Impact of Fans Location on the Cooling Efficiency of IT Servers

Sadegh Khalili [1], Mohammad Tradat [1], Husam Alissa [2], Cheng Chen [3], Bahgat Sammakia [1]

[1] Department of Mechanical Engineering, Binghamton University-SUNY, NY 13905, USA
[2] Microsoft, WA 98052, USA
[3] Facebook, CA 94536, USA
Email: skhalil6@binghamton.edu

ABSTRACT

The role of data centers in modern life has expanded rapidly over the past decades. In addition, this expansion has resulted in a significant increase in the share of data centers in total energy consumption of the world. Thus, reliability and energy efficiency have become a common concern in data centers. Information technology equipment (ITE) and cooling infrastructure are the largest power consumer in data centers. The cooling power in a data center depends on the amount of heat dissipated by ITE. Therefore, the thermal design of the ITE impacts not only the ITE power but also affects the infrastructure power and has a significant role in the overall efficiency of data centers. This paper studies the impact of fans location and airflow balancing on the thermal performance and power of a server. A detailed computational fluid dynamic (CFD) model of the server is built, calibrated and validated using experimental test results. Next, impacts of moving fans to the rear side of the chassis on the flow rate and temperature of components are investigated. Special attention is given to controlling airflow through power supplies.

KEYWORDS: Data center, Server, Thermal design, CFD, Airflow balancing, Internal recirculation, Power supply.

NOMENCLATURE

AFC	active flow curve
BMC	baseboard management controller
CFD	computational fluid dynamics
C_P	specific heat capacity, J/kg°C
FS	fan speed, rpm
HDD	hard disk drive
IPMI	intelligent platform management interface
ITE	information technology equipment
PDU	power distribution unit
PSU	power supply unit
Q	power, W
RU	rack unit
SRI	server recirculation index
T	temperature, °C (°F)
$\dot{V}$	flow rate, m^3/s (cfm)

Greek symbols

ρ	mass density (kg/m^3)

Superscripts

avg	average

INTRODUCTION

A large part of daily activities in modern societies relies on continues access to information technology (IT), e.g. banking, government, traffic routing, telecommunication, education, internet-of-things services, etc. As a result, data centers have become mission-critical facilities for many aspects of today's life. Any interruptions in critical data center operations can lead to loss of data or service and can be extremely costly without an effective backup strategy. In a recent cost analysis of data center outages by Ponemon and Vertiv [1], IT equipment failure is found to cause the highest outage cost between all the primary root causes of the outages. This demonstrates the importance of the reliability of IT equipment in data centers. Besides, the energy consumption in data centers has increased due to building new data centers along with the increase in power density in the existing data centers to propel the fast-growing demand for online services. A recent study by Shehabi et al. [1] estimates that energy use in the U.S. is anticipated to increase by 4% from 2014-2020, resulting in 73 billion kWh power consumption in 2020. The ITE (servers, storage, networking etc.) and infrastructure (cooling system, lighting, un-interruptible power supply system, etc.) are the two largest components of total energy consumption in data centers [2]. Typically, the cooling systems consume about a third to half of the total power in enterprise data centers [3]. Cooling efficiency can be improved by implementing practices such as deploying containment systems, blanking bypass and recirculation paths, and installing directional floor tiles [4–7]. Depending on the location of the data center, free cooling can be utilized to decrease this share when the ambient temperature is sufficiently lower than the facility temperature [8–10]. Implementation of liquid and two-phase cooling are alternative ways for improving cooling efficiency in data centers due to much higher specific heats (latent heat for two-phase cooling) and higher heat transfer coefficients compared to air [8,11]. Currently, most enterprise data centers implement air-cooling [12] because of its proven high reliability, compatibility with the environment and lower initial and maintenance costs.

Shehabi et al. [2] anticipated a continues 3% annual growth in server shipments through 2020. Also, the rising demand for cloud computing in the world has increased the motivations for high-density data centers. This has urged engineers to develop high-performance server designs that are also energy efficient. An analytical energy model by Vertiv [13] demonstrated that 1 W of power savings at the server component level creates a total of 2.84 W of savings in the facility energy consumption. Therefore, improving the cooling efficiency of servers is a crucial part of any effort toward an energy efficient data center. The server fans are responsible for drawing in cold air through front panel to cool the internal components such as hard disk drives (HDDs), CPUs, memory modules, etc. Typically, the temperature of the components is a function of their utilization and can vary in time. The baseboard management controller (BMC) is responsible for monitoring the physical variables such as the temperature of components, the speed of chassis fans, power supply voltage and operating system functions using built-in sensors, and taking required actions for maintaining operating conditions of server components. Wang et. al [14] proposed a multi-input multi-output fan controller that utilizes thermal models to predict the server temperatures and adjust fans speed accordingly. In most servers, an internal temperature sensor is implanted in the control-panel board in front of the server chassis to monitor the inlet air temperature.

The chassis dimensions and component locations inside it are among important parameters in a server thermal design. Frachtenberg et al. [15] described the impact of changes made in the design of power supply unit (PSU), motherboard, chassis and the thermal design of a Facebook server on the total cost of ownership (TCO) of a data center. They suggested spreading the hot components side by side in a 1.5 rack unit (RU) chassis to maximize the compute-capacity of a rack per cooling-energy ratio. In addition, they also moved the hottest components (i.e. processors and memory) closer to the server's inlet to supply the coldest air to hottest components first. Further improvement can be achieved by implementing air ducting within a server to route cooling air within a server and decrease bypass of cold air over the critical components. Mani et al. [16] studied the impact of ducting on the fan power in a high-end web server. They showed up to 39% savings in fan power can be achieved via improving airflow distribution inside the server by redesigning the implemented ducting. Dhadve et al. [17] conducted a similar study on an Open Compute Big Sur server. Vuckovic and Depret [18] compared three local cooling solutions, namely traditional heatsink, remote heatsink utilizing heat pipes and direct liquid cooling. Khalili et al. [19] studied the impact of chassis perforations and internal recirculation on the thermal performance of three commercial 2 RU servers experimentally. They found that an internal recirculation and location of sensors relative to chassis perforations can cause a discrepancy between the readings of IPMI and external sensors (false high IPMI temperature). A false high IPMI temperature can lead to an elevated speed of chassis fans and consequently, higher power consumption, a shorter lifetime of fans and lower reliability of the server. In addition, they reported that a high false IPMI temperature can cause CPU throttling due to an open loop safety action taken by the BMC. Furthermore, their results showed that the intake flow rate can be increased by 25% by sealing the gap between fans and covering perforations on the sides of a tested chassis.

