

John Edward Fernandes

Mechanical and Aerospace Engineering Department,
University of Texas at Arlington,
P.O. Box 19023,
Arlington, TX 76013
e-mail: john.fernandes@mavs.uta.edu

Richard Eiland

Mechanical and Aerospace Engineering Department,
University of Texas at Arlington,
P.O. Box 19023,
Arlington, TX 76013
e-mail: richard.eiland@mavs.uta.edu

Bharath Nagendran

Amazon Lab, 126
Thermal Engineer Enterprise,
Sunnyvale, CA 94089
e-mail: bhanagen@amazon.com

Veerendra Mulay

Facebook Inc.,
Menlo Park, CA 425081
e-mail: vmulay@fb.com

Dereje Agonafer¹

Fellow ASME
Mechanical and Aerospace Engineering Department,
University of Texas at Arlington,
P.O. Box 19023,
Arlington, TX, 76013
e-mail: agonafer@uta.edu

Effectiveness of Rack-Level Fans—Part II: Control Strategies and System Redundancy

Fan efficiency is known to increase with size. In part I of this study, savings in server fan power on the order of 50% were reported by replacing server-enclosed 60 mm fans with a rear-mounted wall of larger fans (80 mm or 120 mm in size). A methodology for row-wise control of such rack-level fans, with the purpose of simulating an actual product, is previewed and savings comparable to part I are reported. Performance under real-life scenarios such as nonuniform computational loads and fan failure is investigated. Each rack-level setup has distinct advantages. Selecting between configurations would necessitate a compromise between efficiency, redundancy, and cost. [DOI: 10.1115/1.4038014]

Introduction

The growth and dependence of global commerce, social interaction, news sources, and other industries on information technology (IT) systems over the last decade has contributed to the rise of large data centers. These facilities are responsible for a significant portion of national and global energy consumption. Such are the implications that the national energy usage by data centers more than doubled between 2000 and 2005 and, it was projected that consumption would continue to rise over the course of the following five years [1]. By 2010, it was reported that data centers accounted for around 2% (between 1.7% and 2.2%) of the total national electricity consumption [2]. In the United States, this figure continues to grow with usage increasing by 8.7% between 2011 and 2012 while projected growth is expected to be around 9.8% over the following year. Electricity usage worldwide is reported to be around 1.8% or a corresponding power consumption of 322 TWh [3]. With this industry continuing to grow and strain the national electricity grid, there is a need to target energy savings within the data center.

Power usage effectiveness is a common metric used to gauge energy efficiency of operation of the facility and is the ratio of total facility power to IT equipment power. A recent survey [4] reported that the average power usage effectiveness is around 2.9 with only 20% of surveyed facilities recording a value of less than 2.0. Thus, majority of data centers in North America operate

extremely inefficiently with IT equipment accounting for less than half the total power consumption (support infrastructure responsible for majority of the remaining). It is, therefore, imperative to reduce data center power consumption and operating cost by increasing the efficiencies of power distribution and cooling systems. With a sizeable portion (around 30%) of typical data center power consumption attributed to cooling [5], which is categorized as a parasitic load, it has become vital that energy savings and efficiencies be pursued in these components at various levels within the data center facility [6].

In this study, focus is placed at the server and rack levels. Traditional servers are configured to include all sub-systems such as compute, memory, storage, networking, and cooling within a single chassis. Common rack-mount units generally have a low profile of 1 U ($U = 1.75$ in), which accommodate small 40 mm fans. Manufacturers [7] and standards associations [8] have published data that encourage designers to opt for larger fans to increase their peak total efficiency. However, since fans are generally selected based on server profile, there exist opportunities to instead consolidate larger fans at the rear of a rack to increase savings.

Over the past few years, original equipment manufacturers [9], semiconductor device manufacturers, and hyper-scale data center owners [10] have been promoting the concept of “rack disaggregation.” This refers to separation of resources or subsystems that are traditionally included in a server, into individual modules at the rack. This includes compute, storage, networking, power distribution, and cooling, and endeavors to make the rack the fundamental building block of a data center. Increased distance between IT components is countered by introduction of silicon photonics [11,12]. The primary advantage of such a deployment is

¹Corresponding author.

Contributed by the Electronic and Photonic Packaging Division of ASME for publication in the JOURNAL OF ELECTRONIC PACKAGING. Manuscript received February 19, 2017; final manuscript received September 21, 2017; published online October 27, 2017. Assoc. Editor: Pradip Dutta.

the ability to change or refresh subsystems at different frequencies. In addition, disaggregation promotes dematerialization [13], capable of significant environmental impact, through reduction in printed circuit board sizes and sheet metal otherwise used for server chassis. In particular, disaggregation of cooling at the rack is synonymous with the focus of this study.

