the tPUE metric which looks at the productive IT equipment power and eliminates the bias in the PUE metric introduced by including the server fan power effects. Also by quantitatively comparing the gains of under-provisioning the CAC, at 78 °F SAT the PUE metric is reduced by 5.6% while the tPUE metric is reduced by only 1.8%. It has been seen that temperatures sensed internal to the IT equipment by IPMI sensors are more dramatically affected by under-provisioning. Until the significance of these variations is investigated, it should be understood that choosing energy saving by under-provisioning might impact the reliability of IT equipment.

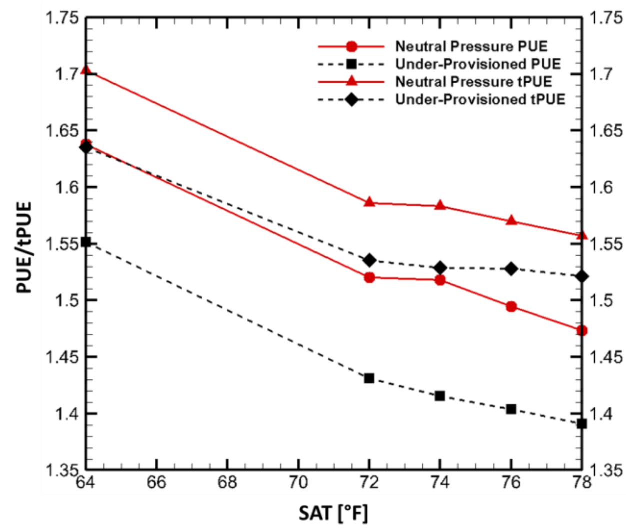

Figure 11. PUE/tPUE comparison.

CONCLUDING REMARKS

This study presented an experimentally based analysis and comparison between two different approaches of environmental data collection; one using a discrete sensor network, and another using available data from installed IT equipment through their IPMI protocol. Despite maintaining neutral pressure inside the contained cold aisle, the IPMI inlet temperature data showed very large variation between the servers with the variation reaching as high as 11 °F. On the other hand, discrete sensor data provided more reliable and consistent results with a measurement variation inside the cold aisle as low as 2 °F which makes discrete sensors a more appropriate to control the cooling air temperature available in the cold aisle for the IT equipment. Given these differences, and without a clear understanding of the significance of the IPMI variations, the study focused on how industry common methods for cooling efficiency management and control can be affected by the data collection approach. It was shown that using IPMI inlet temperature data might lead to unnecessarily lower cooling set points, which in turn minimizes the potential cooling energy savings. It was shown in one case that using discrete sensor data for control provides 20% more energy savings than when using IPMI inlet temperature data.

The behavior of the three classes of servers installed in the tested aisle was evaluated by looking at the behavior of their fan speed with changing SAT set points. The results clearly indicated how different servers may behave despite having the same temperature and pressure environment. For instance, the HP ProLiant DL385 servers reported the same average fan speed for all SAT set points tested at neutral containment pressure. However, the Dell PowerEdge 2950 servers' average fan speed ramped up by 8% between a SAT of 72 °F and 78 °F, at neutral containment pressure. These results highlight how important it is for data center operators to understand the mix of IT equipment installed before initiating a change in the data center control set points.

Finally, the PUE metric was used to evaluate the effect of set point changes on the overall efficiency of the data center. The results demonstrated how the use of the PUE metric can sometimes provide misleading results, and in many cases, overestimate the efficiency of a data center. This was attributed mainly to the effect of server fan speed, and how an increase in server fan power provides a better PUE metric. Therefore, the study evaluated the use of the tPUE metric, which considers the effect of server fan speed on the overall data center efficiency. The results showed that the tPUE metric provides a more reliable energy efficiency trend for varying set point changes.

