

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

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

Bharath Nagendran

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

Veerendra Mulay

Facebook Inc.,
Senior R&D Engineer,
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 I: Energy Savings Through Consolidation

In general, smaller fans operate at lower efficiencies than larger fans of proportional linear dimensions. In this work, the applicability of replacing smaller, 60 mm baseline fans from within the chassis of web servers with an array of larger, geometrically proportional 80 mm and 120 mm fans consolidated to the back of a rack is experimentally tested. Initial characterization of the selected fans showed that the larger fans operate at double peak total efficiency of the smaller fans. A stack of four servers were used in a laboratory setting to represent a rack of servers. When all four servers were stressed at uniform computational loadings, the 80 mm fans resulted in 50.1–52.6% reduction in total rack fan power compared to the baseline fans. The 120 mm fans showed similar reduction in rack fan power of 47.6–54.0% over the baseline. Since actual data centers rarely operate at uniform computational loading across servers in a rack, a worst case scenario test was conceived in which the array of larger fans were controlled by a single server operating at peak computational workload while the other three in the rack remained idle. Despite significant overcooling in the three idle servers, the 80 mm and 120 mm fan configurations still showed 35% and 34% reduction in total rack fan power compared to the baseline fans. The findings strongly suggest that a rack-level fan scheme in which servers share airflow from an array of consolidated larger fans is superior to traditional chassis fans. [DOI: 10.1115/1.4038235]

Introduction

Significant improvements have been made to the operation of air movers within information technology (IT) equipment over the last decade. In typical server systems built prior to 2005, internal server fans frequently consumed 10–20% of the total server power draw. Advances in fan designs, fan speed control (FSC) with the use of pulse width modulation (PWM) signaling, and improved overall server thermal designs have since brought this number down to 2–4% in typical operation [1,2]. However, as with the data center facility in general, cooling power represents a parasitic load that does not contribute to the useful computational work output of a server and must be minimized or eliminated if possible. Typical servers will utilize axial, dual counter rotating, or centrifugal blower type fans as the primary air moving devices within their chassis.

In general, larger fans of geometrically similar proportions operate with higher efficiencies than their smaller counterparts. Various fan manufacturers are diligent in educating customers on this general principle [3,4]. In fan efficiency rating standard published by Air Movement and Control Association, it is shown that below a threshold limit of impeller diameter peak total efficiency is less for smaller diameter fans [5]. The typical peak operating efficiency will occur at roughly one-third of the maximum static pressure a fan can deliver. Proper system design and the use of

impedance matching techniques can help optimize a system's performance to realize the best possible efficiencies. Work by Holahan and Elison outline the governing relations between sound, flow, and pressure ratios for fans of homologous dimensions and discuss the principles of impedance matching [6]. They also performed a survey of peak total operating efficiency as a function of frame size for fans typically found in rack systems, which illustrates the trend of increased efficiency with size. Kodama et al. mapped the temperature response of computing nodes to the fan speeds in a fan array of a blade chassis [7]. This work highlighted possible imbalances in central processing unit (CPU) temperatures and system efficiency that may exist when a consolidated fan array scheme is not optimized.

A previous study by the authors of this work showed computational fluid dynamics evidence based on empirical data that significant savings in fan energy are possible through consolidation of smaller, internal server fans to a larger rack-based fan configuration [8]. Additional details of this previous study included findings that a failure of a single fan out of the four in the server under study has the same impact on die temperatures, irrespective of the location of the failed fan. These initial results provide a foundation and motivation for the present work.

The purpose of this work is to establish the feasibility and possible savings achievable by moving from smaller, chassis enclosed fans, to larger, more efficient fans at the rack level of an actual server system. To achieve this end, a simulated rack of four servers is set up in a laboratory environment. Three different axial fans sizes are experimentally tested to determine their performance over a range of computational operating conditions. This includes the baseline 60 mm fans original to the servers and two

¹Corresponding author.

Contributed by the Electronic and Photonic Packaging Division of ASME for publication in the JOURNAL OF ELECTRONIC PACKAGING. Manuscript received February 16, 2017; final manuscript received October 10, 2017; published online October 27, 2017. Assoc. Editor: Pradip Dutta.




Fig. 1 Intel-based Open Compute server used for this study. The partition separates airflow of the motherboard fans from that of the power supply (PSU) fan air stream, simplifying the geometry for the analysis of this work.

larger size fans, 80 mm and 120 mm, moved outside of the server chassis and fixed to a rack-level frame. In this manner, when a consolidated larger fan array is used, fans will be responsible for providing cooling airflow to multiple servers in the system as opposed to the baseline 60 mm fan case in which each server has its own independent air stream. Testing over a nonuniform server computational operating loads (heat loads) further highlights the feasibility of the proposed modification in a more representative data center scenario.