The goal of server thermal design is to dissipate the generated heat and maintain the components temperature in a safe range with a minimal expenditure of energy and component cost. Generally, fans can be used to draw air through a chassis or to blow air through it. A blowing fan upstream of the server increases the pressure inside the chassis, which can prevent dust from entering a chassis that is not well-sealed. In addition, a blowing fan introduces turbulence to the airflow which improves the heat transfer characteristics of cooling air. However, in blowing systems, the cooling air heats up as it passes over the fan motor, which decreases its cooling capacity.

Most enterprise servers utilize fans somewhere between the front panel and rear of the server. In these servers, the pressure differential between upstream and downstream of the fans can create a recirculation through chassis fans [19]. If the spaces up- and down-stream of fans are not segregated well, a portion of air, after entering the ITE and cooling hot components, recirculates inside the ITE. Consequently, hot air mixes with the fresh cold air from the intake. This mix of hot and cold air streams has some known consequences such as reduced cooling capacity, hot spot occurrence, and decreased thermodynamic efficiency. As a result, server fans may be operated at a higher speed to supply more cold air to make up for this effect.

In an exhaust fan system, fans are installed downstream of the server and draw air through the chassis. This results in a negative gauge pressure inside the chassis. If the server is located in a dusty or dirty space, this negative pressure can pull dust into the chassis through openings on its side walls. On the other hand, with this design, the cooling air reaches the fan motors after cooling the electronic components which results in cooler components. In addition, this design minimizes hot air recirculation inside the chassis which increases the cooling efficiency. Furthermore, airflow pattern is more predictable, so air can be ducted precisely to where it is needed. Finally, moving fans to the rear sides of a server releases a valuable space in the vicinity of processors.

Most of enterprise servers utilize server fans somewhere between the front and rear of the server to allow space for adding multiple hard disk drive (HDD), various ports, and peripheral component interconnect (PCI) cards. The Open Compute Project (OCP) has introduced some servers in which fans are mounted on the rear side of the servers, e.g. project Olympus, Tioga and Big Basin servers [20,21]. In the Open Rack project [22], all of the cables and interconnects are made from the front of the rack, the ITE is hot-pluggable and serviced from the front of the rack. This design eliminates the need for accessing the rear of the rack and allows moving fans to the back of servers' chassis. Also, service personnel and operators need to spend less time in the hot aisle.

This paper is a follow-up study on [19] and focuses on the impact of location of server fans on power and thermal performance in a server. A detailed CFD model of a 2-RU server is built. The model is calibrated and validated under two different operational conditions. Next, the CFD model is used to investigate the benefits and consequences of moving server fans to the rear of the chassis. Special attention is given to CPU temperatures and airflow rate of power supply units (PSU). Finally, a method for balancing the airflow rate between the chassis and PSU fans is described.

Experimental Results

A Dell PowerEdge 2950 2-RU server is selected for this study. The internal design of the server is shown in Fig. 1a. This server utilizes four parallel fans for cooling the CPUs, chips, memory modules and PCBs. The server fans span about three-quarters of the server's width. Two PSUs provide power redundancy with a maximum power of 750 W each. A dedicated fan is installed at the intake of each PSU. An air baffle (ducting) is utilized to distribute airflow within the server. Two Intel Xeon E5345 CPUs are deployed on the mainboard with a thermal design power (TDP) of 80 W. A direct measurement of CPU power was not possible due to limits in the server's firmware. Although the thermal design power (TDP) of this CPU is reported 80 W by Intel [23], a sustained maximum power dissipation of 108 W is reported in [24] which is more consistent with experimental box power measurements in this study via a power distribution unit (PDU). The difference between the reported CPU power and the TDP may be due to a significant leakage power in this CPU generation (65 nm technology) when operated at temperatures close to its maximum allowed temperature. Eight memory modules are installed on the mainboard. A hard disk drive (HDD) is installed in one of the six HDD bays in front of the server while the rest of bays are filled with HDD blank fillers. Air temperature is measured at one point at the rear of the server, and across the power supplies using T-type thermocouples. In addition, the intelligent platform management interface (IPMI) data including CPU temperatures, inlet air temperature and the speed of chassis fans are collected. Power of fans at maximum speed is measured via a Keysight E3634A DC power supply and fan speeds are measured by a UNI-T UT372 tachometer.

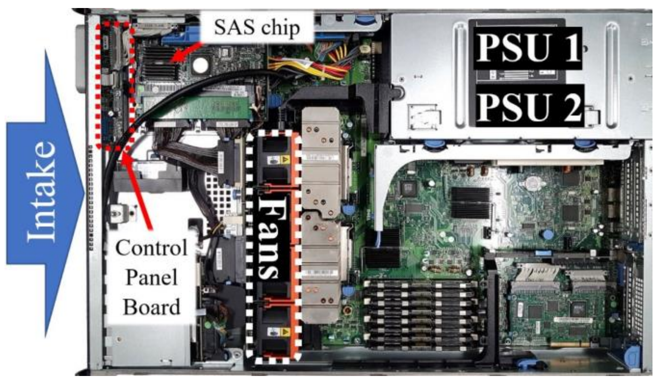

(a)

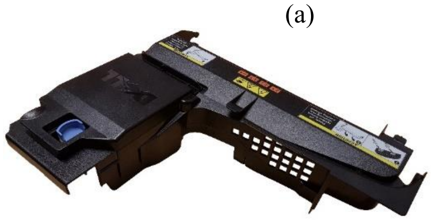

(b)

(c)

Fig. 1: (a) Internal Design of Dell Power Edge 2950 when air baffle is removed, (b) Air baffle, (c) PSU

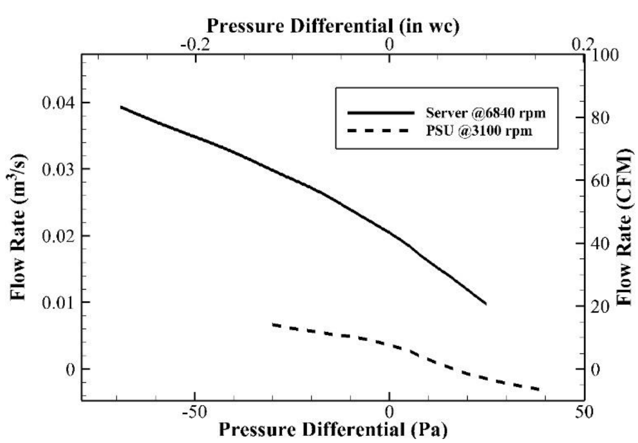

Fig. 2: Active flow curve of the server and PSU.

