Is Hot IT a False Economy? An Analysis of Server and Data Center Energy Efficiency as Temperatures Rise

Stephen Clement*, Kat Burdett*, Nour Rteil*, Astrid Wynne, and Rich Kenny

Abstract—As demand for digital services grows, there is need to improve efficiency and reduce the environmental impact of data centers. The largest energy consumer in any data center is the IT, followed by the systems dedicated to cooling. Aiming to improve efficiency, and driven by metrics like PUE, there is a trend towards running data centers hotter to reduce the cooling energy. There is little research investigating the effect this will have on the IT beyond failure rates. To ensure overall efficiency is improving, we must view the data center as a system of systems, taking a holistic view rather than focusing on individual sub-systems. In this paper we use industry standard benchmarks and a wind-tunnel to profile typical enterprise IT. We analyze the effect of environmental conditions on IT efficiency, showing minor increases in temperature or pressure impact the efficiency of servers. Using an idealized, simulated data center case study we show that the interaction between cooling systems, server behavior and local climate are non-trivial and increasing temperatures has potential to worsen efficiency.

Index Terms—data centers, energy efficiency, sustainable computing, server benchmarking, power consumption, inlet temperature, ASHRAE climate zones, data center cooling, environmental impact

1 Introduction

Data centers are dedicated facilities for housing information communication technology (ICT) hardware. They provide the fundamental infrastructure that underlies our increasingly digital society. In 2020, data centers accounted for roughly 1% of global energy demand with a further 1.1-1.4% consumed by the data transmission networks associated with them [1]. While the growth in energy demand has slowed considerably since the 90% growth boom years of the early 2000s [2], thanks to improvements in energy efficiency, there is still an increasing demand for data center capacity tied to the massive growth in digital utilization. 51% of the world's population had access to the world wide web in 2018, and that number is predicted to grow to 66% by 2023 [3]. The number of connected devices is also expected to rise to an estimated 29.3 billion [3] with the growth of the Internet of Things (IoT) and smart cities [4]. More people and devices online will mean continual growth in the infrastructure that supports this digital environment.

Data centers have a large environmental impact in terms of resources consumed (energy, water, critical raw materials), greenhouse gasses emitted [5], and e-waste generated. One industry accepted practice for reducing energy consumption is widening temperature and humidity ranges, particularly in conjunction with economizers. It is often perceived that operating at higher temperature set-points reduces cooling energy and consequently overall energy consumption. But the effect of this on the IT equipment, specifically servers, has not been well examined and any change in IT load impacts the amount of cooling required.

This paper studies the impact temperature and pressure have on servers and DC cooling energy and the interconnection between the two. We aim to determine whether there is always a net benefit to energy savings in raising inlet temperature across all climate zones. Our primary contributions are novel environmental testing of servers using a wind tunnel, a multi-generational analysis of server behavior under different environments and applying this profiling data to a global case study.

We begin by analyzing the power and performance of 1U enterprise rack servers with five different CPU generations at varying temperatures and pressures to determine the impact environmental conditions can have on IT. We then extrapolate these findings to a simulated, idealized data center based in London to evaluate the overall effect on a DC and compare different cooling systems. Finally, we extend this simulation to identical data centers in other geographical regions to assess the relationship between the DC and its local climate.

The rest of this paper is arranged as follows. Section 2 describes the rationale behind this work, emerging trends in data centers and the industry's reluctance to extend operating temperatures. Section 3 provides an overview of the methodology used to test the servers. Section 4 presents the findings of our experiments, followed by a case study presented in section 5 examining the interplay of IT and cooling energy in a DC. This is expanded on in section 6 to include the climate in 23 other locations worldwide. Finally, section 7 discusses conclusions and future work.

- *Stephen Clement and Kat Burdett are with Techbuyer, Harrogate, U.K Email:{s.clement,k.burdett}@techbuyer.com*
- *Nour Rteil is with Interact, Harrogate, U.K Email:nour@interactdc.com*
- *Astrid Wynne and Rich Kenny are with Techbuyer, Harrogate, U.K Email:{a.wynne,r.kenny}@techbuyer.com*

Manuscript received 29 June 2022; revised 25 November 2022; revised 21 April 2023.
** Authors contributed equally to this work*
Corresponding author: Stephen Clement

2 Background

Since its inception in 2006 and promotion by the Green Grid in 2007, PUE has been the primary metric used by data centers to continuously measure and improve their site's infrastructural energy efficiency [6]. PUE is a measure of the ratio of energy consumed by IT hardware against energy consumed by the entire data center and officially become part of an ISO standard in 2016 [7]. Today, some hyper-scale data centers are able to achieve yearly PUEs as low as 1.1 across all their sites [8] [9], while the average PUE globally is down from 2.5 in 2007 to 1.59 in 2020 [10].

Some of these efficiency gains are attained by the implementation of best practices in cooling and airflow management such as those specified in The Green Grid Data Center Maturity Model [11] and the EU Code of Conduct on Data center Efficiency [12]. Traditionally, data centers have used air-based cooling systems based on Direct Expansion (DX) or chiller systems. A DX system is comprised of a Computer Room Air Conditioning (CRAC) and air-cooled condenser as a means of removing heat. A chiller cooling solution utilizes a chilled water-glycol mix comprising of a cooling tower or dry cooler for heat rejection which feeds chilled liquid to a Computer Room Air Handling (CRAH). Chiller systems typically incur greater capital expenditures than DX systems but benefit more from economies of scale.

Measures to improve efficiency for air-cooled systems include the use of blanking panels, proper management of raised floor plenums or their alternatives, and implementing aisle containment to prevent the hot and cold air from mixing by creating a physical separation between the upstream (cold) and downstream (hot) portions of a server. This can be achieved with room-based CRAC/CRAH units, although in-row units are becoming more common in aisle-contained DCs also [13].

As well as minimizing the increase in air temperature between the outlet of the CRAC or CRAH unit and the inlet of the servers, the separation of the aisles allows for a pressure gradient to be formed across the servers, potentially aiding in the flow of cool air through the servers [14]. This pressure differential can be minimal, requiring the server fans themselves to drive flow, or it can be raised to differentials of 15-25Pa – putting a greater load on the CRAC but taking the load off the server fans to drive the flow [15].