Preliminary work [14] predicted savings of up to 55% in cooling power by replacing smaller, chassis enclosed (60 mm) fans with larger rack-mount units for a stack of four servers. The present studies (parts I and II) advance this work by experimentally validating maximum possible savings through deployment of 80 mm and 120 mm fans. This paper previews a methodology for implementation of a control system that replicates the in-built scheme for modulation of chassis fan speeds. Thus, with minor modifications, row-wise control of rack-level fans is executed with input from each server in the stack. Performance of larger fans under different rack loads and failure conditions is reported, and savings in power over the baseline configuration are quantified.

Experimental Setup and Procedures

Setup

Server Under Study. Figure 1 shows an intel-based open compute server [15,16] used in this study, similar to that employed in Ref. [17]. This 1.5 U rack-mount unit has two central processing units (CPUs), each with a rated thermal design power of 95W. Four 60 mm direct current fans are installed within the chassis to provide cooling to the motherboard and its critical components (CPU and memory). Both processors represent the principal heat load within the system and their temperatures drive a native algorithm that controls speed of the fans using a pulse width modulation (PWM) signal. It is to be noted that, as seen in Fig. 1, a sheet metal partition isolates the flow through the motherboard portion of the server from flow through the power supply unit (PSU) and hard drive seen at the right. In this study, focus is placed on the application of larger rack-mount fans for cooling the motherboard section only. Henceforth, for the sake of simplicity, figures depicting the 80 mm and 120 mm fan configurations will be illustrated without the PSU channel.

Test Setup. As previously discussed in Ref. [17], a stack of four servers is considered when evaluating the rack-level solutions. The rear of the stack provides an area of 330 mm \times 333 mm within which the larger fans must be accommodated. Figure 2 shows the




Fig. 1 Intel-based open compute server with air duct removed for visual purposes

fan wall installed at a distance of 25.4 mm from the rear of the stack for both 80 mm (nine units installed in a 3 \times 3 array) and 120 mm (four units installed in a 2 \times 2 array) cases. Table 1 lists specifications of the fans used. For a detailed description of this setup, please refer to part I of this study. To enable cross referencing between the two parts, the naming scheme is maintained with servers termed A–D from the bottom of the stack to the top and each row of fans similarly numbered 1–3 (1 and 2 for the 120 mm configuration). Since the focus of this study is to monitor cooling power consumption, the fans are powered externally as shown in Fig. 3. However, the internal PWM signals from each server are still used to control the fans and are delivered through a control circuit. The PWM signal component of the test setup will be explained in detail in the Controlling the Fans section. Tachometer output from each fan is logged using a data acquisition unit as well as returned to each server to prevent triggering of a failure scenario (running all remaining fans at full speed to prevent shutdown). It is imperative that the ground signal from each server and fan be shared between all monitoring and controlling equipment. Since the fans are not powered by the server, a power meter measures the rack (or stack) IT power consumption from a 277 VAC source. Together, the fan cooling power plus IT power represent the total power consumption of the system. A workstation communicates with all components in the setup and provides a common timestamp for effective data reduction. An ambient condition's logger records air temperature at the inlet to the server stack. Over the duration of testing, the inlet temperature is found to have a maximum variation of ± 1 °C with a mean of 25 °C.

Stressing the Servers. These rack-mount units are configured with CPU and memory resources to function as web servers. Applications usually deployed on these systems are found to utilize relatively more CPU as compared to memory. Table 2 outlines different simulated computational loads setup to run on these servers to mirror operation in a data center. Synthetic load generator *lookbusy* [18] is employed to create loads as previously outlined. A bash script, executed on each server at the start of a test, automates a procedure wherein each load is applied (in order) for a duration of 30 min with 30 min of idling between two loads (with the exception of the initial Idle load that runs for 60 min). This sequence of stressing the server from idling to maximum loads is repeated two more times before a test is concluded to ensure that the results obtained are repeatable. Native Linux tools *mpstat* and *free* are employed to measure CPU utilization and memory usage over the course of each test and ensure loads are per requirement. An internal diagnostic tool provided by the motherboard manufacturer provides readings from each CPU digital temperature sensor. Along with rack (IT) and fan power consumptions, these represent data of primary interest in this study. Steady-state operation is achieved within the first 20 min of each load and individual measurements made within the last 10 min are averaged before reporting. Results presented in this study represent averages over all three runs. Note that reported CPU temperatures, unless mentioned specifically, are the mean of individual readings from the four servers.

Controlling the Fans

Row-Wise Control. As previously stated, one of the primary objectives is to outline a methodology for row-wise control of larger rack-mount fans. By controlling each row independently based on adjacent server loads, cooling power consumption can be minimized as compared to a scheme where all fan speeds are equally modulated based on the maximum load across the stack. To enable such an arrangement, four PWM signals (one from each server) from the stack need to be converted to row-wise input; in this case, three signals for the 80 mm configuration or two for 120 mm counterpart. Contribution of each server's signal to a given row's input needs to be determined to minimize fan




Fig. 2 Simplified depiction of the fan wall installed behind the stack for both (a) 80 mm and (b) 120 mm cases with corresponding names for all primary components

Table 1 Specifications of fans used

Setup (mm)	Frame (mm)	Max. air flow (cfm)	Max. static pressure (in. aq.)	Rated speed (rpm)
60	60 × 60 × 25	37.1	0.62	7600
80	80 × 80 × 38	100.1	1.98	9500
120	120 × 120 × 25	171.0	0.90	5100

power consumption. A simple “zone of influence” test is carried out to determine these parameters.