Experimental Setup and Procedures

Setup

Server Under Study. For the present study, first-generation Intel-based Open Compute servers are used as the platform for testing as seen in Fig. 1. These particular servers have a highly efficient design in terms of their thermal performance, as well as in their mechanical design, power efficiency, and acquisition costs. Detailed description of the server design can be found in Refs. [9] and [10]. The key elements of the thermal design include

wider spacing of the primary heat generating components (CPUs), an air duct that directs flow over these key components, and a larger chassis height (1.5 rack units). The 1.5 U height allows for use of four 60 mm \times 25 mm fans, which are of much higher efficiency than the typical six 40 mm \times 25 mm fans found in commodity 1 U servers. Under normal operation, the four 60 mm fans adjust their speed based on a PWM signal that is dictated by CPU die temperatures. An FSC algorithm is set in the motherboard's BIOS settings to control the CPU die temperatures within a fixed temperature range.

A key feature of these particular servers that enabled simplification in the present analysis is the partitioning of the primary motherboard portion from the power supply (PSU) and hard drive section of the server as seen in Fig. 1. This partition allows for complete compartmentalization of airflow streams between the two sections. In this study, the airflow provided by the fan configurations are primarily to influence the temperature of CPUs. Additionally, the physical dimensions of the motherboard section allow for a square geometry to locate a fan array. This area is roughly 330 mm \times 330 mm if four servers are stacked on top of each other. When analyzing total rack fan power in the work, only the fans drawing air across the motherboard are considered. It is assumed that PSU airflow and power consumption will remain constant, independent from the proposed rack-level fan schemes.

Modified, Rack Level Fan Scenario. In order to evaluate the possible energy savings of implementing a larger, rack-level fan system, simple modifications to the baseline servers were made. A stack of four servers were used to simulate a "rack" scenario in the laboratory setting. A simple wooden frame and acrylic panel were constructed to hold the larger fan array at the back of the servers. A previous computational study by the authors indicated that a minimum distance of 20 mm was needed between the back of the servers and the fan array to ensure uniform distribution of airflow between all four servers [8]. Following this minimum requirement, the fan array for this study was placed 25.4 mm (1 in) from the rear of the servers.

The simulated rack of four servers includes a total of 16 60 mm fans in the baseline scenario. Each of these 16 fans is enclosed within the chassis of their respective server. In the modified case of 80 mm fans, nine total fans are located at the back of the simulated rack. In the modified case of 120 mm fans, four total fans are located at the rear of the simulated rack. Figure 2 depicts these three different scenarios with both a side view and rear view of the simulated rack with the PSU section of the server omitted for simplification.

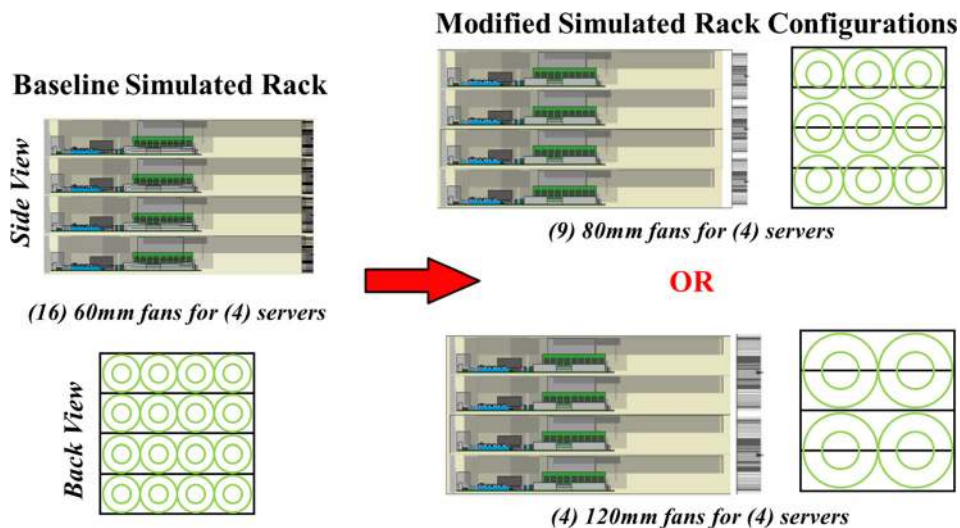


Fig. 2 Depiction of the consolidation of fans to a shared array at the back of the rack

