Impact of Static Pressure Differential Between Supply Air and Return Exhaust on Server Level Performance

Ashwin Siddarth, Richard Eiland, John Edward Fernandes and Dereje Agonafer
Mechanical and Aerospace Engineering,
P.O. Box 19023
University of Texas Arlington,
Arlington, TX, USA, 76010
Email: ashwin.siddarth@mavs.uta.edu

ABSTRACT

Modern Information Technology (IT) servers are typically assumed to operate in quiescent conditions with almost zero static pressure differentials between inlet and exhaust. However, when operating in a data center containment system the IT equipment thermal status is a strong function of the non- homogenous environment of the air space, IT utilization workloads and the overall facility cooling system design. To implement a dynamic and interfaced cooling solution, the interdependencies of variabilities between the chassis, rack and room level must be determined. In this paper, the effect of positive as well as negative static pressure differential between inlet and outlet of servers on thermal performance, fan control schemes, the direction of air flow through the servers as well as fan energy consumption within a server is observed at the chassis level.

In this study, a web server with internal air-flow paths segregated into two separate streams, each having dedicated fan/group of fans within the chassis, is operated over a range of static pressure differential across the server. Experiments were conducted to observe the steady-state temperatures of CPUs and fan power consumption. Furthermore, the server fan speed control scheme's transient response to a typical peak in IT computational workload while operating at negative pressure differentials across the server is reported. The effects of the internal air flow paths within the chassis is studied through experimental testing and simulations for flow visualization.

The results indicate that at higher positive differential pressures across the server, increasing server fans speeds will have minimal impact on the cooling of the system. On the contrary, at lower, negative differential pressure server fan power becomes strongly dependent on operating pressure differential. More importantly, it is shown that an imbalance of flow impedances in internal airflow paths and fan control logic can onset recirculation of exhaust air within the server. For accurate prediction of airflow in cases where negative pressure differential exists, this study proposes an extended fan performance curve instead of a regular fan performance curve to be applied as a fan boundary condition for Computational Fluid Dynamics simulations.

KEY WORDS: Fan performance curve, Server Design, Data Center Cooling

INTRODUCTION

In pursuit of extending the air cooling technology in data centers, the industry is constantly looking for innovations that augment the efficiencies and cooling capacities per rack. Numerous studies have contributed in optimizing airflow provisioning at the rack focusing on facility room architecture and data center air distribution [1][2]. The advent of aisle containment has presented an easy solution to implement and its advantages of airflow separation, improved efficiency and thermal ride-through times are well documented [3]– [5]. High density facility design for rack densities up to 30 kW cabinets in [6] proposed the idea of targeted air supply with velocity less than 400 feet per minute. By maintaining a static pressure differential in a cold aisle and hot aisle configuration the server fans speeds were minimized to improve overall efficiency in an Open Compute data center [7].

Previous studies have worked to understand the interactions of under-floor pressure distributions on airflow rates through perforated floor tiles and its effect on air distribution and temperatures at the rack inlet. Common goals for underfloor pressure management are to reduce the velocity of air flow and hence increase the static pressure of the underfloor. This in turn increases the volume of air passing through perforated floor tiles. In an idealized scenario, the entirety of conditioned air provided by a CRAC unit would travel through IT equipment. In actuality, bypass and leakage of cold air around the server and into the hot aisle decreases the efficiency of the room level cooling system. Additionally, recirculation of hot exhaust air back to the cold aisle raises the inlet temperature to servers, also diminishing cooling efficiency.

Most of the previous studies related to data hall pressurization stop at the rack level focusing on inlet temperatures, rack airflow rates and aisle level leakage flow rates [3], [4], [8], [9]. Alissa et al. developed a generalized testing methodology for IT equipment wherein the fan curve, the system impedance and external impedance effects can be quantified with a single curve (i.e. flow curve) that is governed by fan affinity laws. Utilizing the flow curves as numerical boundary condition, it was shown that actual volume flow rates through IT equipment in a cold aisle containment depends upon its form factor, chassis fan array performance degradation and inherent sensitivity to external impedances [10].

This work seeks to understand the impact of pressurization on the thermal performance, fan control scheme's response, directionality of internal airflow paths and cooling energy consumption within a server.

EXPERIMENTAL SETUP

Figure 1 shows a first-generation Open Compute server, based on Intel motherboard, used in this experiment. The key design features of these servers include larger chassis height of 1.5U rack unit and customized heat sinks for wider spaced primary heat generating components (CPUs).

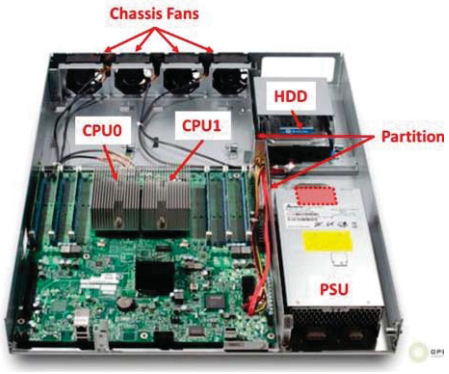

Figure 1: First generation Open Compute server with partitioned airflow paths

The airflow path is compartmentalized with a partitioning for power supply and hard drive as shown in Figure 1. The compartmentalization allowed for discrete analysis of airflow over primary components and apparent interaction of the two airflow paths. In the motherboard (MB) portion, four 60mm x 25mm axial fans are located at the exhaust of the server.