Zone of Influence. Placement of the fan wall behind the stack (with a gap) ensures uniform flow across the rack when all fans are operated at the same speed. This manifests as near-uniform

CPU (or operating) temperatures when all servers are subjected to the same load. For further details, please refer to part I [17]. To study the effect of fan location on cooling across the rack, an individual row must be operated at a different speed compared to others. Variation in CPU temperatures will be indicative of influence each row has in terms of proximity. To enable such a

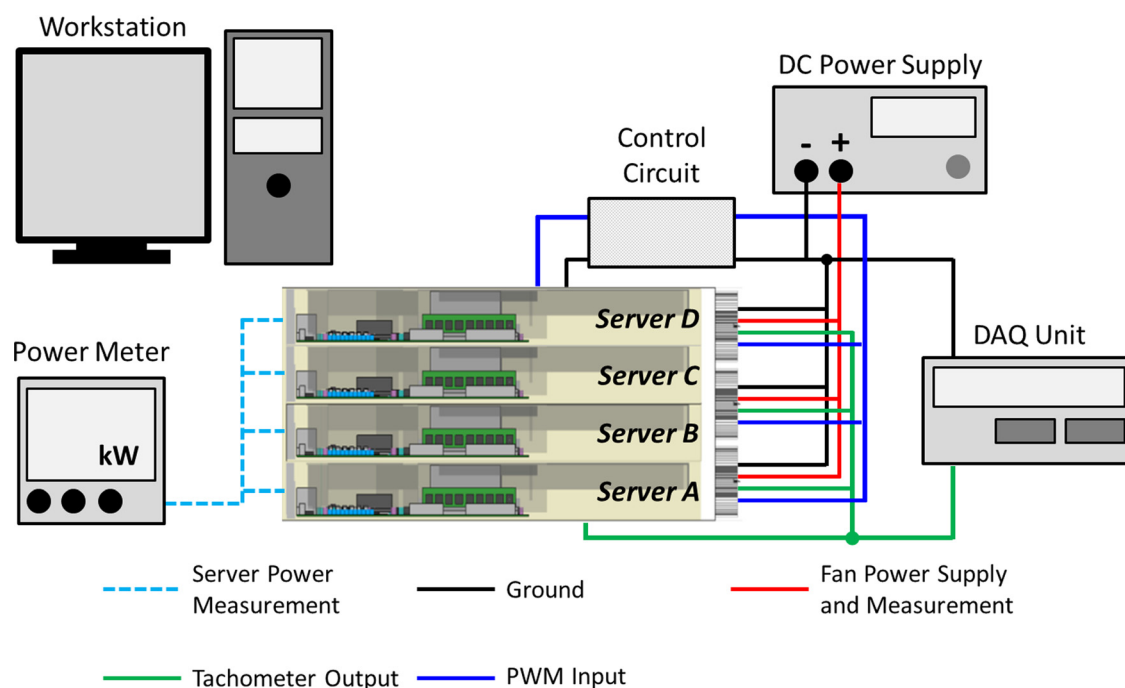


Fig. 3 Simplified depiction of the test setup with control and data acquisition equipment

Table 2 Different computational loads under which servers are stressed

Load (%)	U_{CPU} (%)	Memory usage (MB)
Idle	0–0.2	—
10	10	2000
30	30	2000
50	50	2000
70	70	2000
98	98	2000

comparison, a uniform load is applied to all servers and, in turn, an individual row of fans is operated at a higher duty cycle of 30% with the rest running at 10%. The distribution of operating temperatures across the stack for each test in the 80 mm configuration is shown in Fig. 4(a). For comparison, CPU temperatures for a similar load with all fans operating at 10% are also reported. As expected, it is clear that fans exert greater influence on servers in higher proximity. These CPU temperature readings from Fig. 4(a) must be converted to “weights” to meet cooling requirements for each server. Difference in operating temperatures between each row-wise test and uniform case (all at 10%) forms the basis to determining the aforementioned parameter. This overcooling term is calculated as follows:

$$\Delta T_{i,j} = T_{i,\text{uniform}} - T_{i,j} \quad (1)$$

where the subscript “ i ” corresponds to a given server and “ j ” to a row of fans running at 30% PWM. The $T_{i,\text{uniform}}$ term in the equation represents CPU temperature for server i under uniform fan duty cycles of 10%. The influence (%) can be determined by dividing the overcooling terms in each test by the sum of overcooling across the server stack as follows:

$$I_i = \frac{\Delta T_{i,j}}{\sum_{i=A}^D \Delta T_{i,j}} \quad (2)$$

where I_i is the influence on a given server by row j of fans operating at a higher speed. Influence values for all tests are reported in Fig. 4(b). In an ideal state, values for rows 1 and 3 should mirror each other and distribution for row 2 should be uniform (25% across the stack). However, each distribution is skewed toward server D. This is to be expected, as clearly shown by the CPU temperature variation when all fans are operated at 10% duty cycle in Fig. 4(a). These results provide an indication of the tolerances across the stack, specifically, in terms of difference in CPU