The theoretical operational flow curve of a server can be generated by algebraic subtraction of the impedance and fan curve. However, the theoretical operational flow curve may not give an accurate estimate for airflow through the ITE under different conditions. Nemati et al. [25] demonstrated that the operational flow curve of a server can be significantly different from its experimental active flow curve (AFC) when recirculation is present inside the ITE. The AFC of the server in this paper is extracted by mounting it on a flow bench and measuring the pressure differential across the server at different supplied flow rates. Experimental validation of this approach is discussed in [26]. A similar approach is used for extracting AFC of the PSU. The AFC of PSU will be used in CFD simulations in the next section. This allows modeling the PSUs as a black-box and eliminates the need for capturing geometrical details inside the PSUs.

A fog generator is used to visualize the flow field inside the box when the server lid was replaced with an acrylic sheet. The visualizations demonstrated the presence of a recirculation inside the server and showed that the recirculation expands with backpressure. This expansion is shown in Fig. 3.

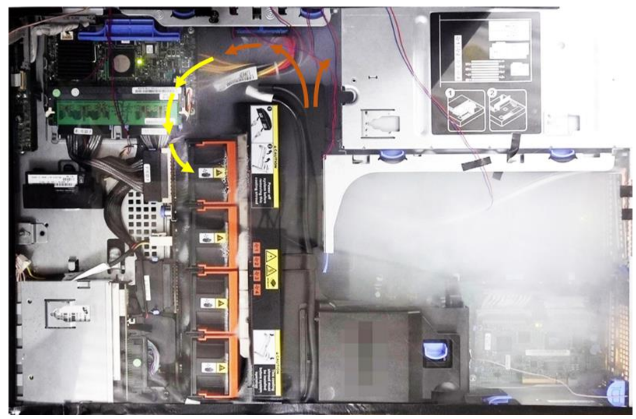

(a) No backpressure

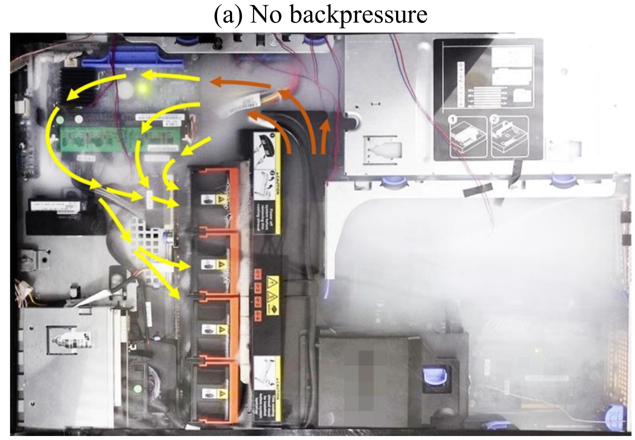

(b) Backpressure = 28 Pa

Fig. 3: Internal flow field at different magnitudes of backpressure

The measured server flow rate, air temperature at the inlet and exit of the server, and IPMI CPU and inlet temperatures under two different environmental and operating conditions are presented in Table 1. In case 1, fans operated at their normal speed (6,800 rpm) and no backpressure was imposed on the server. In case 2, 28 Pa backpressure was imposed on the server and fans operated at approximately 12,000 rpm. The firmware in this server does not allow controlling fan speeds. Instead, there is a push-in button on the chassis that monitors the presence of top cover when the server is operating. By releasing this button, the server fans operate at an elevated speed (~12,000 rpm). It is worth noting that the reported experimental CPU temperatures in Table 1 are time average of reported IPMI and measured data over 2 minutes of steady operation of the server when the CPUs were fully stressed using MPrime (Prime 95) [27]. It should be mentioned that the accuracy of the IPMI inlet air temperature sensor is reported 1°C in the manufacturer datasheet [28]. Finally, the differential temperature across the PSUs is presented in Table 1. In this table, $T_{\text{in-IPMI}}$ is the measured temperature by a built-in temperature sensor on the front control board of the server. The experimental data is used for calibration and validation of a CFD model of the server which will be discussed in the next section.

Table 1: Experimental and CFD results for cases 1 and 2.

	Case 1: No Backpressure FS: 6840 rpm		**Case 2:** Backpressure = 28 Pa FS: 12000 rpm	
	Measured	CFD	Measured	CFD
$\dot{V}_{\text{intake}}$	43 cfm	41.8 cfm	70 cfm	71.2 cfm
T_{supply}	18.5 °C	18.5 °C	19.5 °C	19.5 °C
T_{exit}	30.6 °C	30.0 °C	29.9 °C	30 °C
$T_{\text{in-IPMI}}$	21 °C	20.3 °C	35 °C	34 °C
$T_{\text{CPU-A}}$	64.7 °C	64.5 °C	66.8 °C	66.9 °C
$T_{\text{CPU-B}}$	59.4	58.9 °C	58.9 °C	58.6 °C
ΔT_{CPUs}	5.3 °C	5.6 °C	8 °C	8.3 °C
$\Delta T_{\text{PSU-A}}$	5 °C	4.6 °C	11.0 °C	11.5 °C

CFD Modeling

Current Design

A detailed model of the tested server is built in a commercially available CFD package - 6Sigma ET® - which will be used for investigating the impact of fan location on the flow and thermal performance of the server. Main components of the server such as chassis perforations, air baffle, heat sinks, obstructions, memory modules, and drive bays are modeled in detail. The chassis, heat sink, and air baffle in the actual server and the detailed CFD model are shown in Fig. 4. Physical dimensions are measured using a digital caliper on the actual server. The temperature of the supplied air is set similar to the supply air temperature in the experiments. Heat dissipation of the main components are extracted from data sheets or estimated based on the server total power consumption and applied into the CFD model. The rate of heat dissipation in the active PSU (PSU-A) is estimated based on the measured server power and efficiency of the power supply. For the redundant PSU, the heat dissipation rate is estimated based on its fan's power consumption. The power dissipation of the HDD is assumed 6 W based on the manufacturer data. The power of each memory module is estimated to be 3 W based on the available memory chip power data and number of the chips in the module. The flow curve of server fans is extracted by testing a fan at its maximum speed on the flow bench. The affinity laws are used for estimating power consumption of fans at different rotational speeds based on the measured power at maximum speed. The thermal conductivity of the thermal interface material (TIM) in the CFD model is calibrated by comparing the CPU junction temperature extracted from IPMI data with the mean temperature of the chip in the CFD model in case 1.