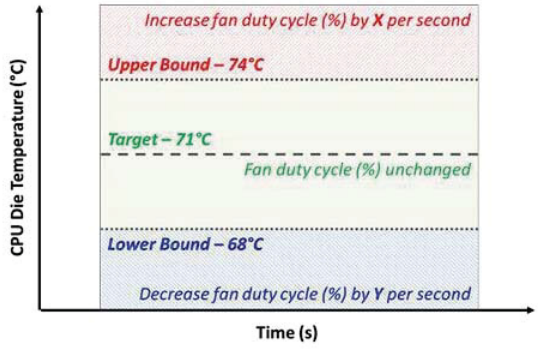

Figure 2: Dead band control limits utilized in the fan speed control algorithm for Open Compute server under study

The two processors represent the principal heat load within the system with a thermal design power of 95W and their temperatures drive the fan speed control algorithm to ensure adequate cooling. The power supply unit (PSU) has an integrated 60mm fan that modulates the airflow based on inlet temperatures, in the PSU portion, and it idles at 30% duty cycle for inlet temperatures below 25□C [11].

All the 60mm DC fans are powered with an external 12VDC power supply that enables measurement of fan power for the rack. The fan speeds are logged using a data acquisition unit sensing off the Hall effect sensor from each fan.

The fan array is controlled through an external pulse width modulation (PWM) signal generator or through the internal motherboard native control algorithm based on a dead band control scheme as shown in Fig 2 [12], [13].

Table 1: Fan speed at various PWM signals for the 60mm axial fan tested

Experimental Results – Fan speed	
PWM Duty Cycle (%)	Fan Speed (rpm)
0	2108
5	2365
10	2721
25	3769
50	5427
75	6839
100	8022

In an operating data center, the airflow provisioning configuration modulates the temperature and pressure variations across the face of the rack. An Air Flow Bench test chamber was used to act as a pseudo cold aisle containment system shown schematically in Figure 3. This is an idealized scenario because it assumes uniform pressure across the face of the servers and flow straighteners within the chamber orient streamlines parallel to the servers. A stack consisting of four Open Compute servers was used to simulate a rack. The servers are named 10 through 40 starting from the bottom of the stack.

The static pressure differential across the servers is measured utilizing the pressure taps located at the inlet of the server. These pressure taps are connected to the unidirectional pressure transducers on the airflow bench and the flow rate and pressure readings are logged using a data acquisition software. Omega OM-EL-USB-1-LCD temperature loggers are used to measure the room ambient temperature and a rack inlet temperature of 24°C ± 1.0°C is observed during testing. To communicate with the four servers operating on Linux operating interface, a desktop workstation is utilized to execute bash command scripts and transfer the data logged on the servers with synchronized time stamps.

When a single server is considered for transient thermal performance measurements, additional thermocouples are installed to measure surface temperature of hard disk drive (HDD), CPU voltage regulators (VRD) and dual-inline memory modules (DIMM). The thermocouples are of T-type and have an accuracy of ±0.5°C. For power measurements, a Yokogawa CW121 power meter is used to measure the current drawn by a server operating at 277VAC.

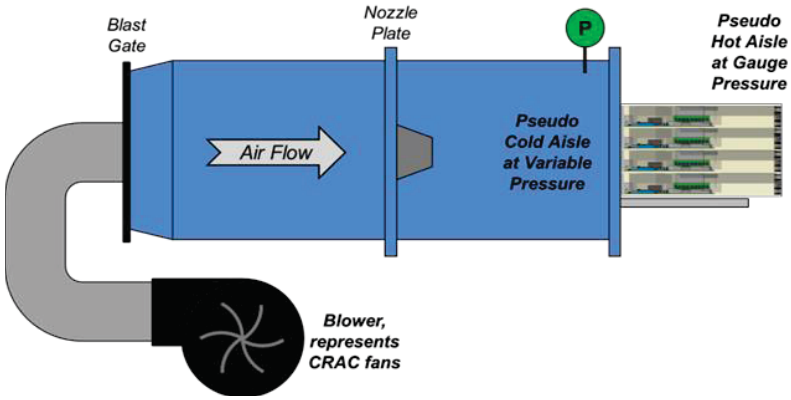

Figure 3: Schematic of experimental setup used to control inlet static pressure to server

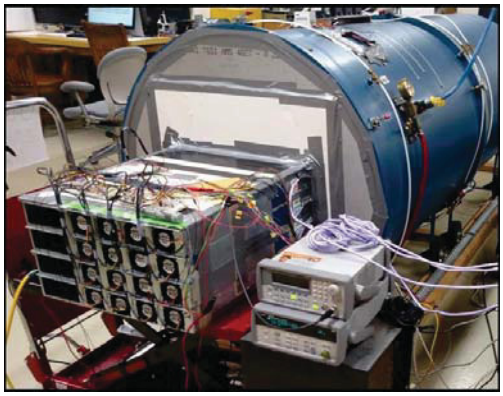

Figure 4: Representative rack attached to the Airflow Test Bench

TESTING PROCEDURE

Steady-state parametric study: For the initial testing, a synthetic computational load is applied to all four servers using a free software package, *lookbusy*. To understand the impact of inlet static pressurization on the thermal performance of the servers, the server fans were first fixed at a constant duty cycle. In this externally controlled and externally powered fan configuration, the effect of pressure on CPU temperatures could be isolated from other fan interactions. While operating at fixed fan speeds, the blower to the Air Flow Bench was adjusted to achieve a desired inlet static pressure. At each server fan speed and blower setting, the server was provided a synthetic computational workload of idle, 30%, and 98% CPU utilization with the *lookbusy* software tool. Each workload was run for 30 minutes and repeated three times in total for repeatability. The results gathered here are taken from the average values over the last 10 minutes of each computational workload as this is when steady state CPU temperatures were achieved.

Following the externally powered, externally controlled test, the server fans were then controlled internally by the motherboard's native fan speed control algorithm while still being powered externally. Again, the Air Flow Bench was set to fixed blower speeds, but now, the inlet static pressure became a function of the server fan speeds as they modulated to maintain the control set point temperature of the dead band operation described in Figure 2. Tests were conducted with the servers operating at 98% CPU utilization only because this was the only workload that generated CPU temperature high enough to trigger the fan speed control. The workload was operated for 60 minutes in this case to allow the fluctuations in fan speed and static pressure to reach steady state conditions. Average values over the last ten minutes of steady state conditions are reported here.