Table 3 Coefficients for each server's duty cycle signal

Setup (mm)	Fan row	Server			
		A	B	C	D
80	3	0.10	0.15	0.25	0.50
	2	0.25	0.25	0.25	0.25
	1	0.50	0.25	0.15	0.10
120	2	0.15	0.20	0.30	0.35
	1	0.35	0.30	0.20	0.15

powers. To counter this effect and to provide a scheme that would function across a multitude of servers in a data center, the weight of each server's PWM output with respect to a given fan row is based on the influences reported in Fig. 4(b) and generalized to provide a mirrored form as listed in Table 3. These coefficients are inputs to the control circuit to process each PWM input from the stack to a row-wise duty cycle output. The Control Circuit and Operation section will describe the control circuit and system in detail.

Control Circuit and Operation. A detailed diagram of the control circuit, using Ref. [19], can be seen in Fig. 5. Each server's output signal is converted to an analog voltage through a low-pass filter. This conversion simplifies the circuit as an analog voltage signal is relatively easier to read in comparison to a PWM input. Each filter consists of a 0.1 μF capacitor and two 33 k Ω resistors in series to provide a cutoff frequency of around 25 Hz. This value is three orders of magnitude lower than the PWM frequency [20] and provides required response as the input frequency does not vary. Analog voltage output ($V_{\text{out},i}$) of the filter is correlated to the duty cycle (DC_i , %) of the input signal as follows:

$$V_{\text{out},i} = \text{DC}_i \times 5 \quad (3)$$

where 5 VDC corresponds to the logic high of the input signal. Analog input ports in an Arduino microcontroller [21] board are used to measure the processed PWM signal from each server. A LabVIEW [22] program reads input to the board and generates PWM output based on the coefficients in Table 3 as follows:

$$\text{DC}_j = \sum_{i=A}^D C_{i,j} \frac{V_{\text{out},i}}{5} \quad (4)$$

where DC_j is the duty cycle signal to fan row j and $C_{i,j}$ is the weight of server i corresponding to row j . Thus, the control system as described in the Controlling the Fans section, is implemented for both 80 mm and 120 mm cases.

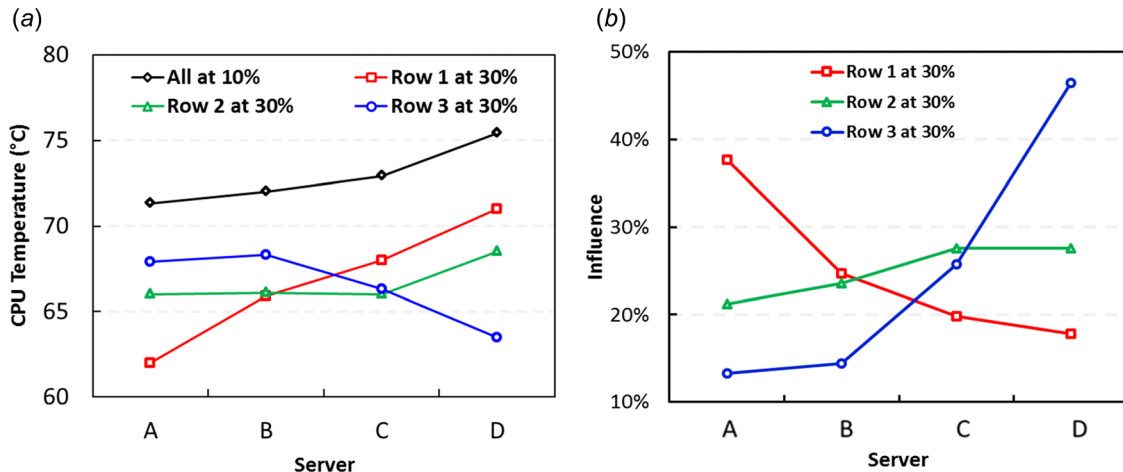


Fig. 4 (a) Operating temperatures under 98% load when fan rows are individually run at higher speeds and (b) influence of each fan row on server cooling; used to select coefficients for the control scheme

Need for a Lower Bound. It is critical to remember that to make comparisons between the baseline and rack-level configurations in terms of cooling power consumption, other critical parameters (average CPU temperature and IT power) must be held constant. This is the basis of part I of the study wherein maximum available savings for a fixed operating temperature are reported. However, when a control system based on server inputs is implemented, achieving an exact target across all setups is near impossible. Therefore, to provide reasonable comparison between configurations, the 80 mm and 120 mm control schemes are modified to ensure that CPU temperatures and rack IT power are either equal to or below values reported for the baseline case (60 mm). This way, savings in cooling power can be reported without need to discuss remaining critical parameters. As seen in Fig. 6, for the 80 mm configuration, a 7.5% lower bound in the PWM output signals is introduced to meet requirements. The original control scheme outlined in Eq. (4) provides greater savings in cooling power as compared to the final. However, higher CPU temperatures cause a corresponding increase in IT power through the effect of leakage current. The control system is modified to account for the lower bound as follows:

$$DC_j = 0.075 + 0.925 \times \sum_{i=A}^D C_{ij} \frac{V_{out,i}}{5} \quad (5)$$