Today, some data centers also utilize water/liquid-based cooling systems within the room and close to the IT, such as direct-to-chip cooling, rear-door heat exchangers and immersion cooling which can be more energy efficient, especially for greater power density – although those are still not widely adopted [16].

Where ambient conditions allow, hyperscalers like Microsoft, Google and Meta are adopting free and evaporative cooling to reduce energy consumption [17] as well as custom IT [18] designed for greater efficiency.

Traditional enterprise data centers lag behind: the industry's average PUE is stagnant around 1.58 [19]. Data center environmental bodies have identified over-cooling to be one of the key reasons why this is still the case, and in recent years, these bodies have worked collectively to increase the allowable operating temperature and humidity ranges in data centers. In 2011, ASHRAE updated its guidelines to define two additional classes of operation, A3 and A4, providing higher allowable temperature boundaries for up to 40°C and 45°C, respectively, as shown in Table. 1 [20].

TABLE 1
ASHRAE Datacenter Equipment Environmental Specifications

	Classes	Dry-Bulb temperature	Humidity range, non-Condensing	Max. Dew Point
Recommended	A1-A4	18-27°C	5.5°C DP to 60% RH and 15°C DP	N/A
Allowable	A1	15-32°C	20% to 80% RH	17°C
	A2	10-35°C	20% to 80% RH	21°C
	A3	5-40°C	-12°C DP & 8% RH to 85% RH	24°C
	A4	5-45°C	-12°C DP & 8% RH to 90% RH	24°C

2.1 Reasons to Keep a Cold DC

Despite ASHRAE's new device classes, the average operating temperature of data centers has hardly changed [21] from 20-24°C. Even when newer equipment that can support wider ranges is installed, ambient temperatures are often not modified to align with them. This reluctance to extend operating ranges is mainly attributed to concerns over IT hardware reliability, reduced response times to manage cooling failures, warranty issues, and potential decreased server efficiency. A lack of research, industry consensus, and historical data compound these concerns.

2.1.1 *IT reliability and failure rates*

In 2011, ASHRAE investigated relative server failure rates when raising server temperature above 20°C based on reliability data from multiple hardware vendors [20]. Continuous operation at 27°C compared to 20°C increased server failure rate by 1.2 times, growing to 1.6 times at 35°C.

In contrast, a 2008 study by Intel [22] using 900 blade servers over a 10-month period found server reliability was not significantly affected by temperature or humidity, even with relatively poor air quality. In a more recent white paper [23], after running Intel data centers at higher temperatures (32.7°C) for eight years (2013-2021), they found no increase in component failure.

Pinheiro et al. [24] studied the effect of temperature on hard disk failures in Google data centers and found a drop in disk failure rates with increased temperatures, except for at very high temperatures (45°C). In 2012, El-Sayed et al. [25] analyzed more data on hard disk replacements at different Google data centers and observed lower disk failure rates than previous models predicted below 50°C. They have also noted no correlation with higher temperatures for Dynamic Random Access Memory (DRAM) failures and node outages.

2.1.2 *Concerns towards design set points, safety margins and hot spots*

With increased inlet temperature, the hot aisle temperature needs to be continuously monitored to ensure that the data center's temperature does not exceed the design parameters for nearby equipment, cables, and hardware. Some rack-based equipment such as power distribution units (PDUs) and network switches have distinct operating envelopes with an upper limit of 45°C [21] which needs to be considered; though these are typically higher than the same operating ranges for servers.

In addition to hot spots, another common concern among data centers is reduced safety margins. Most servers have a critical temperature threshold and will shut down when that threshold is reached to avoid serious damage. Similarly, CRAC units have temperature thresholds and built-in failure modes that force the unit to shut down when that threshold is crossed. As the ambient temperature in a data center increases, equipment will be operating closer to the maximum temperature, reducing the time for graceful shutdown or taking protective measures in the case of critical failure event [26].

Nonetheless, hotspots and safety margin concerns can be controlled if data centers follow the best practice in cooling and airflow management [12]. One of the most effective methods is hot and cold aisles containment. It prevents hot spots and hot air recycling by physically separating the IT exhaust air from the supply air in the DC which can effectively reduce the mixing of hot and cold air and allow for a safe temperature increase.

2.1.3 Decreased Efficiencies

Server power tends to increase with temperature, particularly in older hardware, caused by increased fan activity and an increase in silicon electrical leakage current [27]–[29]. El-Sayed et al. [25] observed a significant increase in power when ambient temperature increased, attributing this mostly to fan power, in an 11th generation Dell PowerEdge R710 2U rack server (released in 2009). In 2018, Wang et al. [30] observed a correlation between CPU temperature and server power by intermittently cooling a CPU in isolation. Later, they profiled a single server's power consumption under load and up to 45°C to generate a temperature-aware model [31]. Both tests showed increased power consumption under hotter inlet temperatures or reduced cooling.

ASHRAE studied the power consumption against temperature for multiple vendors for class A2 and A3 devices [32], showing that server fan power consumption rises notably past the recommended range (25°C). Operating at the top of the allowable range could cause A2 class servers to use up to 20% more power than operating at 20°C. But newer A3 class servers are more efficient and designed for higher inlet temperatures. ASHRAE's data indicate that as temperature rises, the increase in power consumption of a class A3 device is generally only 50% of that of a class A2 device.

The ASHRAE analysis highlights the variation of power consumption among models, vendors and form factors. For instance, 1U rack servers are at a disadvantage in that they tend to be less efficient than larger servers at higher inlet temperatures because of the smaller size and higher rotational speed of their fans [21]. Therefore, these servers are more likely to exhibit higher increases in energy consumption when the inlet temperature rises. Blade servers, on the other hand, typically use larger fans at lower rotational speeds to achieve the same volume of airflow.