Transient response to peak CPU utilization: To understand the impact of pressurization on the server's fan speed control scheme, a single server is tested initially on a table top to establish a baseline. The 'Baseline' arrangement represents the fully-provisioned scenario where the pressure differential across the server and hence the external impedance is zero.

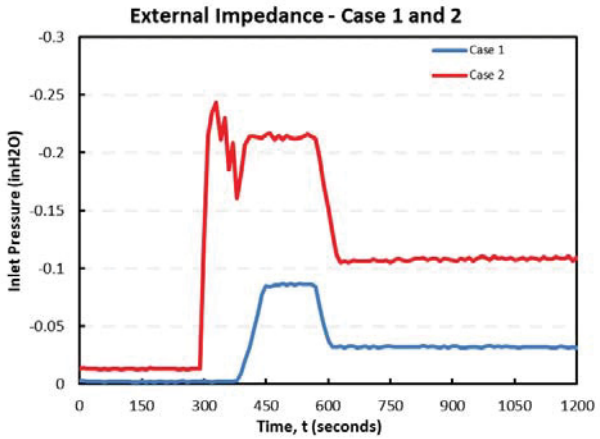

Figure 5: Variation in inlet static pressure when the CPU utilization peaks at 300s for Cases 1 and 2

Subsequently, the server is mounted on the Air Flow Test Bench and additional external impedance is created utilizing different sized nozzle plate opening. Two such arrangements provide the under-provisioned cases 1 and 2. Quantitative values for impedances associated with cases 1 and 2 is given in Figure 5. Cases 1 and 2 represent the under-provisioned scenario with the server deprived of required supply air at the inlet.

The fan control scheme's response to an increase in CPU utilization from idle to 98% utilization is monitored for baseline, cases 1 and 2. MB and PSU fan speeds are

logged continuously. The temperature variation with time for CPUs and VRDs in the MB portion and HDD in the PSU portion is reported. The thermal limitation of VRD and DIMM component is 85ºC. As expected, the server internal fan array's ability to drive the CPU temperatures to target values of dead band control is observed. And steady state is achieved in 15 minutes. The test was repeated three times in total for repeatability.

RESULTS AND DISCUSSION

Externally controlled, externally powered server fans: Figure 6 shows the results for the externally powered, externally controlled tests across the ranges of inlet static pressures studied when the stack of servers 10 to 40 were operated at maximum computational load. Note in the figure that the y-axis crosses at a negative static pressure value. A best curve fit taken from all CPU temperatures across each of the four servers shows the average trend in the data. At the lower fan speeds of 0% and 5%PWM duty cycle, CPU temperatures begin to reach their upper functional limits when the static pressure becomes negative.

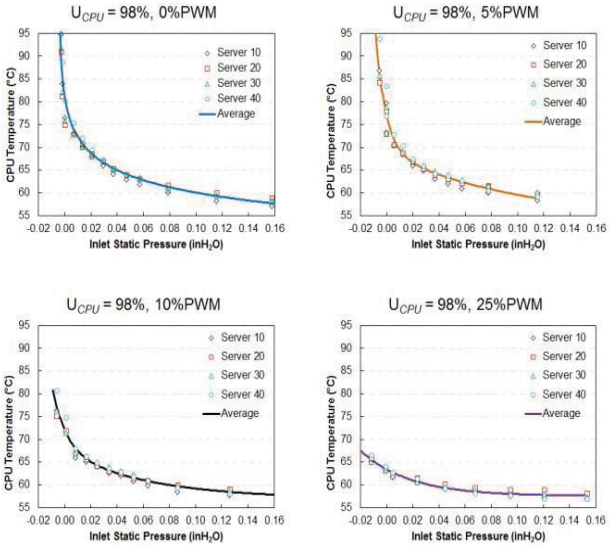

Figure 6: Impact of inlet static pressure on server CPU temperature at fixed fan duty cycles of (a) 0%PWM (b) 5%pwm, (c) 10%PWM and (d) 25%PWM

As seen in Figure 7 these lines of constant fan duty cycles eventually converge. This convergence takes place at higher static pressures for higher fan speeds. As an approximation from the curve fits shown in Figure 6, the die temperatures at 0% and 5%PWM signals are equal at 0.07 inH2O inlet static pressure. CPU temperatures for 0%, 5%, and 10%PWM all converge at 0.11 inH2O inlet static pressure. These curves converge with 25%PWM fan speed at 0.15 inH2O inlet static pressure. This may be expected because, beyond a certain static pressure, the airflow provided by the blower begins to overtake the airflow provided by the internal server fans.

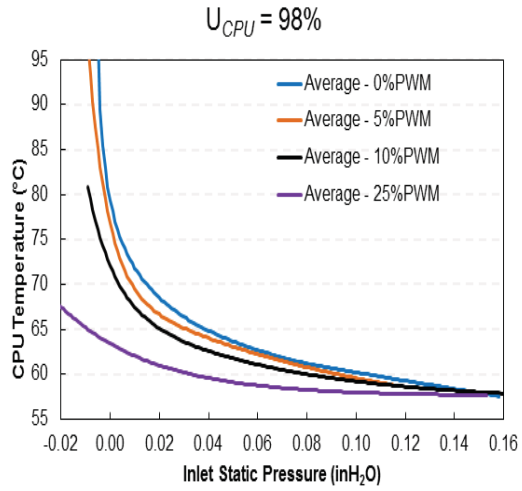

Figure 7: Convergence of cooling performance at positive inlet static pressures with fans operating at 0%, 5%, 10% and 25%PWM fixed duty cycles

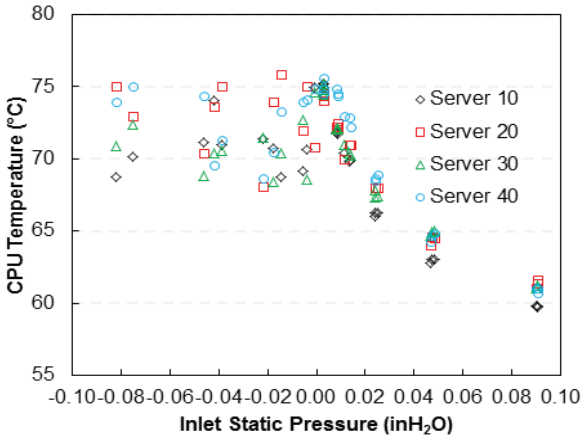

Figure 8: CPU temperatures maintained at or below controls limit threshold across inlet static pressure range with the servers operating at 98% CPU utilization.