The 120 mm fans are inherently designed to produce no change in speed between 10% and idling duty cycles. To account for this and ensure that the output to each row of fans does not lag behind input signals from the rack, the control scheme is modified to account for the lower bound as follows:

$$DC_j = 0.1 + 0.9 \times \sum_{i=A}^D C_{ij} \frac{V_{out,i}}{5} \quad (6)$$

No further modification to this system is required as the fans are found to overcool the servers while idling at 10% PWM irrespective of the load as seen in Fig. 7. Thus, settings for each control system (80 mm and 120 mm cases) are finalized and deployed for further testing and evaluation.

Results and Discussion

The objective of this study is to confirm that substantial savings are available when 60 mm chassis fans are replaced by rack-level configurations with effective control schemes that simulate a product or solution deployed within a data center facility. It is imperative that, in addition to identical utilizations seen in Fig. 7, realistic conditions such as nonuniform workloads and fan failure scenarios must be investigated to reinforce the merits of employing larger fans.

Comparison Under Uniform Loading. It has been shown that despite overcooling the servers with rack-level setups, under uniform utilizations, savings in cooling power are available irrespective of the load. Table 4 summarizes the extent of reduction in fan power available through deployment of 80 mm and 120 mm fans with their designed control method. In either case, it is apparent that savings are maximized when the stack is operated at high workloads. The reason for reduced efficiency at utilizations at or below 30% is because 60 mm fans operate at idling speeds in such conditions. In comparison to the baseline, the 80 mm case reports savings of around 45–51%. Similarly, the 120 mm configuration provides reduction in cooling power of the order of 33–50%. This setup provides reduced savings at lower loads owing to the fact that these fans operate at idling speeds (10% lower bound) across the entire test spectrum and significantly overcool the stack for utilizations below 70%.

Comparisons must be made with results from part I to ensure that the deployed control systems do not deviate significantly from reported maximum available savings (by comparing average CPU temperatures). For the 80 mm case, a maximum deviation of 5% is observed when compared to projected savings of 50–53%. While this reduction can be considered acceptable, it exists because of overcooling at loads between idling to 30% as seen in Fig. 7(b). However, for the 120 mm case, a reduction of 15% is observed in comparison to expected savings of 48–54%. This can be attributed to the fact that in part I, comparisons between the baseline and 120 mm configurations are made with respect to CPU temperatures reported while testing the latter case. It is, therefore, unfair to draw comparisons between results from part I and this study for the 120 mm setup.




Fig. 5 Duty cycle signal from each server is converted into an analog voltage by a low-pass filter and read by the microcontroller board. These inputs are processed by a LabVIEW program and a PWM output is sent to each fan row.




Fig. 6 Comparisons between (a) cooling power, (b) average CPU operating temperature, and (c) IT power demonstrate that a 7.5% lower bound (LB) in 80 mm fan duty cycle is required to ensure that adequate cooling is provided when compared to the 60 mm baseline




Fig. 7 Comparison of (a) cooling power, (b) average operating temperature of the stack, and (c) stack IT power consumption for all final configurations under uniform computational loads

When comparing both rack-level configurations from the perspective of cooling power consumption, it would be understandable to claim that the 80 mm case is more efficient. However, as seen in Fig. 7(c), it is evident that reduction in IT power consumption is available due to overcooling and lower leakage power. Therefore, comparisons can also be made in terms of total power consumption (cooling + IT) as outlined in Table 5. In this case, it is apparent that the benefits of reduced IT power more than compensate for lower savings in fan power. However, it must be noted that increased total savings through overcooling is accompanied by an increase in air flow rates through the stack, which raises cooling power consumption at the facility level. Therefore, attention is focused at rack-level fan performance only.

Comparison Under Nonuniform Load. Since it would be time and resource intensive to test each configuration under all possible workloads observed in a data center facility, a simple test

was conducted to prove that rack-level fans are more efficient than the baseline under nonuniform loads. This involved comparing the 60 mm setup when all servers are idling (lowest fan power consumption) with the larger fan cases when only one server in the stack is operating at maximum utilization. The latter represents an extreme case of nonuniform loading, and server D was chosen as it consistently reported CPU temperatures higher than the remaining units (see Fig. 4(a)). Thus, by reporting savings under such conditions, it would be implicit that rack-level fans under study are superior under nonuniform loads.