Results of the CFD model are compared with experimental data in Table 1. The maximum difference between the CFD and experimental results in Table 1 is 5%. The differences can be due to the accuracy of sensors, and uncertainties in the measurements, power input values and material properties. The AFC of the server model is compared with the experimental AFC in Fig. 5 which shows a maximum error of 6%. The good agreement between results of the CFD model and experimental data gives the confidence to use the calibrated model for evaluation of the server performance under various environmental and operating conditions.

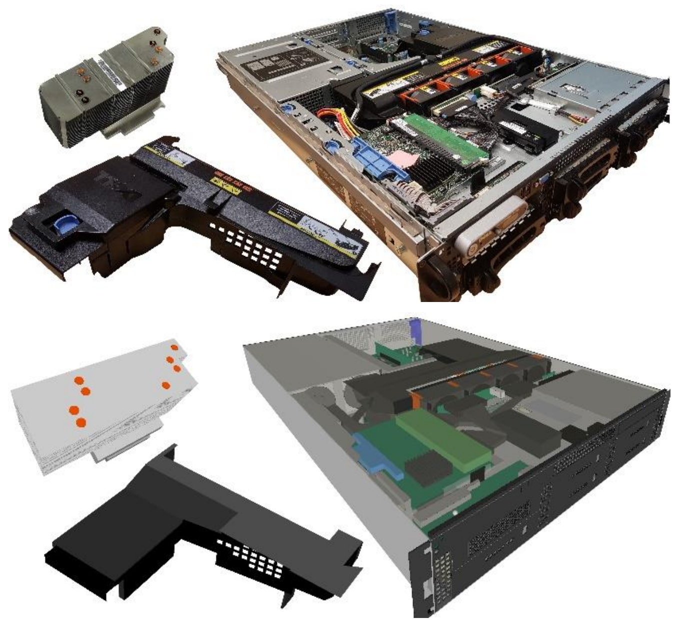

Fig. 4: Actual server and the detailed CFD model

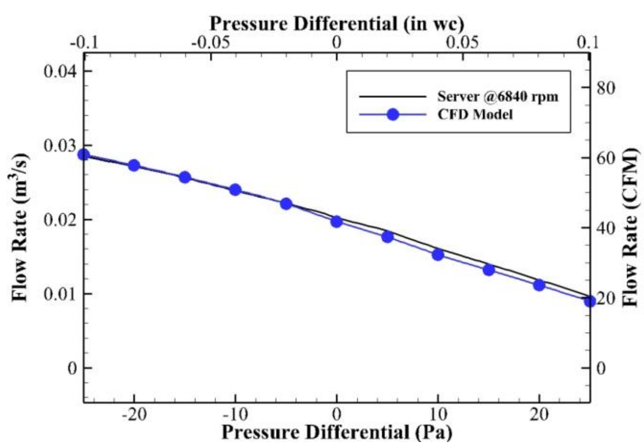

Fig. 5: Active flow curve for the actual server and the CFD model

Ideally, the summation of fans flow rates should be equal to the server's intake flow rate if sides of a chassis are sealed properly. However, CFD results revealed that 7% of the server airflow is supplied through the openings on the sides of the chassis (shown in Fig. 6) in case 1. These openings allow hot air and dust enter the chassis from the internal space of a cabinet that is open to a hot aisle. In addition, a portion of airflow returns to space upstream of the fans and recirculates through the vacant volume next to chassis fans as shown in Fig. 3. This reduces fresh air pick up from an adjacent cold aisle. To quantify the magnitude of this recirculation, server recirculation index (SRI) is defined as

$$\mathrm{SRI} = \frac{\sum_{i=1}^{4} Flow\ Rate_{\text{server fan \#i}} - Flow\ Rate_{\text{intake}}}{\sum_{i=1}^{4} Flow\ Rate_{\text{server fan \#i}}} \tag{1}$$

which is the ratio of the difference between the sum of the flow rate of chassis fans and flow rate of fresh air at the server's intake to the total flow rate of the fans as shown in Eq. (1). In this equation, fresh air is the portion of the server's intake flow rate that is received from a cold aisle in front of the server. A large SRI indicates that a significant portion of fans airflow is recirculated and not utilized efficiently. It should be noted that the flow rate of PSU fans is not included in Eq. (1) since only a portion of their airflow is fresh air from the server's intake and the rest comes from server fans in the current chassis design. Therefore, this equation understates the recirculation ratio in similar chassis designs and the actual value can be higher. Using CFD results, the SRI for cases 1 and 2 is calculated as 29% and 49%, respectively. This means that at least 29% of the flow rate of chassis fans in case 1 is either recirculated inside the server or supplied from the hot aisle. Hot air from a hot aisle can enter a server chassis either through openings on the sides of the chassis or due to reverse flow through PSUs. Figure 7 shows recirculation streamlines for cases 1 and 2. The expansion of recirculation with backpressure is clear in this figure which is consistent with flow visualization in Fig. 3.

Figure 8 shows the variation of SRI and temperature of HDD and SAS controller chip with backpressure in case 1. It is observed that SRI increases with backpressure and reaches 65% when backpressure is 25 Pa. Also, the temperature of the HDD and SAS chip increase significantly with backpressure. This is due to the expansion of the recirculation with backpressure which limits fresh airflow over the HDD and SAS controller chip. In addition, simulations showed that air temperature increases between 0.3 to 0.4°C as it absorbs heat from fans at their normal operating speed (~6840 rpm) which results in a higher temperature of other components downstream of fans.

Fans at the Rear of the Chassis

In this section, chassis fans are moved to the rear side of the chassis in the CFD model. Also, the rear panel of the server and the air baffle are redesigned to accommodate this change. The location and power of the rest of the components remain unchanged. In addition, openings on the side walls of the chassis are removed to improve air pickup through the server's front panel and avoid hot air leakage into the chassis. The developed CFD model is shown in Fig. 9. Modeling results show a 35% increase in fresh air pickup in case 3 compared to case 1 for a similar fans speed. Also, it is found that the speed of chassis fans can be lowered to 4920 rpm in case 3 while achieving a maximum CPU temperature similar to case 1 which indicates a 28% reduction in fans speed. This lower FS translates to 65% saving in chassis fan power. Additionally, a lower FS decreases the noise level of fans and increases its lifetime. Table 2 compares the results of this design with those of case 1 when PSU fans operated at the same speed. Furthermore, the difference between CPU temperatures is decreased significantly in case 3 due to the elimination of recirculation.