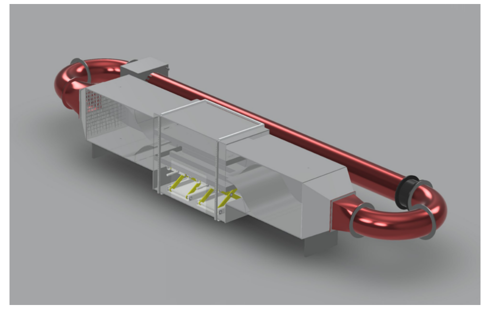

Fig. 1. Low-speed server wind-tunnel render, showing cut-away test chamber and flow conditioning geometry

3 Methodology

We will benchmark servers under varying environmental conditions to determine the impact of pressure and temperature on server efficiency. Using an industry-standard benchmarking tool (SERT) we can monitor power and performance during the benchmark and generate an aggregate efficiency score. To control and vary the environmental conditions we build an experimental low-speed wind-tunnel to emulate the range of temperatures a server might experience in a data center.

3.1 Wind-Tunnel

The low-speed, closed-loop wind-tunnel was designed to simulate the various thermal conditions seen within a data center and push that envelope beyond ASHRAE-recommended standards. The tunnel, shown in Fig. 1, is based upon the design proposed in [33] and is comprised of the test chamber, a heating and cooling system, flow conditioning geometry, and a fan to modify the pressure drop across the server under test (SUT).

Servers are placed in the center of the test chamber such that the only passage air will take is through the server. The test chamber geometry is shaped so that the air flow through the chamber and into the front of the server is laminar. The depth and width of the chamber are variable and can accommodate servers and blade chassis from 1U to 12U. A flow contraction, shown in Fig. 2, based on the design by Bell and Mehta [34] was utilized upstream of the server, with accommodations made to ensure the laminar airflow is maintained for any size of server the tunnel can accommodate.

Maintaining a 5$^{\text{th}}$-order contraction shape was computationally and experimentally determined by Bell and Mehta to be the best in the range of configurations and geometries tested to minimize boundary layer separation, pressure drop and ensure a laminar flow entered the test chamber. It does this by maintaining the 5$^{\text{th}}$-order polynomial:

$$f(x) = H_i - (H_i - H_e)(6x^5 - 15x^4 + 10x^3) \qquad (1)$$

where H_e is exit height, or the height of the server being tested, and H_i is inlet height, the full height of the test chamber.

The contraction is made from a fixed length of aluminium sheeting. Given the variable nature of the exit height H_e, the position of H_i must move laterally along the

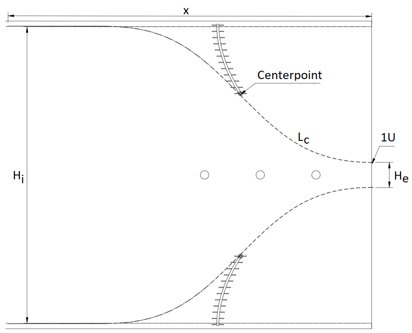

Fig. 2. Variable depth contraction cross section showing the relationship between inlet height (H_i), outlet height (H_e), contraction length (L_c). H_e varies in steps of 1U to accommodate servers of different dimensions.

length of the test chamber. To maintain the 5th-order polynomial in equation 1, the center position of the aluminium sheet sits on a curved slider.

Bell and Mehta determined the minimal acceptable ratio between contraction length and inlet height, a, to be 0.89.

$$a = L_c/H_i$$

To determine the length of the aluminium contraction sheet and the position of both the end point at H_i and the center-point, we treat the contraction sheet as a Bezier curve numerically approximated with Gauss quadrature for a 5th-order polynomial:

$$\int_{-1}^{1} \sqrt{\frac{dx}{dt}^2 + \frac{dy}{dt}^2} dt \equiv \int_{-1}^{1} f(t)dt \equiv \sum_{i=1}^{n} c_i f(t_i)$$

This is transformed from a range of -1 to 1 to a range of 0 to aH_i and combined with the Gauss quadrature with equation 1 leading to:

$$\int_{0}^{ah} \sqrt{\frac{25(-12+\mathrm{u})^2\mathrm{x}^4(-\mathrm{aH}+\mathrm{x})^4))}{16\mathrm{a}^{10}\mathrm{H}^8} + 1} dx$$

Where H_i is 12U or 534mm based on the height of the tunnel. This was in turn determined by a combination of the standardized ducting size, the expansion ratio of the WAD, and the depth of 12U of servers, which lead to an L_c of 553mm providing center-point and H_i locations along the length test chamber.

Under ideal circumstances, deviations in the cross-section of the tunnel should be kept between 10-15°to minimize any pressure drop or boundary wall separation [35], [36] – practically, this would result in an inordinately long test chamber, so a Wide Angle Diffuser (WAD) was installed upstream of the contraction and server.

Typically, a WAD would use honeycomb screens to enforce laminarization on the flow, but we opted for a series of overlapping rows and columns of rods as designed by Barratt and Kim [37], as this allowed for a greater rate of expansion while also minimizing pressure drop and turbulence.

The effectiveness of the WAD is determined by the porosity ϵ of the honeycomb, ideally, to minimize turbulence $\epsilon \geq 0.99$.

$$\epsilon = 1 - 0.907 \left(\frac{d}{S}\right)^2$$

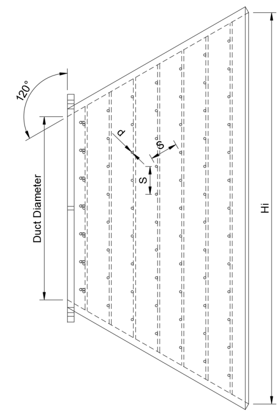

Fig. 3. Banked Wide Angled Diffuser (WAD) cross-section demonstrating expansion ratio and column/row spacing used to achieve laminar flow going into the test section.

where d is the diameter in millimeters of rod and S is spacing between rows or columns. Given the physical limitations of the very small diameter columns and rows, a 3mm steel wire and 38mm spacing was chosen, resulting in $\epsilon = 0.994$. The wire was placed 32mm apart along the length of the WAD, exactly out of phase to maintain the 38mm spacing between rods and columns.

Inside the test chamber, the temperature is measured at each end (upstream and downstream of the server) and a pressure differential is measured across the server. The temperature in the tunnel is maintained by a liquid-air heat exchanger attached to an air handling unit. A constant pressure difference across the server analogous to data center aisle containment is maintained by a variable speed fan, placed upstream of the server on the return ducting.