This paper aims to provide empirical insight to data center owner/operator as well as IT manufacturer on how to monitor the data center environment and IT cooling status. The following guidelines and observations can be inferred from this study:

1) Discrete IT inlet temperature (also referred to as external or infrastructural) sensors are vital to monitor the data center cooling performance and health.
2) Accuracy of IPMI IT inlet temperature (also referred to as internal) can vary. This depends on the server generation and proximity of sensor to active heat dissipating components.
3) Ideally, the reading of the IPMI sensor and the discrete should be close at cases of over and neutral provisioned contained aisles. The difference between IPMI and discrete increases due to global or local containment under-provisioning. When this happens, the IPMI reading increases since hot air inside the server move closer to the frontal area of the chassis. This also shows the importance of infrastructural pressure instrumentation in containment in understanding IPMI outputs.
4) For cases when IPMI sensors are preheated, choosing the cooling setpoint based on the discrete sensors can decrease PUE and better the data center energy savings.
5) For cases of pressure imbalances in contained aisles, IPMI are an effective early alarm tool for the server internal components temperature. The difference between discrete and IPMI inlet temperature can be used as an indicator as well. It is important to note here that the IPMI sensors (IT inlet, CPU...) play an important role in controlling the server's fan speed.
6) Discrete sensors are a necessary tool to establish calibrated computational fluid dynamics models for data center trouble shooting and capacity planning.

ACKNOWLEDGMENTS

We would like to acknowledge, Panduit Corporation, Future Facilities Ltd. and Vertiv. We would also like to thank the ES2 Partner Universities for their support and advice. This work is supported by NSF IUCRC Award No. IIP-1134867 and MRI Award No. CNS1040666.

REFERENCES

[1] A. Shehabi et al., "United States Data Center Energy Usage Report," 2016.

[2] Huigui Rong, Haomin Zhang , Sheng Xiao , Canbing Li, Chunhua Hu," Optimizing energy consumption for data centers," Renewable and Sustainable Energy Reviews, vol. 58, pp. 674-691, 2016.

[3] Morgan Tatchell-Evans, Nik Kapur, Jonathan Summers, Harvey Thompson, Dan Oldham," An experimental and theoretical investigation of the extent of bypass air within data centres employing aisle containment, and its impact on power consumption," Applied Energy, vol. 186, pp. 457-469, 2017.

[4] D. Garday and D. Costello, "Air-Cooled High-Performance Data Centers: Case Studies and Best Methods", Intel Corporation, White Paper, November 2006.

[5] J. Niemann, K. Brown and V. Avelar, "Impact of Hot and Cold Aisle Containment on Data Center Temperature and Efficiency", Schneider Electric Data Center Science Center, White Paper 135 Rev. 2, pp. 1-14, October 2011.

[6] R. Zhou, Z. Wang, C.E. Bash, and A. McReynolds, "Modeling and control for cooling management of data centers with hot aisle containment", Proceedings of ASME 2011International Mechanical Engineering Congress & Exposition (IMECE 2011), Denver, Colorado, USA, November 11-17.

[7] M. Martin, M. Khattar, and M. Germagian, "High-Density Heat Containment", ASHRAE Journal, vol. 49, no. 12, pp. 38-43, December 2007.

[8] "Data Center 2020: Delivering high density in the Data Center; efficiently and reliably", Intel Corporation, White Paper DataCenter 2020, 2011.

[9] https://journal.uptimeinstitute.com/2014-data-center-industry-survey/.

[10] Shrivastava, S. K., Calder, A. R., B. G., Ibrahim, M.," Quantitative Comparison of Air Containment Systems," ITHERM, San Diego, CA, USA, 2012.

[11] Patterson, M. K., Martin, R., von Oehsen, J., Pepin, J., Joshi, Y., Arghode, V. K., Steinbrecher, R., King, J., "A Field Investigation Into The Limits Of High-Density Air-Cooling," InterPack 13, Burlingame, CA, USA, 2013.

[12] Shrivastava, S. K., Ibrahim, M., " Benefit Of Cold Aisle Containment During Cooling Failure," InterPack 13, Burlingame, CA, USA, 2013.