Table 1 Specifications for the selected fans used in this study

Setup	Frame (mm)	Max. airflow (cfm)	Max. static pressure (in H ₂ O)	Rated speed (rpm)
60 mm	60 × 60 × 25	37.1	0.62	7600
80 mm	80 × 80 × 38	100.1	1.98	9500
120 mm	120 × 120 × 25	171.0	0.90	5100

Table 1 outlines the specifications for each fan size selected for this work. Selection of the larger fans (80 mm and 120 mm) was based on a requirement that they meet or exceed the airflow rate and static pressure delivered by the smaller 60 mm fans for a given rack. An in-depth discussion of the selection process is provided in Ref. [8]. Detailed characterization of the fans was performed on an Air Flow Bench test chamber to generate the fan curves and confirm the expected efficiency increase with larger fans. Details of this testing procedure can be found in Refs. [8] and [11]. Using the Air Flow Bench test chamber, at each operating point on the fan curve, the fan's total electrical power draw by a voltage direct current (VDC) power supply was recorded. Fan efficiency is given as the ratio of the output hydraulic power to electrical input power by below equation:

$$\eta = \frac{\dot{Q} * \Delta P}{V * I} \quad (1)$$

Where hydraulic power is the product of the volume flow rate and differential pressure and electrical power is the product of input voltage by current drawn. Figure 3 shows the efficiency curves from each of the three fan sizes when operated at maximum speed. The flow rate is given in terms of the flow rate for the simulated rack, which is a parallel configuration of 16 60 mm fans, nine 80 mm fans, or four 120 mm fans. It can be seen that the peak efficiency of the 60 mm, 80 mm, and 120 mm fans are 14.8%, 29.3%, and 29.3%, respectively. There is a possibility to improve the overall design of the 60 mm fan and thereby realize energy savings at the rack level. However, the roughly double efficiency of the larger fans should manifest in significant reduction in cooling power consumption for a given airflow volume.




Fig. 3 Peak total efficiencies for the 60 mm, 80 mm, and 120 mm at their maximum fan speeds in the rack configuration

Simulated Rack Setup. Figure 4 shows the laboratory test setup of the simulated rack of four servers in the baseline configuration with 16 60 mm internal server fans in place. The servers are labeled A through D with A being the bottom server and D on top. In order to accurately evaluate the cooling energy savings of the proposed fan modifications, the power delivered to the fans is decoupled from the rest of the server. This is done by powering the fans with an external 12VDC power supply source (Agilent E3633A). A control circuit breadboard with four pin headers for each fan mimics the connections fans would make to the motherboard. The four pins serve the functions of ground, power, tachometer sensing, and PWM control. The tachometer output signal is sent to an Agilent 34972A data acquisition unit to record fan speeds in rotations per minute. The control signal is directed by either an external function generator (Arduino microcontroller) which provides a fixed PWM signal or via the servers' internal fan speed control algorithm, which is relayed from the motherboard. It is critical that all these signals (power, sense, and control) maintain a common ground source for proper operation. To assess the overall efficiency of the rack, the remaining server power consumption (IT load) is recorded by a Yokogawa CW121 power




Fig. 4 Laboratory test setup of four servers stacked to represent a rack. Pertinent testing equipment is identified accordingly.

Table 2 Computational workload stressing conditions applied to the servers

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

meter, which measures the 277VAC and current delivered to the whole rack, as well as the power of two of individual servers B and D. The ambient temperature of the laboratory environment is monitored throughout testing with Omega OM-EL-USB-1-LCD temperature loggers, and a rack inlet temperature of $25^{\circ}\text{C} \pm 1.0^{\circ}\text{C}$ is observed during all tests. A desktop workstation is used to communicate to each of the four servers and execute a command script, as well as log all data with common timestamps for effective data reduction.

Procedures and Computational Loadings. For the initial testing, a synthetic computational load is applied to all four servers using a free software package, LOOKBUSY [12]. This program allows for stressing of individual computational subsystems: CPU, memory, and network input/output. The particular servers under test typically operate computationally intensive workloads and as such, focus is given to stressing the CPUs, with only moderate memory usage allocated. The native Linux commands *mpstat* and *free* are used to monitor CPU utilization and memory usage, respectively. Table 2 lists the various computational loadings that are ran on the servers. These loads represent the range of possible computational operating conditions a server may experience in actual service, from idle to medium to maximum utilizations. A complete thermal solution should be efficient across all these operating conditions. An inbuilt software program from the motherboard manufacturer reports the critical temperatures of the CPUs by accessing the on-die digital temperature sensors. For the tests conducted, each computational workload is applied for a 30-min period with 30 mins of idling between each load. Steady-state temperatures are achieved within 20 mins from the start of each applied load, and the last 10 mins are averaged for reporting here. Each cycle of computational loadings is applied two more times to ensure repeatability of the measurements.