3.2 SERT Benchmark

SERT is a benchmarking tool created by the Standard Performance Evaluation Corporation (SPEC) [38]. It is a benchmark divided into three workloads targeting the primary components of a server: CPU, memory and storage. An efficiency score is calculated for each benchmark as a normalized performance score per Watt of power used. These are amalgamated into a SERT metric using a 65:30:5 weighing ratio of CPU to memory and storage workloads that gives a coarse but directly comparable measurement for the efficiency of a server.

3.3 Test details

Our previous research highlighted the substantial impact CPU generation has on the server's overall energy efficiency [39]. For this test, we have chosen a selection of typical enterprise servers of each generation spanning a wider range (2012-2017) to determine what effect ambient conditions have, if any, on newer hardware generations. The configurations tested are shown in Table. 2. We chose to test multiple generations of servers to determine if there were any trends within the wider industry reflected at the server level. In order to compare between generations, the specifications must be similar; however, architectural differences between generations of CPU mean that the technical

TABLE 2
Server specifications to be tested

Chassis (Released)	CPU (Released)	Memory		Disk	Power Supply
		DIMMs	Total Capacity		
PowerEdge R620 (2012)	Intel Xeon E5-2690 (Q1'12)	8×DDR3 1600MHz	64GB	500 GB SAS	1+1 Redundant 750W PSU
	Intel Xeon E5-2690 v2 (Q3'13)				
PowerEdge R630 (2014)	Intel Xeon E5-2690 v3 (Q3'14)	4×DDR4 2133Mhz			
	Intel Xeon E5-2699 v4 (Q1'16)				
PowerEdge R640 (2017)	Intel Xeon Silver 4116 (Q3'17)				
	Intel Xeon Gold 6148 (Q3'17)				

specifications like clock speed are not directly comparable in a way that indicates performance. To counteract this, we chose CPUs that occupy similar positions in their respective product ranges for their generation. We have both low and high-specification versions of the most recent servers, the R640 with Xeon scalable CPUs. For each of the servers tested, the total memory capacity, disk and power supply configuration remain the same, but memory generation and speeds change between those servers supporting DDR3 and DDR4.

Temperature variations for the servers are kept within the recommended environmental envelope defined by the American Society of Heating, Refrigerating and Air-Conditioning Engineers (ASHRAE) [20]. The PowerEdge R600 servers typically fall into the ASHRAE A2 classification per the Table. 1. It is possible for them to operate as A3 devices given lower specification levels and certain restrictions. In our case only the PowerEdge R640 with a Intel Xeon Silver 4116 could fall into this category.

To account for any fluctuations in SERT test results, benchmark tests are repeated a minimum of three times per scenario in order to avoid any spurious results. Using the wind-tunnel, each server was benchmarked at three different temperature ranges (20-23.3°C, 23-26.6°C and 26.6-30°C) and two differential pressures (<5Pa and >15Pa) across the server. The temperature regions were chosen to span SERT's valid ambient temperature range rather than the higher ASHRAE A2 to avoid invalidating the tests and to determine whether the relatively minor changes to the temperature that a DC might be able to achieve without retrofitting specific hardware would have any significant impact. The pressure values were chosen to be at the extremes that could be experienced by the server in an aisle contained environment. Either hot and cold aisles are kept around the same pressure with just enough pressurization in the cold aisle to avoid any flow from the hot side, or the cold aisle can be pressurized, potentially taking some load off the server fans. At low pressure differentials in the wind-tunnel, the server fans play a greater role in driving airflow, relegating the main fan to compensate for the back-pressure of the tunnel geometry.

4 Testing Results

Fig. 4 shows the SERT efficiency score of each server configuration for the environmental conditions tested. It is evident that server efficiency has improved significantly over the years, with SERT scores increasing with each generation. Each PowerEdge server chassis (R620, R630, R640) is represented twice in Fig. 4. All hardware in the chassis is identical except for the CPU. This demonstrates both how component choice can impact efficiency and the significant impact the CPU has on the server. As inlet temperatures increase, server efficiency decreases for every server tested; 1-2% from 20°C to 30°C (p-value < 0.01). In contrast, changes in pressure differential across the server had no significant effect on the efficiency score.

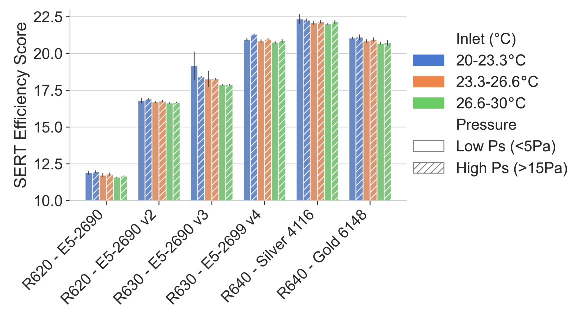

Fig. 4. SERT Server efficiency scores for each server at inlet temperatures of 20-23.3°C, 23.3-26.6°C, and 26.6-30°C.

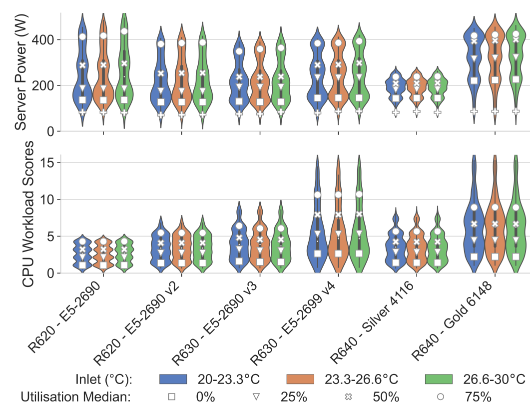

Fig. 5. The distributions of server power and normalized CPU workload scores for each server tested within the temperature ranges and the median power and score for each utilization.