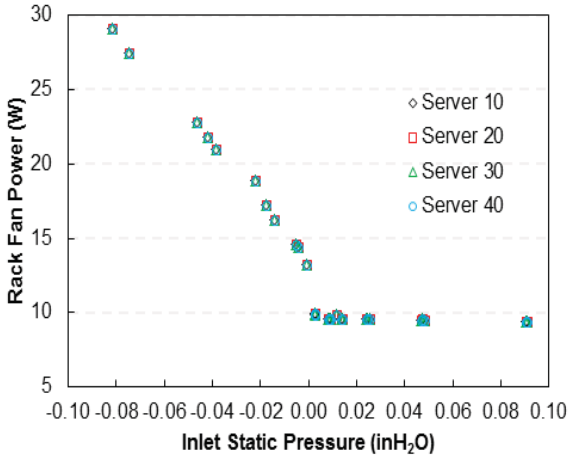

Figure 9: Fan power as a function of inlet static pressure

Internally Controlled, Externally Powered Server Fans: When the server fans are controlled internally, it is expected that the CPU temperatures will not exceed their target values specified in the dead band control. This is seen in Figure 8 with the steady state CPU temperatures across the range of inlet static pressures studied. At lower inlet static pressure values (particularly in the negative region) CPU temperatures are seen to fluctuate around the control range of 68°C to 74°C. At higher inlet static pressures, the air flow provided by the blower over powers the internal server fans and more consistent temperatures are observed. In terms of actual fan power (MB fans) consumed by the rack, a piecewise linear trend is observed in Figure 9. Below a threshold, fan power increases steadily with negative static pressure. This increase in demand is caused by the need to maintain temperatures within the specified target range despite the increase in external impedances. Above 0.003 inH2O, the rack fan power flattens out since the fan control algorithm no longer needs to engage and the fans are remaining stable at idle speeds.

To better understand the fan control scheme's response to external impedances, the results from the transient thermal testing of a single server for baseline, Case 1 and 2 scenarios is reported here. As shown in Figure 5, the external impedance is greater for Case 2 than Case 1. The server's fan control scheme drives the component temperatures as programmed, however any external impedances shall increase the operating temperature of the server. When the external impedance increases, the server component temperatures will become a function of the pressure differential across the server. Figure 10 and 11 shows CPU and CPU VRDs are observed to be at higher surface temperatures when steady-state is achieved for Case 1 and 2.

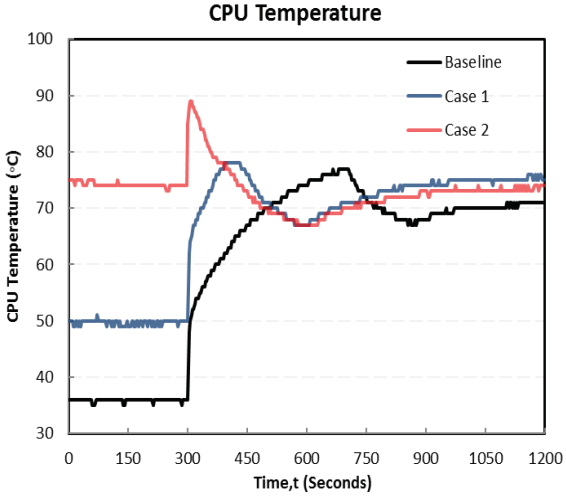

Figure 10: CPU Temperature variation when CPU utilization is increased from idle to 98% utilization at 300s

Any subsequent step increase in CPU utilization necessitates higher MB fan speeds. As seen in Figure 10 and 12, when the computational load peaks the CPU temperatures sharply increases, and fan speed control scheme engages and ramps up the MB fans. The PSU fan is expected to idle at 30% duty cycle as the inlet air temperature doesn't exceed 25ºC.

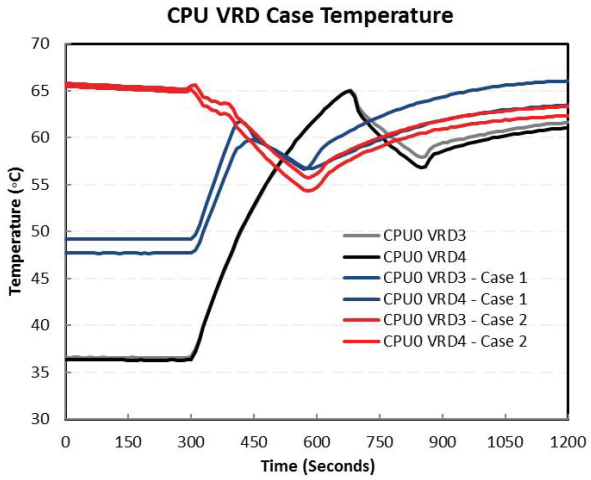

Figure 11: CPU Voltage regulator case temperatures upon a step CPU Utilization increase to 98% utilization