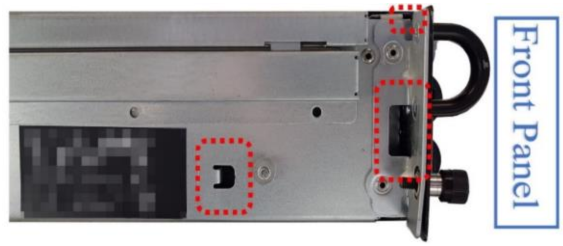

Fig. 6: Openings on the side of the chassis

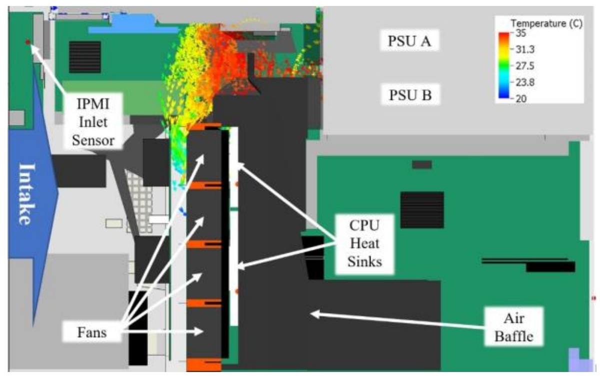

(a) No backpressure, $FS_{chassis}$= 6840 rpm

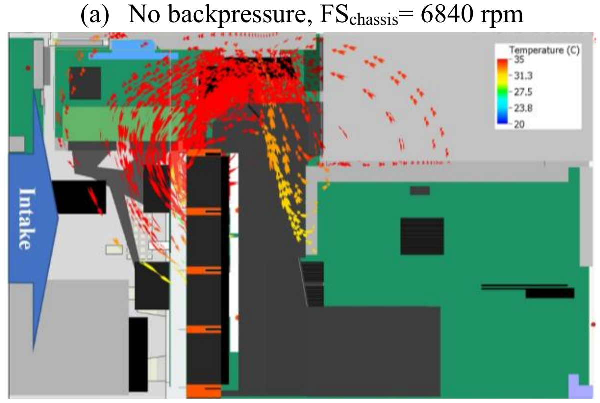

(b) Backpressure = 28 Pa, $FS_{chassis}$=12000 rpm

Fig. 7: Recirculation in the CFD model

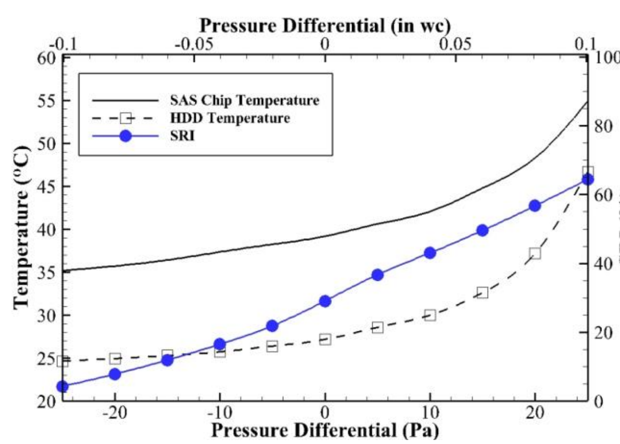

Fig. 8: Impact of backpressure on the temperature of HDD and SAS chip, and SRI

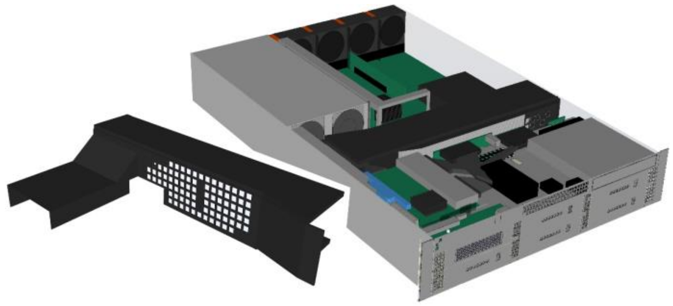

Fig. 9: CFD model of the server with fans moved to the rear

Table 2: Comparison of CFD results between new and original designs

	Case 1: Original Design	**Case 3:** Chassis with Fans at Rear	
FS $_{chassis}$	6840 rpm	6840 rpm	4920 rpm
$\dot{V}_{intake}$	43 cfm	58.3 cfm	43.5 cfm
T_{out}	30.6 ℃	29.9 ℃	33.9 ℃
$T_{in\text{-}IPMI}$	21 ℃	19.7 ℃	20.2 ℃
$T_{CPU\text{-}A}$	64.7 ℃	61.2 ℃	64.8 ℃
$T_{CPU\text{-}B}$	59.4	61.3 ℃	65.2 ℃
ΔT_{CPUs}	5.3 ℃	0.1 ℃	0.4 ℃

In case 3, the fans of the chassis and PSUs compete for air because of shared spaces up- and downstream of them. Generally, fan speed (FS) of a PSU is either fixed or controlled based on the temperature of its internal components. In addition, server fans are larger and can generate larger pressure differentials typically. Therefore, the generated negative pressure upstream of the chassis fans can limit airflow rate through power supplies and even lead to reverse flow through PSUs in designs that are similar to case 3. For instance, CFD results revealed a reverse flow through PSUs – as shown in Fig. 10 - when fans of the chassis and PSUs were operated at similar speeds to case 1. On the other hand, a positive flow was achieved through PSUs when chassis fans operated at 4920 rpm which demonstrates the impact of chassis fan speed on the flow rate through PSUs. Reverse flow through a PSU can increase SRI, create hot spots, cause overheating and decrease the lifetime of the PSU. It is worth mentioning that a reverse flow may not be the most critical situation for a PSU. In fact, zero or very low airflow is more dangerous than a reverse flow and can cause overheating rapidly. In addition, reaching a zero-airflow rate requires smaller backpressure compared to the required backpressure for reversing the flow. Implementing a wall that channels the intake of PSUs to the intake of the server, similar to the design proposed in [15], can minimize interaction between chassis fans and PSUs via segregating airflow paths corresponding to the chassis and PSU fans. However, complete segregation limits an effective utilization of internal space of the chassis or may not be possible due to the presence of PCBs, wiring and HDD bays in the front of the server. Instead, a control system can be utilized to balance airflow inside the chassis and assure sufficient cooling airflow for PSUs. To achieve this goal, the FS of PSUs should be controlled in a way that both chassis and PSU fans generate similar pressure differentials. This can be done by adjusting FS of PSUs through BMC when chassis fans speed changes. A method for finding an appropriate PSU FS for case 3 will be discussed in the next section.