4.1 Efficiency Detail

The two influences on efficiency scores are performance and power consumption. Fig. 5 shows the power consumption

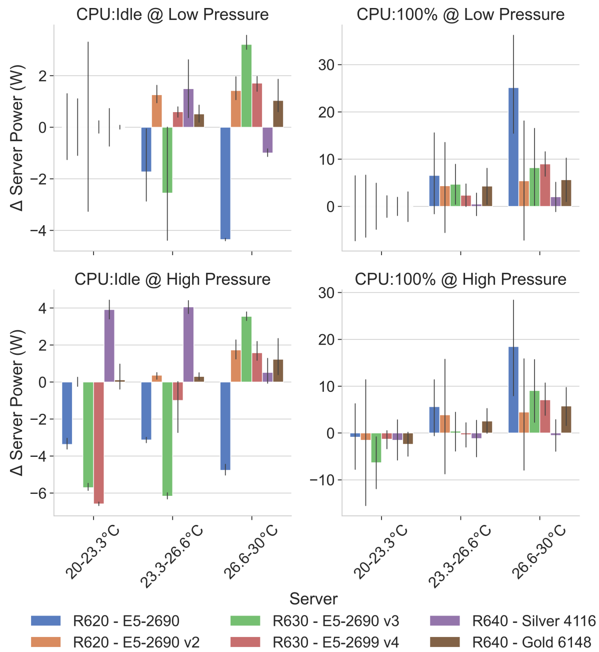

Fig. 6. Mean and standard deviation for the change in power due to the environment for each server at extreme utilizations (idle and 100%). The reference point for each server is the mean power during low pressure tests at 20-23.3°C. Separate reference points are used for each utilization.

and CPU workload performance for each server and the temperature range limited to the CPU workload. The frequency of each measurement is displayed as the width of the plot, giving an indication of the underlying distribution. The figure shows that the distributions of both power and score are multi-modal due to the server power and performance varying with target utilization in the benchmark, the median power, and score for each target utilization are also shown in the figure.

Fig. 5 shows the decline in efficiency as temperature increases is driven by increased power consumption, not reduced performance. The performance of each server increases proportionally to utilization, but power does not. In the older servers, the rate of power increase is greater at higher utilization, whereas the newer R640 server had larger power increases from 0% to 25%. Server power for earlier generations is convex with respect to utilization and concave in the later generations. Considering the relative change in performance, this indicates that aside from environmental considerations, workload allocation strategies between generations of servers offers an opportunity for optimization.

The effects of temperature and pressure are most exaggerated at the extremes of utilization: when the server is idling or running at 100%. Fig. 6 shows the change in power consumption for each server at both extremes compared to the mean power at 20-23.3°C at low pressure. Increasing inlet temperatures resulted in increased power consumption for every server under load. At idle, the impact is much less though still trending upwards at higher temperatures. Additionally, every higher pressure scenario has slightly improved power consumption against its low-pressure counterpart, though the effect is much smaller than temperature.

The R640 Xeon Silver 4116 is the only A3-class server and shows minimal impact of temperature or pressure under load. It is much more efficient and less sensitive to the environment than the higher power Gold 6148 (Fig. 5) despite its lower performance due to very low power consumption. This indicates that lower power density servers coupled with higher temperatures could realize savings in the cooling system without incurring the reductions in efficiency seen in other servers.

4.2 Effects on CPU

We can observe how the CPU temperatures respond to changes in the environment to give an indication of whether setpoint changes are influencing CPU behavior. If environmental conditions were to affect the lifespan of the CPU, we would expect to see increased CPU temperatures over the range. Additionally, one potential reason why a hotter environment will cause increased power consumption is due to increased leakage current in the CPU at higher CPU temperatures.

Fig. 7 shows the power and temperatures of the CPU during benchmarks reported by the internal registers. There are distinct effects due to the temperature and pressure of the environment, but these are relatively small compared to increasing utilization. Additionally, the changes in CPU temperatures do not perfectly mirror the changes in CPU power and neither does CPU power completely mirror the changes in overall server power (Fig. 6). Therefore, the increase in power consumption is not wholly attributed to CPU leakage current, and there are other components in the server also contributing to the increase.

Every CPU remained below its specified maximum temperature (TCase Max), with the exception of the E5-2690 and E5-2699 v4 at 100%, but there is no increase at higher ambient temperatures. Since there are no changes in performance either (Fig. 5) we can conclude this did not incur any throttling.

4.3 Internal Server Cooling System

There was no change in CPU temperatures with ambient conditions, indicating the server cooling system has increased fan speeds. This is confirmed in Fig. 8, where fan speed response to increasing utilization is shown in RPM and as a percentage difference normalized against the mean for each utilization in 20-23.3°C at low pressure. A portion of the increased server power can be attributed to increased fan activity; in some cases, fans were spinning 25-50% faster than the base case. Due to the cubic relationship between fan speed and power, this could mean up to a 300% increase in fan power alone. Here, we can see a more exaggerated impact due to pressure than seen elsewhere, since with a higher pressure differential across the server the fans do not need to spin as quickly to generate the airflow. Notably, the A3 server (R640 - Silver 4116) had almost no fan response compared to the other servers.

Changes in fan speed and internal temperatures also impact the change in air temperature across the server, known as delta-T. We noted a typical 1-4°C range on each of the server's delta-T, though the relationship between

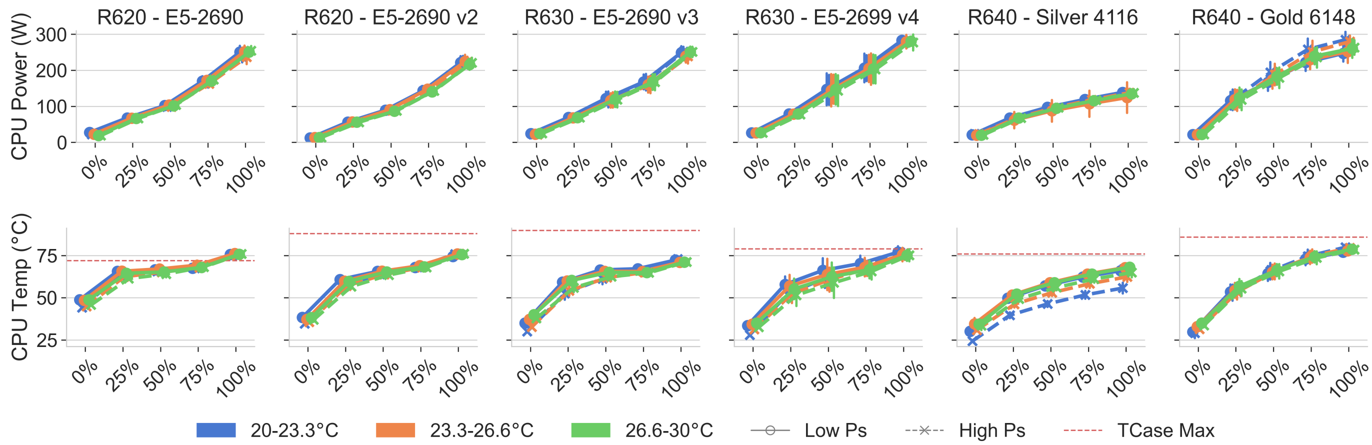

Fig. 7. Mean and 95% confidence intervals for CPU power and temperature as reported by internal CPU performance registers for each target benchmark utilization under the different environmental conditions.