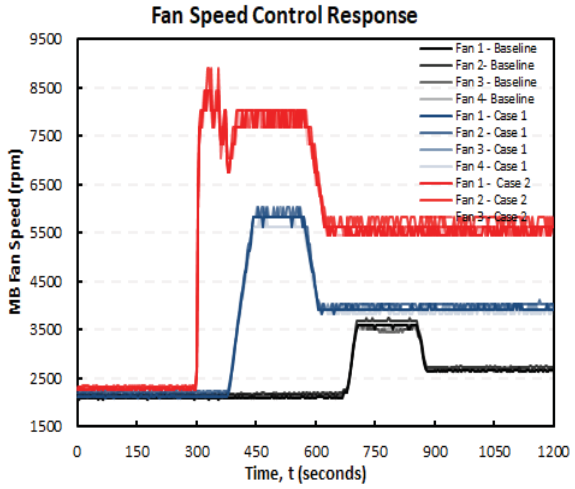

Figure 12: Increase in MB fan speeds in response to step increase in CPU Utilization

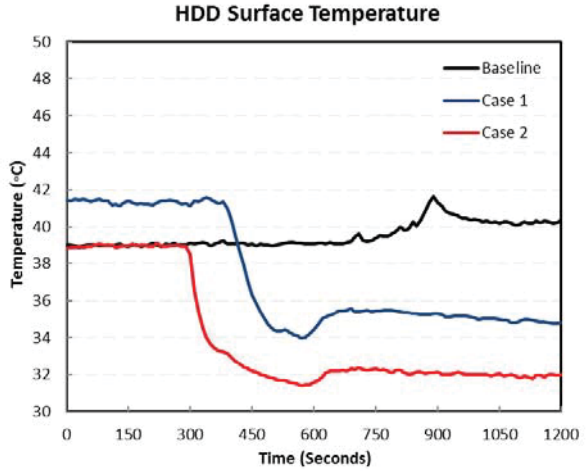

Figure 13: Reduction in HDD surface temperature for under- provisioned scenarios

When the inlet pressure is negative, PSU and MB fans behave differently, especially when the PSU fan is operating at a lower fan speed and the MB fans are at full speed. With MB fan at 100%PWM and PSU fan at 30%PWM, Figure 14 shows that PSU fan speed can decrease significantly while the average MB fan speed increases at negative inlet pressures. This demonstrates the possibility of PSU fan being overpowered by the MB fan array. A reversal of flow through the PSU portion is suspected as the two parallel flow paths within the server consists of a fan/group of fans generating more pressure than the other due to the disparity in fan control logic and fan loads. Also, a reduction in HDD surface temperature is observed due to increased forced convection from reversed flow, as shown in Figure 13, for case 1 and 2 as soon as the MB fan speeds ramp up.

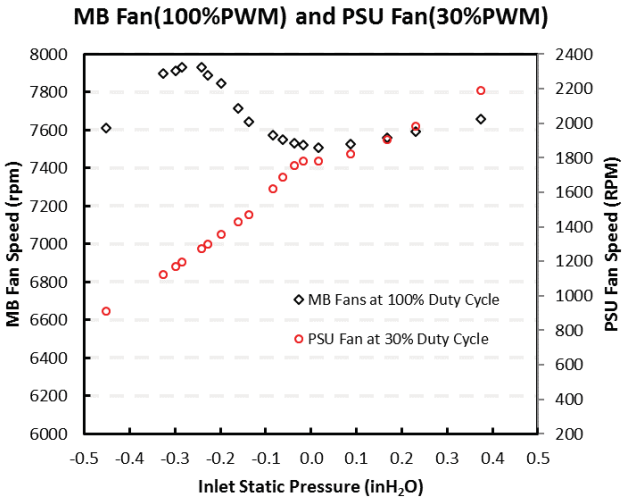

Figure 14: PSU and MB fan speeds over a range of inlet static pressures at fixed fan duty cycles

To further understand the characteristic relation between flow rate and pressure drop across the server, we can obtain a flow curve with internal fans operating at a fixed fan speed. The experimental testing procedure applied can be found in [10]. Figure 15 shows the experimental flow curve for the Open Compute server under study for 100% and 50% fan duty cycles. The Y-axis shows the inlet static pressure of the server, where negative inlet pressure represents external impedance. The flow curve provides the server flow rate for a fully-provisioned scenario at the free delivery condition. This represents the 'Baseline' scenario discussed in the previous section titled transient response to peak CPU utilization. The power curve given in Figure 15 shows that as the direction of the airflow reverses the fan power increases significantly. MB and PSU flow paths have different internal system impedances and different operating fan speeds. By isolating the two airflow paths within the server, we can isolate the effects of MB and PSU fans and separately understand the influence of each.

Figure 16 shows the flow curve for just the MB portion with MB fans at 100% PWM duty cycle and the PSU section completely blocked off. The flow curve for MB portion is comparatively similar to the flow curve with all fans (MB and PSU fans) at 100% duty cycle in the first quadrant. However, when the inlet pressure is positive the PSU portion acts as a leakage path into the server outlet. The flow curve for the server with MB fans at 100% and PSU fans at 30% duty cycle follows a similar trend and the reduction in PSU fan speed has offset the curve as seen in Figure 16. To visualize the direction of airflow for MB and PSU sections, a Computational Fluid Dynamics (CFD) tool can be used. However, the fan performance curves generally applied as a boundary condition cannot model the airflow in in reverse direction when applicable. To address this requirement, an extended fan performance curve is proposed in the following section.