The following key terms will be used through this work when referring to the manner in which fans are powered and controlled and are defined as follows:

Fan power:

- Internally powered—fans receive their 12VDC power signal directly from the motherboard fan header.
- Externally powered—fans receive their 12VDC power signal from a DC power supply external to the servers.

Fan control:

- Internally controlled—fans receive their speed control signal directly from the motherboard fan header with the server’s native fan speed control algorithm in effect.
- Externally controlled—Fan speeds are dictated from an outside source such as an Arduino microcontroller circuit. This may be either a fixed PWM signal or a more detailed control scheme driven by the internal CPU temperatures.

Results and Discussion

In order to assess the fan energy savings possible by consolidating the internal server fans to a rack-level array, the cooling performance of each configuration is established. This is done by

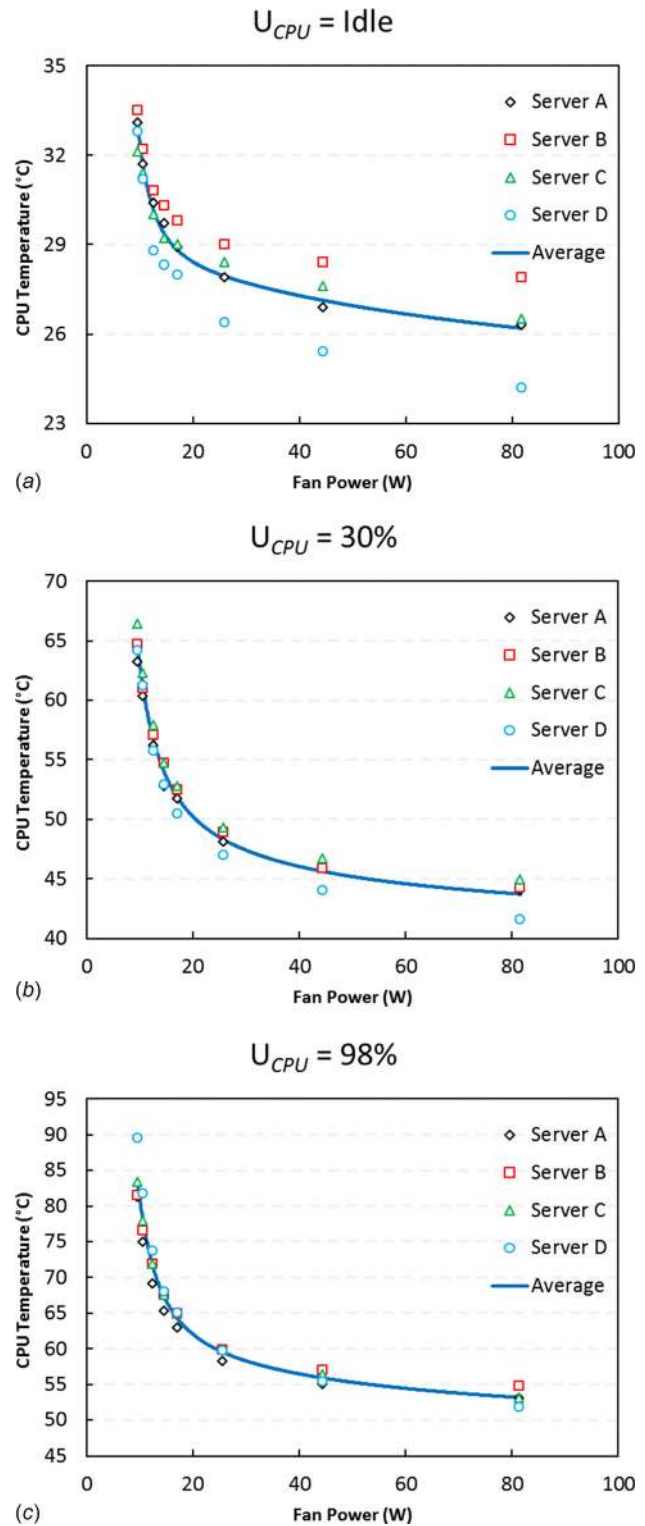


Fig. 5 Relation between total rack fan power and CPU die temperatures for each of the four servers when operated at uniform computational workloads of idle, 30%, and 98% CPU utilization in the baseline 60 mm fan configuration

recording the temperature of the servers’ critical components, the CPUs, as a function of fan power. Fan power is varied by operating the fans over their range of speeds by adjusting the fixed PWM signal delivered externally from the microcontroller board.