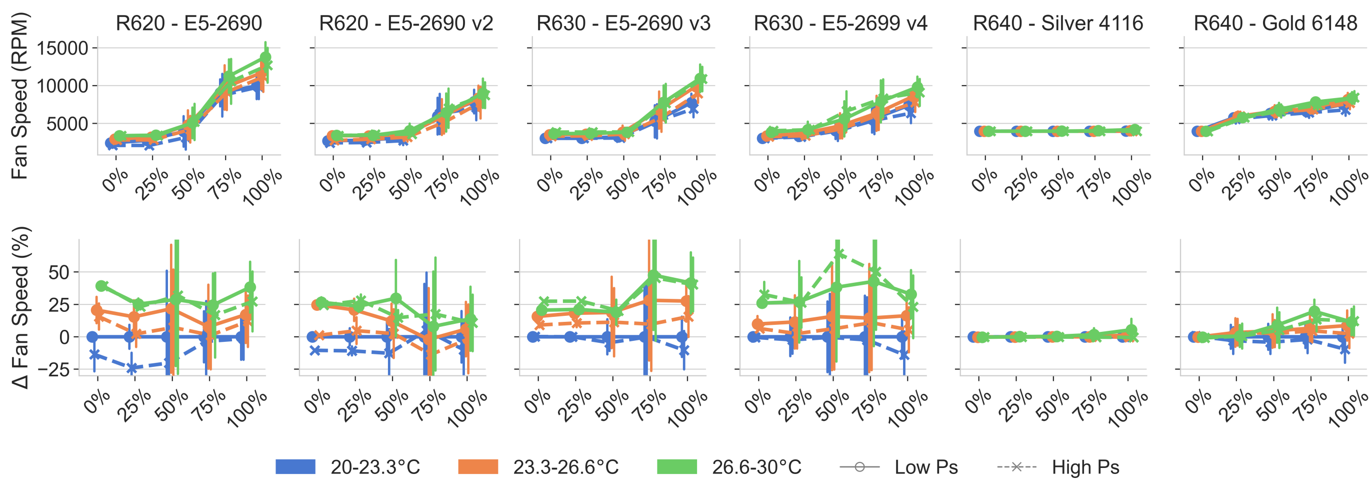

Fig. 8. Mean point estimate and 95% confidence intervals of absolute fan speed and percentage change at different target benchmark utilizations. Percentage change is calculated using the mean at each utilization for each server at 20-23.3°C and ≤5Pa pressure as a reference point.

CPU temperature, CPU power and fans speed is complex. For the newer servers, R640 and R630 Xeon 2699 v4 the delta-T decreased, indicating the smaller heat flux between the heatsink and that air was impacting cooling to some degree. Meanwhile, older servers, particularly the oldest R620 E5-2690 specification, exhibited the largest increases in delta-T (>4°C) due to the significant increase in power consumption of that server in the hotter environment.

4.4 Benchmark findings

Overall, we see that age, specification, and utilization all impact the server's efficiency and power consumption; these effects are compounded to a lesser but still significant extent by environmental factors like temperature and pressurization. Utilization is the most significant factor overall, and each generation behaves differently as this increases. The other factors compound this behavior in non-trivial ways, but a general trend is power consumption increases with ambient temperature by 2-6% at the highest temperatures, and a saving of 2-4% is available by utilizing higher pressure differentials across servers in aisle contained DCs.

5 Case Study: London

The second half of this paper deals with the real-world implications of the findings of section 4. If running data centers at higher inlet temperatures do not significantly impact failure rates or performance then the overall energy efficiency improvement depends on the saving in the cooling system outweighing the increased power from IT. Higher inlet temperatures open the door for massive savings and the capability to run data centers in a wider range of climates.

5.1 Methodology

A simulated idealized homogeneous data center, loosely based on DC3 at LDeX [40], [41] was the basis of our analysis. The Green Grid's Liquid Cooling Total Cost of Ownership Calculation Tool [42] was used to estimate the cooling costs as they account for IT power consumption, power density, room setpoint, different types of air cooling, and local climate. There are few tools available that can perform this analysis, this one was chosen for its ability to estimate these attributes and because its calculations are

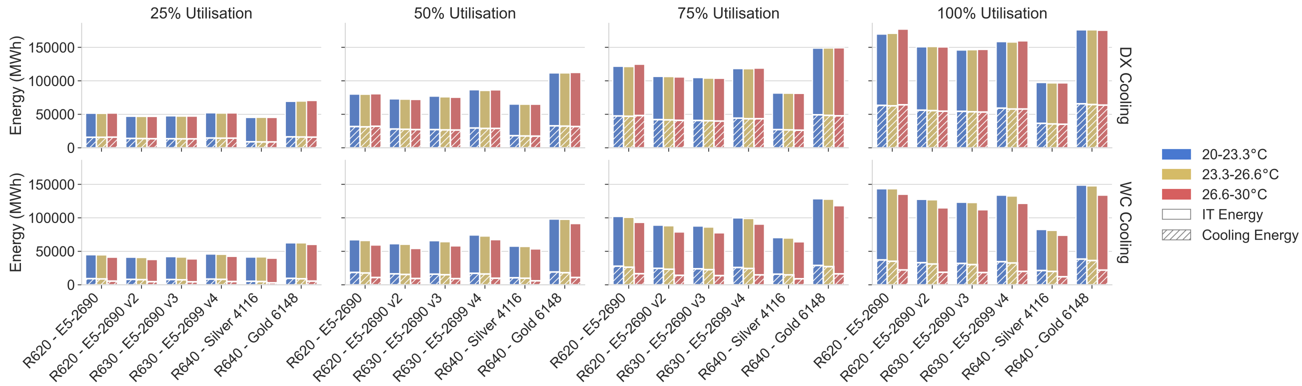

Fig. 9. Cooling and IT energy combined to show total annualized energy for a model data center in London using different servers, cooling technologies (DX:Top, Chiller:Bottom), and set points. In all cases, IT uses proportionally more energy than cooling and with DX cooling (top row) the increase in IT load often outpaces any decrease in cooling when increasing the temperature.

based on historical trends and established scientific principles.