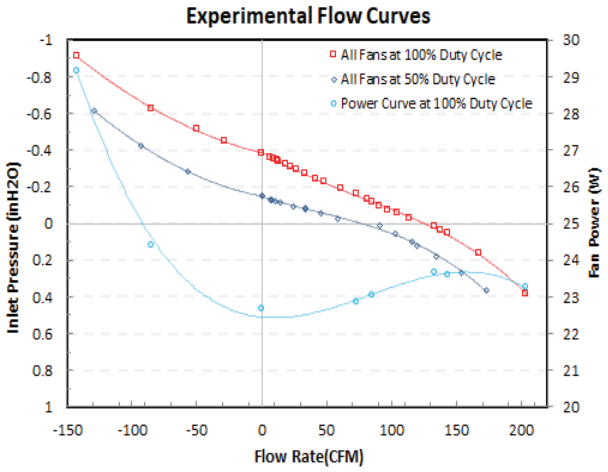

Figure 15: Experimental flow curves for Open Compute Server under study

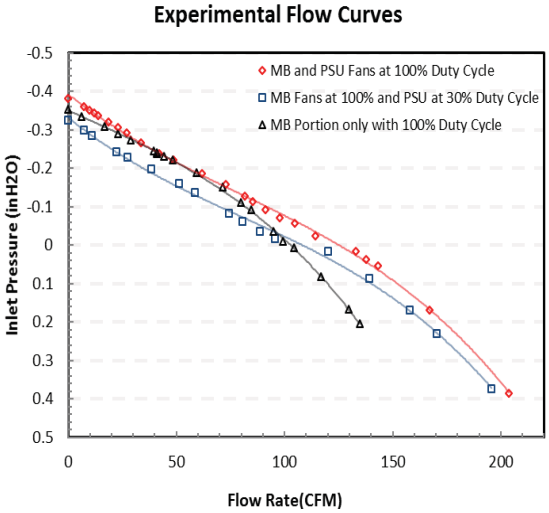

Figure 16: Experimental flow curves for 100% duty cycle MB fans but different PSU fan duty cycles and MB portion

Extended Fan Curve: Airflow provided by an internal fan depends on the operating pressure differential across itself and the nature of impedances (internal and external) at the fan inlet and outlet regions. A typical fan performance curve provides pressure characteristic between the fan shut

off and fan free delivery condition in the first quadrant. However, the second quadrant performance can be the pressure versus reversed fan flow rate. The pressure characteristic can also be further extended in to the fourth quadrant [14].

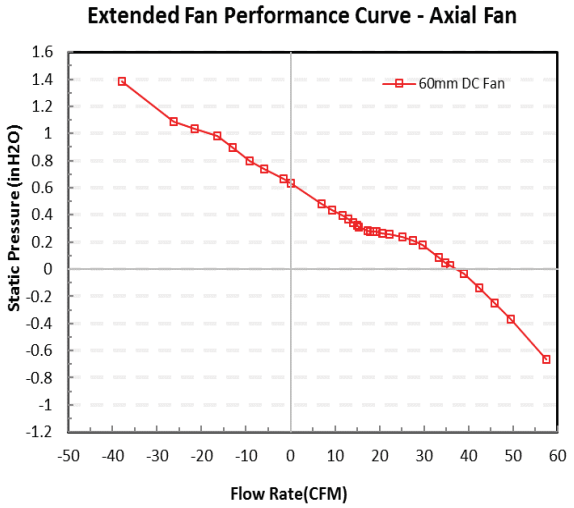

Figure 17: Fan performance in Second and Fourth Quadrant for the 60mm axial fan

An extended fan performance curve, characterized experimentally with the fan operating at 100%PWM, for the 60mm DC fans installed in the server is shown in Figure 17. The static pressure in the ordinate is the fan backpressure. Figure 18 shows that upon flow reversal through a fan, there is a drastic reduction in operating fan speed and a significant increase in the fan power. Therefore, any internal fan operating with backpressure (negative inlet pressure) within the server will suffer performance degradation. And such a scenario can occur in servers consisting of fan arrays/fans with distinctive airflow paths within the enclosure system.

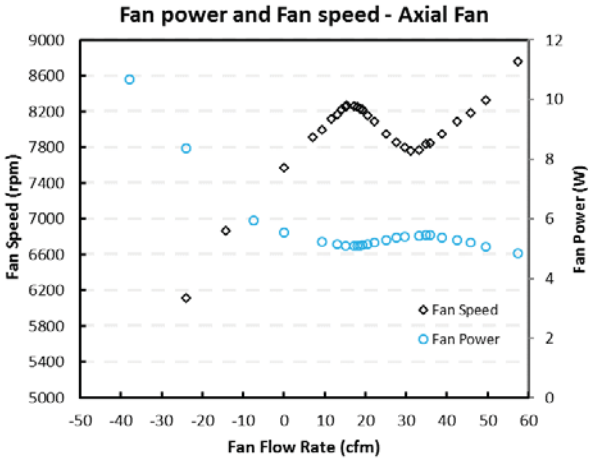

Figure 18: Fan speed and power consumption of a 60mm axial fan

An extended fan performance curve can be used to co-relate flow rate through the server to pressure drop in cases where negative inlet pressure could exist. When the local pressure differential across the 60mm axial fan exceeds 0.61 inH2O, flow reversal onsets. Also, it can be used as a boundary condition in CFD studies to enable prediction of directionality of airflow in an enclosure system with fan/group of fans operating in parallel and cases where negative inlet pressure can exist. However, it should be noted that an extended fan performance curve doesn't account for the reduction in fan performance due to fan-system affects. The experimental flow curve [10] does account for fan-system affects, in addition to external impedances, but it cannot be used for visualizing server internal flow paths as it is generally prescribed as a boundary condition at the server inlet or outlet.

CFD simulations of the detailed model of the server under study, with extended fan curves as boundary conditions for the internal fans, is carried out for different airflow provisioning cases. For further details of the detailed model of the server, please refer to [15]. From the experimental flow curve in Figure 15, fully-provisioned case is at the free delivery condition of 125CFM. Additionally, 20%, 50% and 120% of fully-provisioned case is specified as a fixed inlet flow rate to the server. The MB fans are at maximum speed while the PSU fan is at 30% duty cycle (~4000rpm).