A steady-state temperature is established at fixed fan speeds over the range of operating speeds and computational loadings. In this externally powered, externally controlled test, precise

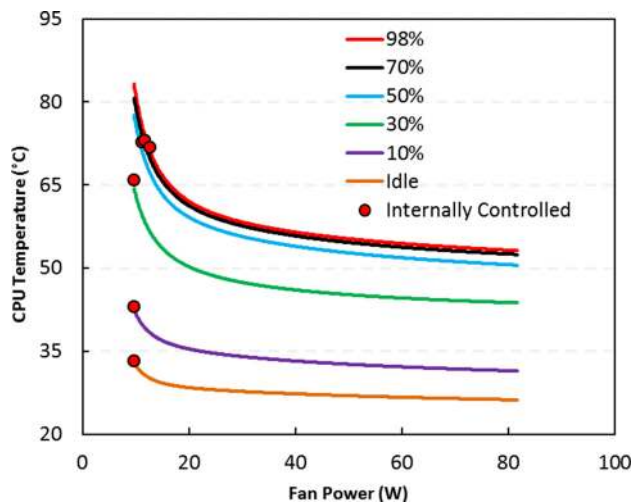


Fig. 6 Location of the discrete system operating temperatures when the fans are internally controlled lie directly on the curve fit lines of the externally controlled tests across all computational loadings

operation of fans at discrete power levels can be achieved. For these tests, all four servers were stressed with uniform computational loading across the rack. Figure 5 shows the critical relation between CPU die temperature and fan power when the 60 mm fans were employed at the idle, 30%, and 98% CPU utilization loadings. Fan power is given as the cumulative fan power of the rack (ignoring the PSU fan power as discussed previously) and an average CPU temperature of all four servers is shown by the curve fit line.

As should be expected, when the fan speeds are increased, there is a corresponding increase in fan power and subsequent reduction in CPU die temperature. As is typically seen in conventional chip packaging architectures, there is a sharp decrease in temperatures at the left leading portion of the curve due to significant reduction in the convection resistance between the heat sink and the air stream. The temperatures begin to level off toward the end of the curve as conduction resistance in the heat sink begins to dominate the heat flow path when convective resistance is minimized.

As a check for the performance that the servers may experience during normal operation, the server fans are operated in an externally powered, internally controlled configuration. In this configuration, the native fan speed control FSC signal from the servers is sent from the motherboard to the control circuit powering the fans. Each row of four fans receives the individual PWM control signal from their respective server's motherboard. When the native FSC is in effect, the fans fluctuate in speed to achieve a target CPU temperature which is within a deadband range. The average of this fan speed and power over the last 10 mins of each computational load is reported here. It can be seen from Fig. 6 that these discrete operating points lie directly on top of the previously generated curves at each computational loading conditions.

To assess the possible fan energy savings in the larger 80 mm and 120 mm fan configurations, similar externally powered, externally controlled tests were conducted when the larger fan arrays were fitted to the back of the rack. Again, the fans were operated over a range of fixed power levels by adjusting the PWM signal and hence speed. A comparison of the average CPU temperature as a function of fan power for the larger 80 mm and 120 mm fans to the baseline 60 mm fans at idle, 30%, and 98% CPU utilizations is shown in Fig. 7.

Across all computational loadings, to achieve a given CPU temperature requires significantly less rack fan power when larger 80 mm or 120 mm fans are employed. This indicates that more air-flow is delivered to the servers for a fixed rack fan power. There is no distinct superiority between the 80 mm and 120 mm fan