Using all six server types considered in section 4, the model data center was comprised of 700 homogeneous rack cabinets, with power density ranging from 10kW per cabinet to 18.4kW per cabinet depending on the server. This puts the scenarios considered in the median to upper quartile of power density as reported by Uptime Institute [43]. The total IT power consumption is estimated at 7-13MW for a 300,000 sqft area facility. These six homogeneous data centers were then considered with DX or chiller cooling technology at set-points of 21.6°C, 25°C, and 28.3°C, taking the respective power consumptions and delta-Ts recorded for those servers at the set-points into account to calculate IT power consumption and power density for sizing the cooling systems.

The relationship between IT power consumption and cooling power consumption for these various temperature and hardware scenarios was considered at loads of 25%, 50%, 75%, and 100%, with the cooling systems themselves being sized based on 100% IT consumption proportionally scaled to match load percentage. All experimentally derived figures used in this analysis were of the low pressure results, as this seemed most analogous to the capacity of the Green Grid tool.

A total of 144 scenarios were evaluated across four variables:

1) 6 Server configurations: See Table. 2.
2) 3 Temperature set-points: 21.6°C, 25°C, 28.6°C
3) 4 Utilization rates: 25%, 50%, 75% and 100%
4) 2 Cooling systems: air-cooled DX, water-cooled chiller with water economizer

5.2 Observations

Like for an individual server, as utilization goes up at the DC level, so does the IT energy consumption and thus energy for the cooling systems (Fig. 9). As in section 4 the impact of utilization and temperature varies with server type. The small increase in power consumption with temperature set-point seen at the server level compounds when considered for the DC and has varying levels of impact on cooling requirements. In Fig. 9, we see the overall impact of increasing set-points on the DCs energy consumption. Particularly with DX cooling, we see very little change or even increases in total power for the DC as inlet temperature set-point increases.

5.2.1 Difference in cooling system behavior between DX and Chiller

When comparing DX and chiller cooling systems for this case study, the chiller system consumes considerably less energy than DX. Chiller cooling systems typically incur a greater capital expenditure but conversely benefit greatly from economies of scale; the larger the DC, the smaller that capital expenditure is as a proportion of operating costs.

From a performance standpoint, the behavior of the chiller is more susceptible to changes in temperature set-point than the DX system. As the temperature setpoint rises, the energy consumed by the chiller decreases at a greater rate of change than the DX, leading to a knock-on effect of a greater decrease in total facility energy consumption also.

Fig. 10 illustrates the change in energy consumption of the DX and chiller cooling systems as well as IT as a result of raising the setpoint to 28°C. The net effect is shown as a solid line for each method of cooling and for all servers. In many cases, for DX cooling, any savings from reductions in cooling power are negated by the increased IT energy (although significantly this is not always the case, such as the R640 - Gold 6148 at 100% utilization, or the R630 E5-2690 v3 at 50% utilization, both of which result in a net facility saving.) On the other hand, increasing the temperature set point for the chiller is such a large saving that the net total is always less than at 21.6°C. This speaks to the complexity of the interplay between IT and cooling systems in DC environments and how an awareness of the behavior of the IT is imperative to predicting the requirements of the systems in place to cool it.

6 Case Study: Impact in Global markets

As the demand for internet access grows globally, so does the requirement for its underlying infrastructure. Emerging markets in South America, Africa, and Asia face unique challenges afforded by the varied climates of their high population centers, but they also benefit from existing research and lessons learned in their creation and development.

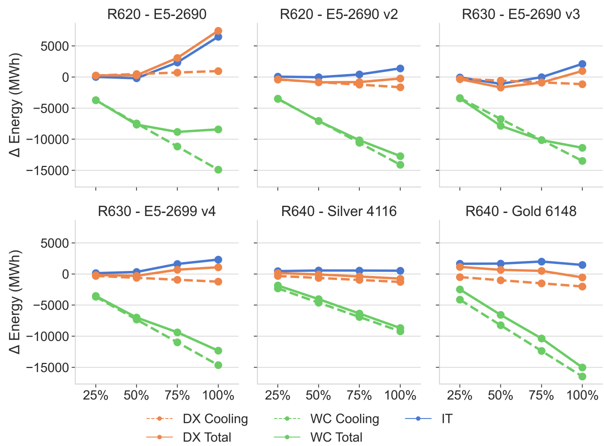

Fig. 10. Change in annual energy consumption of IT and cooling infrastructure (DX and WC) at varying total DC utilization when increasing the setpoint from 21.6°C to 28.6°C.

This portion of the case study considers transporting the previously discussed model DC to 24 locations across these three continents to analyze how energy consumption varies with climate. The locations were chosen to focus on both emerging markets as well as established markets in hotter climates.

ASHRAE International Climate Zones were developed by ASHRAE for Standard 90.1-2007 [44] to combine seasonal temperature and humidity information for a geographical location into a classification system for use in building management. The original release had a limited number of locations cataloged, but this database was built upon in future releases. For the purposes of this case study, standard 90.1-2016 and its database of 6,443 global locations was used to maintain consistency with the version used in the Green Grid tool. We selected 24 locations included in the 90.1-2016 Standard [45], and these are shown in Fig. 11 and Table. 3.