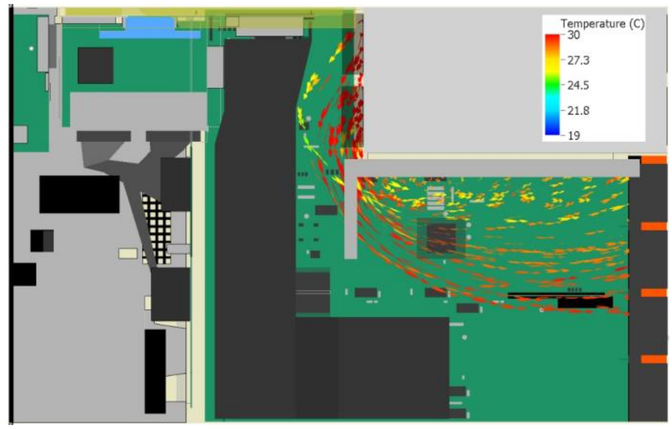

Fig. 10: Reverse flow through PSUs due to a high differential pressure created by chassis fans at 6840 rpm.

Airflow Balancing

In this section, an analytical method for balancing airflow inside a chassis is introduced and tested numerically. This method assures that the desired flow rate is delivered to PSUs for various given intake flow rates. The required flow rate for cooling a PSU under different workload can be determined experimentally or estimated via Eq. (2)

$$\dot{V} = Q\big/\left(\rho C_p \Delta T\right)_{air} \qquad (2)$$

in which Q is the generated heat, and ρ, C_p and ΔT are density, specific heat capacity and the average desired temperature rise of air, respectively. The required airflow rate for cooling the server can be estimated in a similar way. It should be noted that Eq. (2) only estimates the net flow rate required for dissipating the generated heat assuming uniform temperature at the inlet and exit of the chassis. Equation 2 does not provide any insight into the temperature of components. In addition, an internal recirculation or bypass of air over critical components may lead to overheated components even when sufficient air is supplied to the chassis according to Eq. (2). Therefore, it is recommended to utilize CFD simulation or experimental tests to find the actual required net flow rate for maintaining safe operating temperature for various components under different workloads which is out of the scope of this paper. Therefore, it is assumed that the required flow rate for cooling the PSUs and total flow rate of the server are known.

Air encounters resistance as it flows through a chassis. This flow impedance is proportional to the square of the stream velocity. An impedance curve for a chassis is typically parabolic and can be generated by plotting the static pressure of the air across the chassis against the airflow rate through it. The impedance curve for the chassis of case 3 is shown in Fig. 11 which is generated using a CFD model. This curve can be utilized for determining the required pressure differential for achieving a desired airflow rate through the chassis. In case 3, it is assumed that the chassis fans and PSUs operate in parallel. The total flow rate of parallel fans under equal pressure differential is equal to the sum of the flow rate of individual fans that operate under the same pressure differential. The equivalent flow curve for the system of fans in this design can be generated by adding the flow rate of the four chassis fans to the PSUs as shown in Fig. 12. If the impedance curve and the equivalent flow curve are superimposed, the point of intersection represents the actual operating point for the server (point A). Points B and C show operating points of chassis fans and PSUs. The flow rate of each PSU and chassis fan can be found by dividing the total airflow by the number of the corresponding devices in the graph. The accuracy of this method depends on the accuracy of the assumption that the chassis fans and PSUs operate in parallel. To investigate the accuracy of this assumption in case 3, three different scenarios are considered in which fans speed are calculated to achieve the respective desired PSU and server's intake flow rates. In the first scenario, 60 cfm and 4 cfm are considered as the server's intake and PSU desired flow rates, respectively. Therefore, a total of 8 cfm should be delivered by two PSUs when the chassis fans should deliver the remaining 52 cfm. The required pressure differential to achieve 60 cfm through the chassis is 64 Pa from the chassis impedance curve. Therefore, each PSU and chassis fan should generate 64 Pa for delivering the desired flow rates. With this information, the required operating speed of the fans can be back-calculated by implementing affinity fan laws and accurate curve fitting for the flow curves. The value of design parameters and results from the corresponding CFD model are presented in Table 3. This process is repeated for the desired intake flow rates of 30 cfm and 90 cfm in scenarios 2 and 3, respectively and the results are presented in Table 3. Different values for the desired PSU flow rate are chosen to examine the versatility of the method. The maximum difference between the desired values and CFD results in Table 3 is 5%. It should be noted that PSUs in this chassis are long and their intake is located far from the chassis fans, i.e. PSU fans are closer to the server intake. This weakens the assumption of parallel operating fans and results in a slightly higher flow rate than the corresponding desired flow rate. In addition, numerical errors in curve fittings and CFD simulations can be other causes of the difference. Overall, the results show successful implementation of airflow management via pressure balancing on this chassis. The operating points of the server, PSUs and chassis fans for scenario 1 are shown in Fig. 12.

Special attention should be given to the actual operating point of a fan to achieve reliable and predictable operation. In axial fans, a steep change of flow rate with pressure occurs at the middle range of the fan's flow curve. It is generally not advisable to operate fans in this area as a small change in pressure can cause a significant change in the fan flow rate. In fans with high blade angles, a stalling condition on the blade airfoil can occur in this region which is called the centrax point. Operating a fan at a lower flow rate compared to the centrax point may cause pressure surges that can lead to oscillations in the fan flow rate and undesirable flow rates. Therefore, the type and size of the fans should be chosen carefully. Generally, it is recommended to run the fan at lower pressure differentials to increase the fan's lifetime and avoid pressure surges during a system start-up. Finally, experimental tests should be carried to ensure proper cooling of the server components and verify the design.