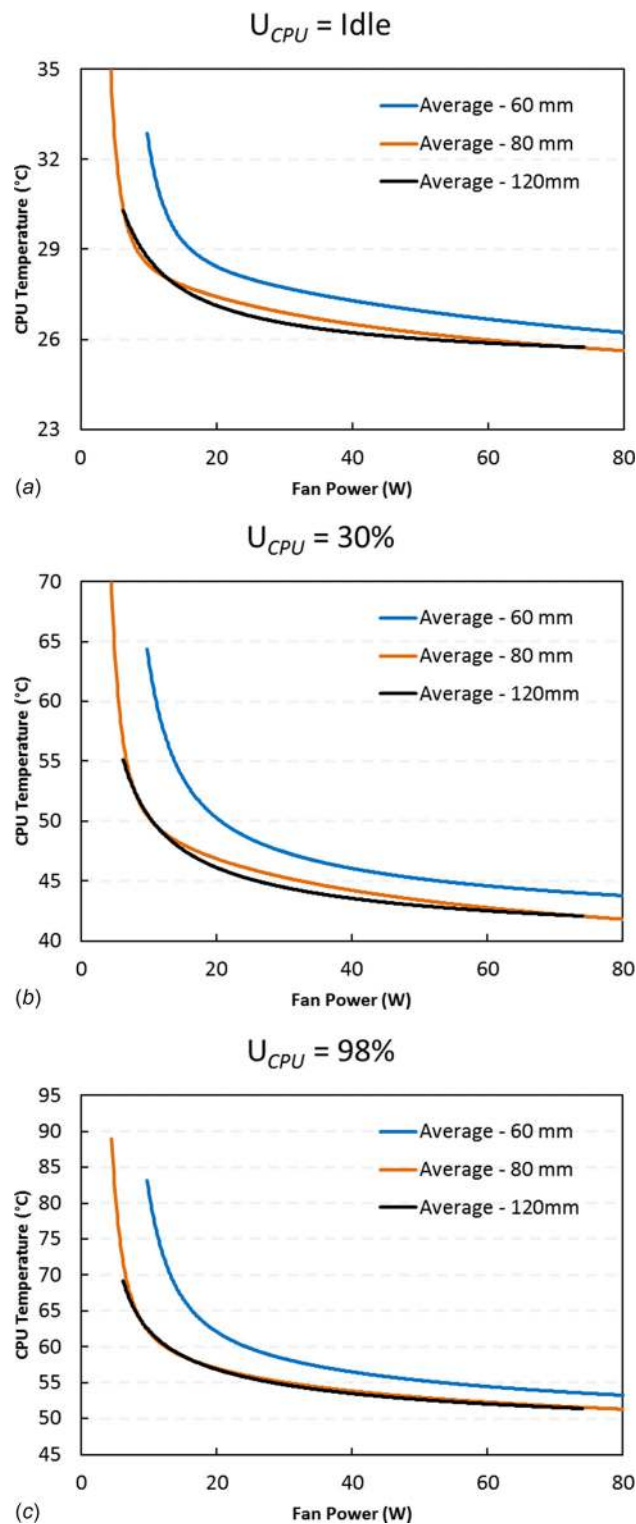


Fig. 7 Comparison between cooling performance of 60 mm, 80 mm, and 120 mm fans at idle, 30%, and 98% CPU utilizations

performance as both average curve fit lines intersect and cross paths depending on computational loading and fan power. To obtain a more precise understanding of the fan energy savings possible with the rack-level fan approach, comparison at specific operating temperatures can be made by looking at the fan power required to reach the CPU temperatures achieved during the 60 mm externally powered, internally controlled test. Using the curve fit models for the 80 mm fans, fan power at the discrete

Table 3 Comparison between required rack fan power for 60 mm and 80 mm fans to achieve target operating temperatures obtained during the externally powered, internally controlled 60 mm fan case

U_{CPU} (%)	CPU DT (°C)	60 mm (W)	80 mm (W)	% Savings
IDLE	33.2	9.59	4.78	50.2%
10	43.1	9.58	4.78	50.1%
30	66.0	9.56	4.76	50.2%
50	72.7	11.19	5.31	52.6%
70	73.2	11.67	5.70	51.2%
98	71.8	12.66	6.08	52.0%

Table 4 Comparison between required rack fan power for 60 mm and 120 mm fans to achieve target operating temperatures obtained when the 120 mm fans are operated at idle speeds

U_{CPU} (%)	CPU DT (°C)	60 mm (W)	120 mm (W)	% Savings
IDLE	30.3	12.42	6.51	47.6%
10	38.1	13.08	6.49	50.4%
30	54.9	13.86	6.40	53.8%
50	65.4	13.62	6.36	53.3%
70	67.6	13.78	6.34	54.0%
98	69.0	13.66	6.33	53.7%

operating temperatures was extracted and the savings shown in Table 3 for all computational loadings were tested. Savings from 50.1% to 52.6% were achieved with the array of nine 80 mm fans when compared to the baseline 60 mm fan case.