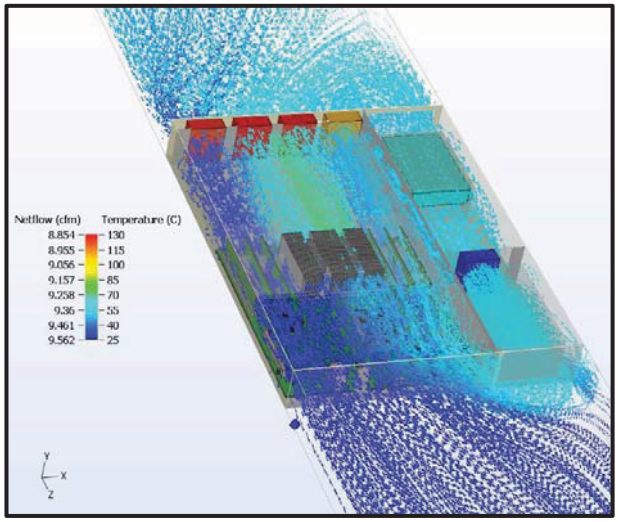

Figure 19: Temperature streamline plot and net flowrate of fans for 20% provisioning of Open Compute Server with an extended fan curve

Figure 19 shows the temperature streamlines for the model utilizing extended fan curves as a fan boundary condition. The results show that for the PSU fan, a net flow in the reverse flow direction can be observed. This provides a possible explanation for the reduction in HDD surface temperature observed in Figure 13. The reversed flow in this scenario provided cooling effect and reduced the HDD surface temperatures. Therefore, air recirculation within the server is clearly observed at 20% provisioning when extended fan curves are used. On the other hand, the PSU fan flow rate is zero for the regular fan curve model. As shown in Figure 20, the net PSU fan flow rate is zero for

20% airflow provisioning when a regular fan curve is applied as a fan boundary condition.

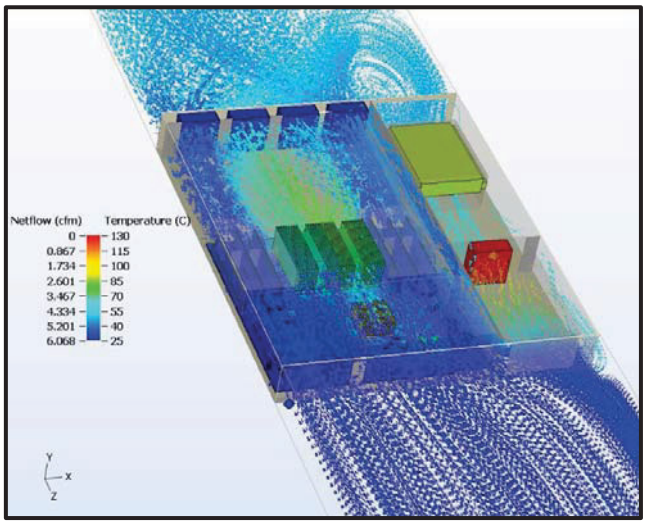

Figure 20: Temperature streamline plot and net flowrate of fans for 20% provisioning of Open Compute Server with a regular fan curve

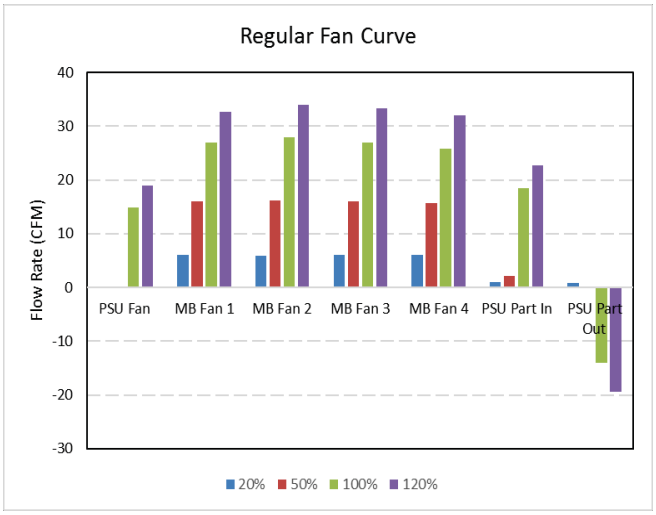

Figure 21: Fan flow rates and PSU portion net flow rates for various provisioning cases using a regular fan curve as boundary condition

Finally, the CFD results for the different airflow provisioning cases are summarized and shown in Figure 21 and Figure 22 where a regular fan curve and an extended fan curve is used as fan boundary condition respectively. MB and PSU fan flow rates and the net flow rate in and out of the PSU portion is reported. Comparing the 20% provisioning scenario in Figure 21 with Figure 22, it can be observed that the extended fan curve accounts for the reverse flow through the PSU portion. There is an onset of internal recirculation due to the negative inlet pressure at the PSU fan for 20% provisioning. The additional supply air leakage through the PSU portion for 120% provisioning scenario is also evident in the extended fan curve model as seen in Figure 21.

In Figure 21, the PSU fan flow is again estimated to be zero for the 50% provisioning scenario while using a regular fan curve. For scenarios when inlet pressure at the PSU fan is negative, the regular fan curve fails to account for the heat infiltration back into the PSU due to the recirculating exhaust air. The occurrence of such flow paths during under-provisioned flow conditions can lead to component failure and/or reduction in the efficiency of the overall cooling system design. In addition to that, the internal recirculation can affect the local temperature measurements within the server.