Table 6 outlines the results from these tests. It is observed that both control schemes provide substantial savings (around 35%) in fan power when compared to the baseline. In addition, these results support the need for row-wise control as, for the 80 mm case, 5% reduction in fan power is observed when compared to a similar test conducted in part I; wherein all fans are controlled by server D. Since the fan control does not engage for the 120 mm

Table 4 Savings in cooling power under uniform loads for the server stack

Load (%)	P_{fan} (W)			Savings, W (%)	
	60 mm	80 mm	120 mm	80 mm	120 mm
Idle	9.59	5.26	6.43	4.33 (45.2)	3.16 (33.0)
10	9.58	5.23	6.39	4.35 (45.4)	3.19 (33.3)
30	9.56	5.21	6.33	4.35 (45.5)	3.23 (33.8)
50	11.19	5.47	6.27	5.72 (51.1)	4.92 (44.0)
70	11.67	5.73	6.27	5.94 (50.9)	5.40 (46.3)
98	12.66	6.41	6.26	6.25 (49.4)	6.40 (50.6)

Table 5 Total savings (IT + cooling) under uniform loads for the server stack

Load (%)	P_{total} (W)			Savings (W)	
	60 mm	80 mm	120 mm	80 mm	120 mm
Idle	348.5	343.0	342.7	5.53	5.74
10	455.9	450.6	448.1	5.25	7.75
30	696.9	682.5	672.7	14.35	24.18
50	841.9	837.1	826.1	4.82	15.81
70	878.4	872.7	867.9	5.64	10.50
98	894.1	886.1	883.6	7.95	10.50

Table 6 Savings under a non-uniform load

Configuration (mm)	Server loads	P_{fan} (W)	Savings (%)
60	All servers idling	9.59	—
80	Servers A–C idling; D at 98% load	5.87	38.8
120	Servers A–C idling; D at 98% load	6.33	34.0




Fig. 8 Locations of simulated individual fan failures for (a) 60 mm, (b) 80 mm, and (c) 120 mm configurations

configuration, similar savings are reported in both part I and this study.

Fan Failure Study. Along with other server components such as hard drives and memory, failure of fans commonly occurs in data center facilities. To account for the same, thermal engineers are required to design cooling systems for servers that ensure uptime even when a fan fails. Therefore, it is imperative that all configurations under study are tested while simulating failure to ensure no detriment to performance of the system such as throttling, increase in power consumption, etc. Figure 8 provides an illustration of locations of single fan failures for all configurations. The diagonal pattern is chosen to reduce testing time as power consumption under failure is found to be independent of location [14]. For the rack-level setups, preliminary trials were found to be in agreement as well. For each configuration, results from all tests are averaged and reported in Table 7. It is observed that failure of a single fan has a marginal effect on power consumption for the baseline and 80 mm cases. However, for the 120 mm setup, a 33% increase in fan power is observed under simulated failure. This can be attributed to lower available redundancy caused by the restriction in available area at the rear of the stack that limits the number of units to four. It is important to note that, while the difference in performance between 80 mm and 120 mm cases is apparent, the frequency of failures and time between failure and replacement are equally important when making a decision between configurations.

Further Discussion. As previously discussed, multiple factors must be considered before drawing conclusions from a fan failure study. One such statistic is the life expectancy of selected fans. Manufacturers publish L_{10} (time at which 10% of the fans will fail) and mean time to failure along with specifications to aid selection based on product lifecycle. Based on operating conditions of the fan, life expectancy can be scaled up or down using an

acceleration factor of 1.5 for every 10 °C change in temperature [23]

$$AF = 1.5^{[(T_{test} - T_{use})/10]} \quad (7)$$

where T_{test} is the temperature at which manufacturers conduct tests to report data and T_{use} is the temperature at which the fan will be operated in a system. Data centers are mission critical facilities and, as such, component failure rates of 10% may be considered unacceptable. Therefore, L_{10} values reported by manufacturers need to be scaled to more stringent requirements such as L_1 (time at which 1% of the fans will fail) [24].

$$L_{10} = L_1 \times \left(\frac{\ln(0.99)}{\ln(0.90)} \right)^{1/\beta} \quad (8)$$

where β is the Weibull shape parameter and chosen to be three in this case [23,25]. L_1 values for all three fans as a function of operating temperature are plotted in Fig. 9. Over 80 mm fans have 22.5% greater life expectancy than 60 mm and 120 mm counterparts. However, this does not translate to lower failures than the 120 mm units. Considering there are more than twice the number of parts in the 80 mm fan wall, more failures would be expected




Fig. 9 Life expectancy of chosen fans as a function of operating temperature

Table 7 Performance under simulated fan failure

Test condition	P_{fan} (W)		
	60 mm	80 mm	120 mm
No failures	12.66	6.41	6.26
One fan failed	12.81	6.23	8.33
Penalty (%)	1.19	−2.81	33.07

due to the sheer difference in quantity. In addition, since 120 mm fans operate at a fixed duty cycle (at ambient temperatures of 25 °C or below), and changes in speed are detrimental to fan life, lower failure rates can be expected.