The same comparison between the 120 mm fan array and 60 mm fans could not be made because, even at idle speeds, the 120 mm fans cooled the CPUs to below the discrete operating temperatures of the 60 mm externally powered, internally controlled test. To accommodate this, a comparison is made by extracting the fan power from the 60 mm fan curve fit models at the highest CPU temperatures achieved with the 120 mm fans. Table 4 shows this comparison across all computational loadings tested. A wider range of savings is seen with the 120 mm fan array, spanning 47.6% up to 54.0%. In either case of the 80 mm or 120 mm fan arrays, savings are in line with the improved peak total efficiencies and predicted savings of Fig. 3.

Comparison Under Nonuniform Load. Actual data centers rarely, if ever, operate with uniform computational loading across all servers in a rack. Results from the Results and Discussion section represent an idealized condition and best case scenario for possible savings. In order to test the limits of the validity of a rack-level fan configuration, a simple test was conducted to evaluate the savings achievable in a worst possible condition. In this situation, the worst possible situation for the larger, rack-level fan array would be if all fans in the system are controlled by a single

Table 5 Comparison of rack fan power between the “worst case scenario” of the rack-level fans that the lowest energy state possible for the baseline 60 mm fans

Test case	Total rack fan power (W)
60 mm “best case scenario”	9.59
80 mm “worst case scenario”	6.20
120 mm “worst case scenario”	6.35

server operating at maximum computational load (and hence temperature) while all other servers remain at idle loading, as seen in Fig. 8(b). This would constitute the most inefficient use of fan energy possible because all fans in the array would operate at higher speeds. The shared airflow of the servers would lead to overcooling (and unnecessarily increased rack fan power) of the three idle servers. This can also be thought of as the simplest possible control scheme that may be developed for a rack-level fan configuration in which the maximum temperature of the rack dictates the speed of all the fans in the array. A comparison can be made to the best possible case for the 60 mm fans, in which all servers in the rack are operating at idle computational loadings. This represents the lowest possible energy state for the 60 mm fans because, as seen in the externally powered, internally controlled baseline test, the fans will operate at their idle speeds. This configuration is shown in Fig. 8(a).

To test this notion, server D, which showed consistently higher temperatures (as seen in Fig. 5(c)), was operated at 98% CPU utilization. The fan control signal from the motherboard in server D was sent to all nine and four fans in the 80 mm and 120 mm arrays, respectively. A comparison between rack fan powers for these three cases is shown in Table 5. Even in the worst possible situation in which three servers are overcooled, the 80 mm fans still provide a 35.3% savings in rack fan power compared to the lowest possible energy state of the 60 mm fans. Even at idle speeds, the 120 mm fans overcool to the point that the native fan speed control algorithm does not engage and rack fan power savings of 33.8% are achieved. These extreme situations indicate that a rack-level fan configuration provides an opportunity for significant fan energy savings when compared to smaller, chassis fans. The savings presented here also do not take into account any savings in IT energy that may be realized by the overcooling of components and hence, reduced leakage power in the silicon devices. More detailed description of these impacts are discussed in a follow-on paper to this work.

Further Discussion. Along with reduction in operating rack fan power, consolidation of smaller internal server fans to a rack-level array may provide additional benefits to data center owners and operators as discussed here.

Clearly, the number of fans in the system may be substantially reduced, from either 16 to nine or four in the particular hardware used for this study. Although larger fans may have a slightly higher first cost per unit, the overall system cost may be reduced.




Fig. 8 Depiction of (a) the “best case scenario” for the 60 mm fans in which all servers are operating at idle computational load and (b) the “worst case scenario” for the rack-level fans in which a single, high computationally loaded server dictates the speed of all the fans in the array

Table 6 Comparison of relative costs associated with the fans used in this study

Fan size	Relative cost per fan unit	Quantity per “rack”	Relative cost per “rack”
60 mm	0.37	16	1.00
80 mm	0.49	9	0.76
120 mm	1.00	4	0.68

Table 6 compares the relative individual unit costs and relative cost per rack for 60 mm, 80 mm, and 120 mm fans used in this study. Prices were taken from online electronics component distributor. The results here are not exhaustive of all possible fans that may be selected for this application and do not include consideration for price breaks for larger quantity orders. However, as a first estimate, this may provide an indication of the capital savings that may be achieved. Additional physical infrastructure may need to be put in place at the rack level to house an array of fans, as well as any control circuitry, which may have additional associated costs. Part II of this work will explore how the control scheme of rack-level fans may be further designed and optimized to realize maximum fan power savings from such a configuration.

An additional benefit of the reduced number of fans in the system is fewer points of failure. Part II of this work will investigate the impact of failure and redundancy in a rack-level fan configuration.