Table 3: The desired and achieved flow rates with pressure balancing

	Scenario 1		Scenario 2		Scenario 3	
	Desired	CFD	Desired	CFD	Desired	CFD
$\dot{V}_{\text{intake}}$	60 cfm	57 cfm	30 cfm	29 cfm	90 cfm	86 cfm
$\dot{V}_{\text{PSU}}^{\text{avg}}$	6 cfm	6.2 cfm	3.5 cfm	3.6 cfm	10 cfm	10.5 cfm
FS_{PSU}	6070 rpm		4060 rpm		11040 rpm	
$\text{FS}_{\text{Chassis}}$	7410 rpm		3210 rpm		8790 rpm	

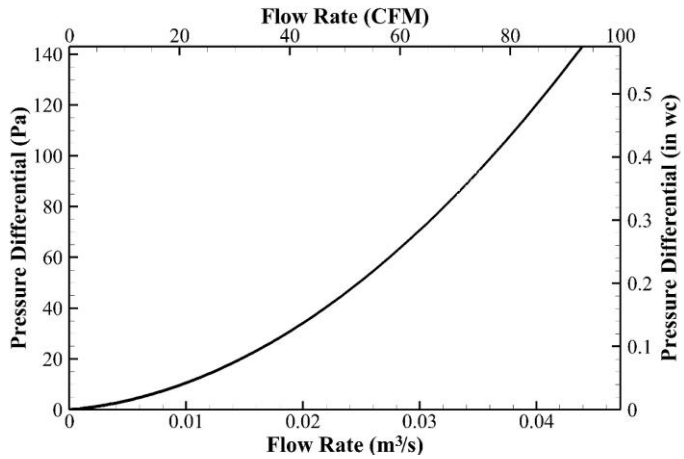

Fig. 11: Airflow impedance curve for the server with rear mounted fans

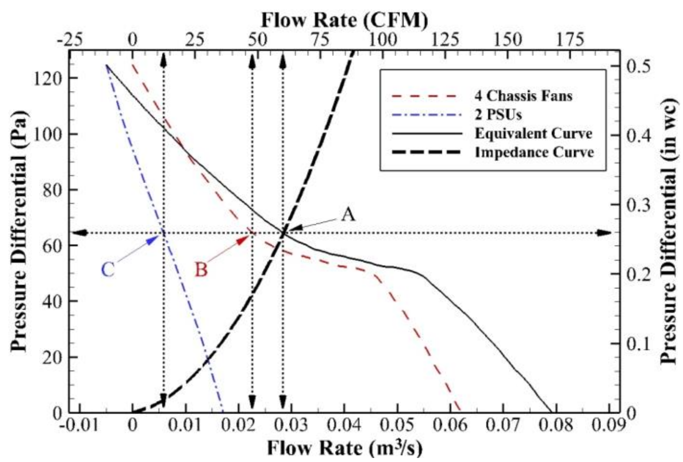

Fig. 12: Operating points of the server, PSUs and server fans in scenario 1 where points A, B and C represent the operating points for the server, the chassis fans and the PSUs respectively.

Lastly, in case 4, pressure balancing is implemented for the server in case 3. The results are compared with cases 1 and 3 in Table 4. Although the described method is tested numerically, experimental testing is required to ensure reliable operation of the airflow system.

It should be mentioned that the accuracy of this method depends on the accuracy of curve fittings and experimental flow curve of the fans when the assumption of parallel operating fans is valid. It is also noteworthy that the distance of the operating point of the fans from the stall region can affect the accuracy of the method. Alternatively, hot-pluggable fans similar to chassis fans can be installed downstream of fanless PSUs. As a result, the complexity of the airflow system decreases, and all fans can be operated at similar speeds.

Table 4: Airflow and thermal improvements due to the implementation of airflow balancing

	Case 1: Original design	**Case 3:** w/o airflow balancing	**Case 4:** w/ airflow balancing
$FS_{chassis}$	6840 rpm	6840 rpm	6840 rpm
FS_{PSU}	5000 rpm	5000 rpm	7410 rpm
$\dot{V}_{intake}$	43 cfm	58.3 cfm	61.9 cfm
$\dot{V}_{PSU}^{avg}$	7.2 cfm	-3.4 cfm	4.5 cfm
$T_{CPU\text{-}A}$	64.7 ℃	61.2 ℃	58.3 ℃
$T_{CPU\text{-}B}$	59.4	61.3 ℃	58.7 ℃
ΔT_{CPUs}	5.3 ℃	0.1 ℃	0.4 ℃

Conclusions

A detailed CFD model of a commercial 2-RU server was built, calibrated and experimentally validated to investigate the effects of implementing chassis fans at the rear side of the server. It was observed that an internal recirculation due to the fans arrangement in the chassis may occur and can decrease cooling efficiency and cause hot spots inside the server. The developed CFD model was used for investigating the benefits and consequences of moving the chassis fans to the rear of the server. The main advantage of this design is the elimination of internal recirculation in the server. It was shown that the power of chassis fans can be decreased by 65% via moving the fans to the rear of the server without any significant increase in the temperature of the main components. In addition to the energy savings, operating the fans at a lower speed reduces fan acoustic noise level which improves the working environment in data centers. Furthermore, dissipated heat by the fan's motor does not affect the temperature of other components. Although the new design offers multiple benefits, implementation of this design requires sealing of chassis sidewalls and careful airflow balancing inside the chassis to avoid overheated PSUs or a reverse flow through PSUs. A method for balancing airflow between chassis and PSU fans is described and tested successfully. Alternatively, hot-pluggable fans similar to chassis fans can be used downstream of fanless PSUs. Finally, it is worth mentioning that moving the server fans to the rear of the chassis limits the available space for rear access ports and PCI cards. However, this potential issue can be addressed by certain changes in the hardware (similar to OCP servers) or using extension cables to gather the ports at a designated area on the rear side of the server.

Acknowledgments

The authors would like to thank Prof. Kanad Ghose and Anuroop Desu from Binghamton University Computer Science and Data Center Group. 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-1738793.

References

[1] Ponemon Institute LLC, 2016, *2016 Cost of Data Center Outages*.

[2] Shehabi, A., Smith, S., Sartor, D., Brown, R., Herrlin, M., Koomey, J., Masanet, E., and Lintner, W., 2016, *United States Data Center Energy Usage Report*.

[3] Zhou, R., Wang, Z., Bash, C. E., and McReynolds, A., 2011, "Modeling and Control for Cooling Management of Data Centers with Hot Aisle Containment," ASME 2011 Int. Mech. Eng. Congr. Expo., pp. 739–746.