Previous discussions in the results section have made valid arguments for choosing either rack-level configuration over the baseline case. However, making a selection between the larger fans is not straightforward. A decision between configurations would require consideration of advantages outlined herein as well as factors that exist beyond the scope of this study. Accommodation of all parameters would necessitate a study of total cost of ownership. Setup and results from this study would factor in both capital and operational expenditure, based on which a decision could be made. Regardless, based on reported observations, maximizing savings through deployment of rack-level fans can be achieved through a compromise between efficiency, redundancy, and cost. These terms are ultimately dependent on the size of fans selected.

Conclusions

Configurations of larger fans setup to mirror a product that could be deployed in a data center facility have been shown to provide greater efficiency in cooling web servers traditionally configured to use smaller (60 mm) chassis fans. In union with [14] and part I [17] of this study, a detailed methodology that outlines selection of fans, prediction and validation of savings, and setup for server-dependent operation has been previewed. A control scheme was implemented for both 80 mm and 120 mm cases that delivered savings in cooling power consumption of the order of 45–51% and 33–50%, respectively, under uniform rack load. Testing both rack-level setups under highly nonuniform load reports savings of around 35% when compared to baseline (60 mm) fan power for idling utilization. Simulation of fan failure showed marginal penalties in performance for both 60 mm and 80 mm cases. However, a 33% increase in cooling power was observed for the 120 mm configuration and attributed to lack of available redundancy (fan count). A summary of advantages in support of either rack-level setup is included below.

Advantages of 80 mm Configuration

- Consistent savings in fan power across a spectrum of loads
- No penalty or increase in cooling power for single fan failure
- 22.5% greater life expectancy per fan
- Less overcooling corresponds to lower air flow requirement, which influences facility-level cooling power consumption.

Advantages of 120 mm Configuration

- Greater overcooling gives rise to decrease in IT power consumption due to reduced leakage power. These savings more than compensate for relatively higher fan power at lower loads.
- At ambient temperatures of 25 °C and below, fans will idle (10% duty cycle). Lack of change in speed will increase life expectancy.
- With less than half the number of fans (4) in the wall, lower number of total failures is expected.

Analysis of total cost of ownership would be required to make an informed decision between the two rack-level solutions. A final selection would represent a compromise between efficiency, redundancy, and cost.

Acknowledgment

The authors thank Jacob Na for providing information on life expectancy for the selected fans. Any opinions, findings, and conclusions or recommendations expressed in this material are those

of the author(s) and do not necessarily reflect the views of the National Science Foundation.

Funding Data

- Directorate for Engineering (Grant No. 1134821).

Nomenclature

C	= server coefficient or weight
i	= server name, A to D
I	= influence, %
j	= fan row number, 1 to 3
L_1	= time at which 1% of fans will fail, hours
L_{10}	= time at which 10% of fans will fail, hours
T	= CPU temperature, °C
T_{test}	= temperature at which fine life test is conducted, °C
T_{use}	= air temperature at inlet side of fan, °C
V_{out}	= DC analog output voltage from low-pass filter, V