[13] H. A. Alissa, K. Nemati, B. Sammakia, M. Seymour, R. Tipton, D. Mendo, D. W. Demetriou, K. Schneebeli,"Chip to Chiller Experimental Cooling Failure Analysis of Data Centers: The Interaction Between IT and Facility," in IEEE Transactions on Components, Packaging and Manufacturing Technology, vol. 6, no. 9, pp. 1361-1378, Sept. 2016.

[14] Makwana, Y. U., Calder, A. R., Shrivastava, S. K., "Benefits of Properly Sealing a Cold Aisle Containment System," ITHERM, Orlando, FL, USA, 2014.

[15] Sundaralingam, V., Arghode, V. K., Joshi, U., " Experimental Characterization of Cold Aisle Containment for Data Centers," SEMI-THERM 29, San Jose, CA, USA, 2013.

[16] B. Muralidharan, Saurabh K. Shrivastava, Mahmoud Ibrahim, Sami A. Alkharabsheh, B. G. Sammakia, "Impact of Cold Aisle Containment on Thermal Performance of Data

Center," Proceedings of the ASME 2013 International Technical Conference and Exhibition on Packaging and Integration of Electronic and Photonic Microsystems InterPACK2013, July 16-18, 2013, Burlingame, CA, USA, IPACK2013-73201.

[17] " Impact of Air Containment Systems Reducing Energy Consumption in the Data Center", Panduit Corporation, White Paper, June 2012.

[18] Chao Liu, "Using real-time data to improve efficiency and reduce total cost of ownership", Intel Corporation, White Paper, 2014.

[19] Patterson M. K., Weidmann R., Lebercht M., Mair M., Libby R. M., "An Investigation Into Cooling System Control Strategies for Data Center Airflow Containment Architectures," Proceedings of the ASME 2011 International Technical Conference and Exhibition on Packaging and Integration of Electronic and Photonic Microsystems InterPACK2011, July 6-8, 2011, Portland, Oregon, USA, IPACK2011-52090.

[20] Nishi A., Rego Ch. W., Sandeep A., Zhou Sh., Shrivastava, S., "Real time monitoring and availability of server airflow for efficient data center cooling," Semiconductor Thermal Measurement and Management Symposium (SEMI-THERM), 2013 29th Annual IEEE.

[21] Alissa, H.A., Nemati, K., Sammakia, B.G., Schneebeli, K., Schmidt, R.R. and Seymour, M.J., 2016. Chip to facility ramifications of containment solution on it airflow and uptime. IEEE Transactions on Components, Packaging and Manufacturing Technology, 6(1), pp.67-78.

[22] Tradat, M.I., Alissa, H.A., Nemati, K., Khalili, S., Sammakia, B.G., Seymour, M.J. and Tipton, R., 2017, March. Impact of elevated temperature on data center operation based on internal and external IT instrumentation. In Thermal Measurement, Modeling & Management Symposium (SEMI-THERM), 2017 33rd (pp. 108-114). IEEE.

[23] H. A. Alissa, K. Nemati, B. Sammakia, K. Ghose, M. Seymour, and R. Schmidt, "Innovative approaches of experimentally guided CFD modeling for data centers," Thermal Measurement, Modeling & Management Symposium (SEMI-THERM), 2015 31st. pp. 176–184, 2015.

[24] Alissa, H. A., Nemati, K., Sammakia, B. G., Wu, T., & Seymour, M. J. (2016, March). Management and predictions of operational changes and growth in mission critical facilities. In Thermal Measurement, Modeling & Management Symposium (SEMI-THERM), 2016 32nd (pp. 178-184). IEEE.

[25] http://www.hardwaresecrets.com/maximum-CPU-temperature/5/.

[26] Patterson M. K.," The Effect of Data Center Temperature on Energy Efficiency," in IEEE, 2008.

[27] Hamilton, J. (2009, July). PUE and total power usage efficiency (tPUE) [Blog Post]. Retrieved from http://perspectives.mvdirona.com/2009/06/pue-and-total-power-usage-efficiency-tpue/