[4] ASHRAE Technical Committee 9.9, 2016, *Thermal Guidelines for Data Processing Environments, 4th Ed*, W. Stephen Comstock, Atlanta, GA, USA.

[5] Khalili, S., Alissa, H., Desu, A., Sammakia, B., and Ghose, K., 2018, "An Experimental Analysis of Hot Aisle Containment Systems," *Proceedings of the 17th InterSociety Conference on Thermal and Thermo-mechanical Phenomena in Electronic Systems, ITherm 2018*, San Diego, CA USA, pp. 748–760.

[6] Makwana, Y. U., Calder, A. R., and Shrivastava, S. K., 2014, "Benefits of Properly Sealing a Cold Aisle Containment System," *Thermomechanical Phenomena in Electronic Systems -Proceedings of the Intersociety Conference ITherm2014*, pp. 793–797.

[7] Khalili, S., Tradat, M., Nemati, K., Seymour, M., and Sammakia, B., 2018, "Impact of Tile Design on the Thermal Performance of Open and Enclosed Aisles," J. Electron. Packag., **140(1)**, pp. 010907-010907-12.

[8] Greenberg, S., Mills, E., Tschudi, B., Rumsey, P., and Myatt, B., 2006, "Best Practices for Data Centers: Results from Benchmarking 22 Data Centers," ACEEE. Proc. 2006 ACEEE Summer Study Energy Effic. Build.

[9] Lee, K.-P., and Chen, H.-L., 2013, "Analysis of Energy Saving Potential of Air-Side Free Cooling for Data Centers in Worldwide Climate Zones," Energy Build., **64**, pp. 103–112.

[10] Zhang, H., Shao, S., Xu, H., Zou, H., and Tian, C., 2014, "Free Cooling of Data Centers: A Review," Renew. Sustain. Energy Rev., **35**, pp. 171–182.

[11] Ohadi, M. M., Dessiatoun, S. V., Choo, K., Pecht, M., and Lawler, J. V., 2012, "A Comparison Analysis of Air, Liquid, and Two-Phase Cooling of Data Centers," Annu. IEEE Semicond. Therm. Meas. Manag. Symp., pp. 58–63.

[12] Capozzoli, A., and Primiceri, G., 2015, "Cooling Systems in Data Centers: State of Art and Emerging Technologies," Energy Procedia, **83**, pp. 484–493.

[13] Vertiv, *Energy Logic: Calculating and Prioritizing Your Data Center IT Efficiency Actions, White Paper*.

[14] Wang, Z., Bash, C., Tolia, N., Marwah, M., Zhu, X., and Ranganathan, P., 2009, "Optimal Fan Speed Control for Thermal Management of Servers," (43604), pp. 709–719.

[15] Frachtenberg, E., Heydari, A., Li, H., Michael, A., Na, J., Nisbet, A., and Sarti, P., 2011, "High-Efficiency Server Design," *Proceedings of 2011 International Conference for High Performance Computing, Networking, Storage and Analysis*, ACM, New York, NY, USA, p. 27:1-27:27.

[16] Mani, D., Fernandes, J., Eiland, R., Agonafer, D., and Mulay, V., 2015, "Improving Ducting to Increase Cooling Performance of High-End Web Servers Subjected to Significant Thermal Shadowing – an Experimental and Computational Study," *2015 31st Thermal Measurement, Modeling & Management Symposium (SEMI-THERM)*, pp. 319–323.

[17] Dhadve, M., Shah, J. M., and Agonafer, D., 2018, "CFD Simulation and Optimization of the Cooling of Open Compute Machine Learning 'Big Sur' Server," *2018 17th IEEE Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm)*, pp. 962–970.
[18] Vuckovic, M., and Depret, N., 2016, "Impacts of Local Cooling Technologies on Air Cooled Data Center Server Performance: Test Data Analysis of Heatsink, Direct Liquid Cooling and Passive 2-Phase Enhanced Air Cooling Based on Loop Heat Pipe," *2016 32nd Thermal Measurement, Modeling & Management Symposium (SEMI-THERM)*, pp. 71–80.
[19] Khalili, S., Alissa, H., Nemati, K., Seymour, M., Curtis, R., Moss, D., and Sammakia, B., 2018, "Impact of Internal Design on the Efficiency Of IT Equipment In A Hot Aisle Containment System - An Experimental Study," *ASME 2018 IPACK*, San Francisco, CA, p. V001T02A009.
[20] "OCP Server Working Projects" [Online]. Available: https://www.opencompute.org/wiki/Server/Working. [Accessed: 02-Jan-2019].
[21] "Server/Project Olympus" [Online]. Available: https://www.opencompute.org/wiki/Server/ProjectOlympus. [Accessed: 02-Jan-2019].
[22] "Open Rack" [Online]. Available: https://en.wikipedia.org/wiki/Open_Rack. [Accessed: 02-Jan-2019].
[23] Intel, "Intel® Xeon® Processor E5345" [Online]. Available: https://ark.intel.com/products/28032/Intel-Xeon-Processor-E5345-8M-Cache-2-33-GHz-1333-MHz-FSB-. [Accessed: 25-Nov-2018].
[24] CPU-World, "Intel Xeon E5345 Specifications" [Online]. Available: http://www.cpu-world.com/CPUs/Xeon/Intel-Xeon E5345 - HH80563QJ0538M (BX80563E5345A - BX80563E5345P).html. [Accessed: 25-Nov-2018].
[25] Nemati, K., Alissa, H. A., Murray, B. T., Sammakia, B. G., and Seymour, M., 2015, "Experimentally Validated Numerical Model of a Fully-Enclosed Hybrid Cooled Server Cabinet," *ASME InterPACK 2015*, p. V001T09A041.
[26] Alissa, H. A., Nemati, K., Sammakia, B., Seymour, M., Schneebeli, K., and Schmidt, R., 2015, "Experimental and Numerical Characterization of a Raised Floor Data Center Using Rapid Operational Flow Curves Model," *ASME InterPACK 2015*, American Society of Mechanical Engineers, p. V001T09A016.
[27] Great Internet Mersenne Prime Search, 2015, "Mprime (Prime95)" [Online]. Available: https://www.mersenne.org/download/.
[28] TMP75 Datasheet, "Texas Instruments" [Online]. Available: http://www.ti.com/lit/ds/symlink/tmp75.pdf.