References

- [1] Brown, R. E., Brown, R., Masanet, E., Nordman, B., Tschudi, B., Shehabi, A., Stanley, J., Koomey, J., Sartor, D., Chan, P., and Loper, J., 2007, "Report to Congress on Server and Data Center Energy Efficiency: Public Law 109-431," Ernest Orlando Lawrence Berkeley National Laboratory, Berkeley, CA, Report No. **LBLN-363E**.
- [2] Koomey, J., 2011, "Growth in Data Center Electricity Use 2005 to 2010," Analytics Press, Berkeley, CA.
- [3] Data Center Dynamics, 2012, "Global Data Center Energy Use to Grow by 19% in 2012," Data Centre Dynamics Ltd., London, UK, accessed Dec. 2, 2014, <http://www.datacenterdynamics.com/content-tracks/powercooling/global-data-center-energy-use-to-grow-by-19-in-2012/33006.fullarticle>
- [4] Digital Realty, 2013, "Digital Realty: North America Campos Survey Results," Digital Realty Trust Inc., San Francisco, CA, accessed Oct. 6, 2017, https://na6.salesforce.com/sfc/p/300000005uRq/a/80000000CpC7/k_RJOcvs31zvPC4hgEz9NMQjNd0m4KjS_CzGO5_ni48
- [5] Pelley, S., Meisner, D., Wenisch, T. F., and VanGilder, J. W., 2009, "Understanding and Abstracting Total Data Center Power," Workshop on Energy-Efficient Design (**WEED**), Austin, TX, June 20.
- [6] Alkharabsheh, S., Fernandes, J., Gebrehiwot, B., Agonafer, D., Ghose, K., Ortega, A., Joshi, Y., and Sammakia, B., 2015, "A Brief Overview of Recent Developments in Thermal Management in Data Centers," **ASME J. Electron. Packag.**, **137**(4), p. 040801.
- [7] Strike, N. D., 2015, "Fan Efficiency, an Increasingly Important Selection Criteria," NMB Technologies Corporation, Tempe, AZ, accessed Jan 23, 2015, http://www.nmbtc.com/fans/white-papers/fan_efficiency_important_selection_criteria/
- [8] Brendel, M., 2012, "Fan Efficiency Grade Classification for Fans," **ASHRAE Annual Meeting**, San Antonio, TX, June 23–27.
- [9] Arlitt, M., Bash, C., Blagoderov, S., Chen, Y., Christian, T., Gmach, D., Hyser, C., Kumari, N., Liu, Z., Marwah, M., McReynolds, A., Patel, C., Shah, A., Wang, Z., and Zhou, R., 2012, "Towards the Design and Operation of Net-Zero Energy Data Centers," 13th IEEE Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems (**ITherm**), San Diego, CA, May 30–June 1.
- [10] Teich, P., and Moorhead, P., 2013, "Intel's Disaggregated Server Rack," Moor Insights & Strategy, accessed Jan. 23, 2017, <http://www.moorinsightsstrategy.com/research-note-intels-disaggregated-server-rack/>
- [11] Han, S., Egi, N., Panda, A., Ratnasamy, S., and Shi, G., and Shenker, S., 2013, "Network Support for Resource Disaggregation in Next-Generation Datacenters," 12th ACM Workshop on Hot Topics in Networks (**HotNets-XII**), College Park, MD, Nov. 21–22.
- [12] Weiss, J., Dangel, R., Hofrichter, J., Horst, F., Jubin, D., Meier, N., La Porta, A., and Offrein, B. J., 2014, "Optical Interconnects for Disaggregated Resources in Future Datacenters," European Conference on Optical Communication (**ECOC**), Cannes, France, Sept. 21–25.
- [13] Chang, J., Meza, J., Ranganathan, P., Shah, A., Shih, R., and Bash, C., 2012, "Totally Green: Evaluating and Designing Servers for Lifecycle Environmental Impact," 17th International Conference on Architectural Support for Programming Languages and Operating Systems (**ASPLOS XVII**), London, UK, Mar. 3–7, pp. 25–36.
- [14] Nagendran, B., Nagaraj, S., Fernandes, J., Eiland, R., Agonafer, D., and Mulay, V., 2014, "Improving Cooling Efficiency of Servers by Replacing Smaller Chassis Enclosed Fans With Larger Rack-Mount Fans," IEEE Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems (**ITherm**), Orlando, FL, May 27–30.
- [15] Furuta, S., 2011, "Server Chassis and Triplet Hardware v1.0," Open Compute Project, accessed Jan. 23, 2015, <http://files.opencompute.org/oc/public.php?service=files&t=97a085fe84e16d7d2e950d883838c9ba>
- [16] Frachtenberg, E., Heydari, A., Li, H., Michael, A., Na, J., Nisbet, A., and Sarti, P., 2011, "High-Efficiency Server Design," International Conference for High

- Performance Computing, Networking, Storage and Analysis (SC), Seattle, WA, Nov. 12–18.
- [17] Eiland, R., Fernandes, J., Nagendran, B., Agonafer, D., and Mulay, V., 2017, “Effectiveness of Rack-Level Fans—Part I: Energy Savings Through Consolidation,” *ASME J. Electron. Packag.*, epub.
 - [18] Carraway, D., 2015, “Lookbusy—A Synthetic Load Generator,” accessed Jan. 23, 2015, <http://www.devin.com/lookbusy/>
 - [19] Knorig, A., and Howell, B., 2010, “Advanced Prototyping With Fritzing,” Fourth International Conference on Tangible, Embedded, and Embodied Interaction (TEL), Cambridge, MA, Jan. 24–27.
 - [20] Intel, 2004, “4-Wire Pulse Width Modulation (PWM) Controlled Fans,” Intel Corporation, Santa Clara, CA, accessed Jan. 23, 2015, http://www.formfactors.org/developer%5Cspecs%5Crev1_2_public.pdf
 - [21] Arduino, 2015, “Arduino Mega 2560,” Arduino, accessed Jan. 23, 2015, <http://arduino.cc/en/Main/ArduinoBoardMega2560>
 - [22] National Instruments, 2015, “LabVIEW Version 15.0.0,” National Instruments, Austin, TX, accessed Jan. 23, 2015, <http://www.ni.com/labview/>
 - [23] Tian, X., 2006, “Cooling Fan Reliability: Failure Criteria, Accelerated Life Testing, Modeling and Qualification,” Annual Reliability and Maintainability Symposium (RAMS), Newport Beach, CA, Jan. 23–26.
 - [24] Sidharth, S., and Sundaram, S., 2004, “A Methodology to Assess Microprocessor Fan Reliability,” Ninth Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems (ITHERM), Las Vegas, NV, June 1–4.
 - [25] Kim, S., Vallarino, C., and Claassen, A., 1996, “Review of Fan Life Evaluation Procedures,” *Int. J. Reliab. Qual. Saf. Eng.*, **3**(1), pp. 75–96.