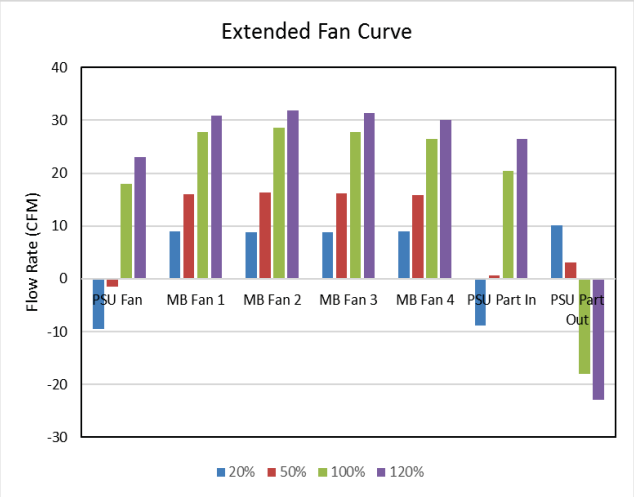

Figure 22: Fan flow rates and PSU portion net flow rates for various provisioning cases using an extended fan curve as boundary condition

SUMMARY AND CONCLUSIONS

In this paper, the effects of a positive and negative static pressure differential between the supply inlet and return exhaust of an Open Compute server operating in a simulated rack setup in an idealized laboratory scenario is studied. A truly optimized system will require dynamic balancing between the fan power within a server and the fan power from blowers within the data center's CRAC units. From the results presented here, room level pressure variations may significantly impact the performance of individual servers, especially if allowed to fall below a certain pressure threshold or to negative pressure values.

The work of [16] proposed a control system utilizing dampers within either the cold or hot aisle to physically restrict airflow and maintain static pressure at specified levels. It is possible in some cases that a large enough static pressure could be built up across a server rack allowing for complete removal of server fans. All airflow through a server may be provided solely though room level airflow from CRAC units.

Other alternatives may include additional, larger fans placed either within the cold aisle or hot aisle. As demonstrated in [17], larger fans will typically operate more efficiently than their smaller counterparts. Existing products

are available on the market that have active floor tiles which include fans coupled the perforated floor tiles of the raised floor. Alternatively, an exhaust fan placed within the return air plenum of the hot aisle may serve similar function.

The overall goal should be to minimize total system fan power. This will likely require various pressure measurements throughout the data hall, either as static pressure at the front of racks in the cold aisle, or ideally, as a differential pressure measurement across a rack. More importantly at the server level, the enclosure thermal design should consider the phenomenon of internal recirculation of exhaust air. As demonstrated in this study, utilizing extended fan curves when available as boundary conditions for internal fans can model the onset of internal recirculation. The extended fan curves can also be used to optimize the internal flow paths within the server and improve the overall server thermal performance.

Acknowledgements

This work is supported by NSF IUCRC Award No. IIP-1738811.

REFERENCES

[1] R. Sharma, C. Bash, and C. Patel, "Dimensionless Parameters for Evaluation of Thermal Design and Performance of Large-scale Data Centers," 8th AIAA/ASME Jt. Thermophys. Heat Transf. Conf., no. June, pp. 1–11, 2002.

[2] P. Kumar and Y. Joshi, "Fundamentals of Data Center Airflow Management.," Springer, pp. 199–236, 2012.

[3] M. Bharath, S. K. Shrivastava, M. Ibrahim, S. A. Alkharabsheh, and B. G. Sammakia, "Impact of Cold Aisle Containment on Thermal Performance of Data Center," InterPACK, pp. 1–5, 2013.

[4] S. K. Shrivastava and M. Ibrahim, "Benefit of Cold Aisle Containment during Cooling Failure," InterPACK, pp. 1–7, 2013.

[5] S. A. Alkharabsheh, S. K. Shrivastava and B. G. Sammakia, "Effect of Cold Aisle Containment Leakage on Flow Rates and Temperatures in a Data Center," InterPACK, pp. 1–9, 2013.

[6] J. Musilli and B. Ellison, "Facilities Design for High Density Data Centers," Intel IT White Pap., no. January, 2012.

[7] E. Frachtenberg, D. Lee, M. Magarelli, V. Mulay, and J. Park, "Thermal design in the open compute datacenter," Intersoc. Conf. Therm. Thermomechanical Phenom. Electron. Syst. ITHERM, vol. 94025, pp. 530–538, 2012.

[8] B. Fakhim, N. Srinarayana, M. Behnia and S. W. Armfield, "Thermal Performance of Data Centers-Rack Level Analysis," Components, Packag. Manuf. Technol. IEEE Trans., vol. 3, no. 5, pp. 792–799, 2013.

[9] H. E. Khalifa, "Optimization of Enclosed Aisle Data Centers Using Bypass Recirculation," J. Electron. Packag., vol. 134, no. June 2012, p.20904, 2012.

[10] H. A. Alissa, K. Nemati, B. G. Sammakia, K. Schneebeli, R. R. Schmidt, and M. J. Seymour, "Chip to Facility Ramifications of Containment Solution on IT Airflow and Uptime," vol. 6, no. 1, pp. 67–78, 2016.

[11] P. Sarti and F. Frankovsky, "700W-SH/450W-SH Power Supply Hardware v1.0, 2012. http://www.opencompute.org/assets/download/Open_Compute_Project_700W_450W_Power_Supply v1.0.pdf.," [Online]

[12] J. Na, "Facebook Server Fan Speed Control Interface," pp. 1–10., Open Compute Project, 2014

[13] J. E. Fernandes, "Minimizing Power Consumption at Module, Server and Rack-Levels within a Data Center Through Design and Energy-Efficient Operation Of Dynamic Cooling Solutions," PhD Dissertation, University of Texas Arlington, 2015.

[14] Jorgenson R, Fan Engineering. Buffalo, New York: Howden Buffalo, Inc.

[15] V. Pandiyan, "Development of Detailed Computational Flow Model of High End Server and Validation Using Experimental Methods", Thesis, University of Texas Arlington, 2012.

[16] D. Kennedy, "Ramification of Server Airflow Leakage in Data Centers with Aisle Containment," Tate Inc., 2012.

[17] B. Nagendran, S. Nagaraj, J. Fernandes, R. Eiland, D. Agonafer, and V. Mulay, "Improving cooling efficiency of servers by replacing smaller chassis enclosed fans with larger rack-mount fans," in Proceedings of the 2014 ITherm Conference, Orlando, FL, USA, May 27-30, 2014.