An emerging paradigm in the data center industry is “rack disaggregation.” The concept involves separating a traditional server’s compute, memory, storage, and input/output into discrete modules. The benefit being that the individual subsystems of a server may be changed and upgraded at different frequencies without having to replace all parts of the system, resulting in less material waste [13]. In this manner, a rack then becomes the fundamental building block of a data center. The concept of rack-level fans aligns with this idea well. As could be seen in the case of the particular 80 mm and 120 mm fans studied here, there is a significant amount of additional cooling capacity available beyond what is required to maintain safe operating temperatures for the servers. This additional capacity may last for multiple refresh cycles of IT equipment as component heat loads and power densities increase. Prudent selection of fans during initial design stages and realistic projections of future server hardware requirements will ensure rack-level fans have their greatest impact.

Conclusions

Consolidation of smaller fans from within servers to a rack-level array of larger fans offers an opportunity to significantly reduce total rack fan power. Initial fan performance characterization showed that the 80 mm and 120 mm fans chosen for this study have peak total efficiencies two times greater than that of the baseline 60 mm fans. When configured into an array at the rear of a rack and at uniform computational loading across four servers, the 80 mm fans offer 50.1–52.3% reduction in total rack fan power at a fixed target temperature. Similar savings of 47.6–54.0% in rack fan power were observed when the 120 mm fans were employed. When a highly nonuniform computational workload was applied across all four servers, savings in rack fan power were found to be 35.3% and 33.8% for the 80 mm and 120 mm fans, respectively, when compared to the lowest possible fan operating energy of the 60 mm fans. Part II of this work will study additional operational issues for a rack-level fan array such as developing a control scheme, as well as redundancy in failure situations.

Acknowledgment

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

- The National Science Foundation Center for Energy Smart Electronic Systems (ES2) (Grant No. IIP-1134867).
- Directorate for Engineering (Grant No. 1134821).

Nomenclature

- I = current (A)
 \dot{Q} = volume flow rate (m^3/s)
 V = voltage (V)
 ΔP = pressure (Pa)
 η = fan efficiency (%)

References

- [1] ASHRAE, 2012, “IT Equipment Thermal Management and Controls,” American Society of Heating, Refrigerating and Air-Conditioning Engineers, Atlanta, GA, Document No. **ASHRAE TC 9.9**.
- [2] Alkharabshah, 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.
- [3] Strike, N. D., 2015, “Fan Efficiency, an Increasingly Important Selection Criteria,” NMB Technologies Corporation, Tempe, AZ, accessed Jan. 21, 2015, http://www.nmbtc.com/fans/white-papers/fan_efficiency_important_selection_criteria/
- [4] Greenheck, 2013, “Understanding Fan Efficiency Grades (FEG),” Greenheck Fan Corp, Schofield, WI, Report No. **CS/104-13**.
- [5] AMCA, 2010, “Energy Efficiency Classification for Fans,” American Movement and Control Association, Arlington Heights, IL, Standard No. **205**.
- [6] Holohan, M. F., and Elison, B. P., 2007, “Fan Laws for Rack Systems,” *ASME Paper No. IPACK2007-33892*.
- [7] Kodama, Y., Itoh, S., Shimizu, T., Sekiguchi, S., Nakamura, H., and Mori, N., 2013, “Imbalance of CPU Temperatures in a Blade System and Its Impact for Power Consumption of Fans,” *Cluster Comput.*, **16**(1), pp. 27–37.
- [8] 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.
- [9] 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.
- [10] Li, H., and Michael, A., 2011, “Intel Motherboard Hardware v1.0,” Open Compute Project, accessed Nov. 22, 2015, <http://opencompute.org/wp/wp-content/uploads/2011/07/Server-Intel-Motherboard-Specifications-v1.pdf>
- [11] AMCA/ASHRAE, 2007, “Laboratory Methods of Testing Fans for Certified Aerodynamic Performance Rating,” Air Movement and Control Association International, Inc./American Society of Heating, Refrigerating and Air Conditioning Engineers, Arlington Heights, IL/Atlanta, GA, ANSI/AMCA Standard No. 210-07, ANSI/ASHRAE Standard No. **51-2007**.
- [12] Carraway, D., 2013, “Lookbusy—A Synthetic Load Generator,” Version 1.4, accessed Oct. 23, 2017, <http://www.devin.com/lookbusy>
- [13] Skillern, R., 2015, “Meet the Future of Data Center Rack Technologies,” Data Center Knowledge, San Francisco, CA, accessed Jan. 21, 2015, <http://www.datacenterknowledge.com/archives/2013/02/20/meet-the-future-of-data-center-rack-technologies/>