We have previously considered the interplay between the use and type of server and the impact this IT load has on cooling energy consumption, for these two types of cooling. By reducing the variation in servers down to their raw IT loads in order to more generally consider trends, we can analyze the other side of this coin - the relationship between IT energy consumption and cooling energy consumption as it relates to the different climates and cooling systems. This is in a sense analogous to a partial PUE for a DC in each of these climates, insofar as the gradient of each climate zones scatter points in Fig. 12 represent the ratio of cooling energy consumption to IT energy consumption. A shallow gradient suggests increasing IT power has a minimal impact on cooling and thus would be analogous to a low or “good” partial PUE, whereas a steep gradient demonstrates any change to IT load results in a significant increase in cooling load or a “bad” partial PUE.

At a glance, Fig. 12 seems to suggest all climates see a shallower gradient for this cooling/IT relationship as setpoint increases, although, in practice there is more nuance to this conclusion.

Fig. 13 shows the change in cooling energy consumption against IT load as the setpoint increases from 21.6 °C for both methods of cooling. For DX cooling, the greatest savings occur when increasing temperature setpoint in the warmest countries, such as those with climate zones 0B (hottest and dry) and 0A (hottest and humid). Conversely, when using a chiller system, the greatest savings to cooling energy consumption occur in the coolest climates.

This can perhaps be explained by considering the way in which these cooling systems maintain a setpoint within the DC. DX systems cool the air for a DC directly using refrigerant cycles within the air handling unit itself and will be running almost constantly against the significant difference in temperature between the setpoint and the ambient air temperature of climates such as Dubai or India. Increasing the setpoint within the DC, in turn, minimizes the temperature gradient a DX system is constantly fighting against and thus lowers the cooling energy requirements. While the same could potentially be true for the refrigerant cycles in a chiller system, the presence of an intermediate coolant – be it water or glycol – means any savings are vastly overshadowed by the benefits of creating a greater temperature gradient between the coolant reservoir and the setpoint. If a water tower full of 10°C coolant that has reached that temperature because of the ambient air temperature now only needs to cool a DC to 28.3°C instead of 21.6°C, the duty cycle of the compressors may drop significantly.

The flip side to this behavior occurs when increasing the setpoint for chiller systems in warmer climates. That some of these climates have average temperatures greater than these setpoints likely means thermal mass of the coolant becomes less of a contributing factor for savings, and in turn the increased IT energy load results in a higher demand for cooling than is saved by increasing setpoint. A similar behavior is likely responsible for the increased cooling seen for 28.3°C setpoint for the singular 4C climate represented in this study – that is to say, any energy savings made decreasing the cooling load by increasing setpoint are negated by the increased cooling load requirements of hotter IT.

It is important to take these savings and losses in context, however, for a DC of this size, we calculate chiller cooling to be consistently more efficient than DX cooling. Fig. 14 shows the absolute annualized energy consumption of the total DC – cooling and IT – for each of the 24 cities considered in this study and for both DX and chiller cooling. This is taken at 50% utilization of a fabricated server that occupies the median power draw of all six studied servers. This is best summarized by comparing the ‘best’ DX cooled example - Puerto Montt at 28.6°C at a total annualized energy cost of 28.07GWh – with the ‘worst’ chiller cooled example – Caracas as 28.6°C at a total annualized energy cost of 19.48GWh. This shows that in every location studied (for a data center of this size with this configuration) the greatest saving to be had would be a change of cooling system from DX to chiller. Even if this idealized data center did exist, such a change would incur heavy capital expenditure to retrofit the scale of which might outweigh the savings available. The important takeaway is not to underestimate the complexity of the interlocking systems that make up a data center – server type and age, load utilization and power density, type of cooling and setpoint, and geographic location are among a number of factors that need to be taken into consideration when analyzing the energy efficiency of a DC.

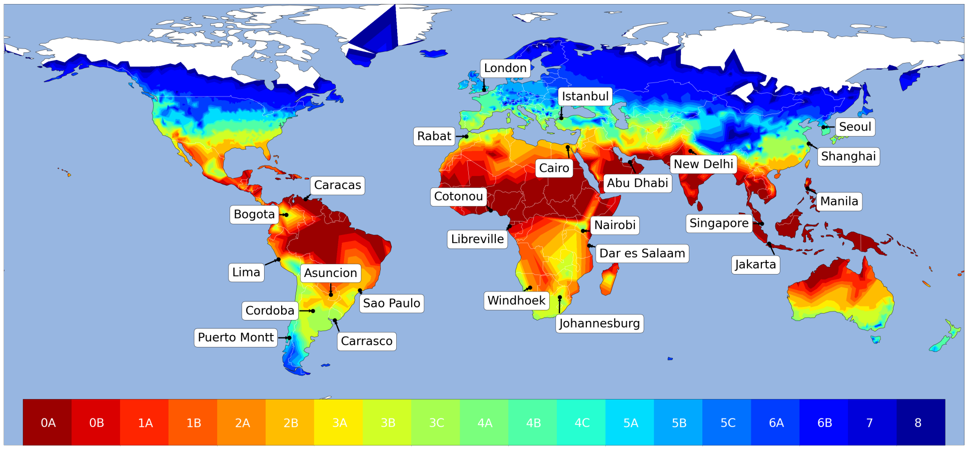

Fig. 11. The ASHRAE Climate Zone data applied to a world map, annotated with the 24 cities considered in the Case Study.

TABLE 3
Evaluated locations for the case study.

Country	City	No. of DCs	ASHRAE zone	ASHRAE max allowed MLC	Country	City	No. of DCs	ASHRAE zone	ASHRAE max allowed MLC
UK	London	196	4A	0.33	Namibia	Windhoek	1	2B	0.36
Chile	Puerto Montt	2	4C	0.32	Brazil	Sao Paulo	64	2A	0.35
Colombia	Bogota	21	3A	0.33	Singapore	Singapore	105	0A	0.37
Turkey	Istanbul	24	3A	0.33	Philippines	Manila	21	0A	0.37
South Korea	Seoul	32	4A	0.33	Benin	Cotonou	2	0A	0.37
Peru	Lima	17	2B	0.36	Indonesia	Jakarta	44	0A	0.37
Kenya	Nairobi	6	3C	0.32	Argentina	Cordoba	2	3A	0.33
South Africa	Johannesburg	31	3A	0.33	Egypt	Cairo	6	2B	0.36
Uruguay	Montevideo	4	3A	0.33	Venezuela	Carcas	3	0B	0.4
China	Shanghai	105	3A	0.33	Paraguay	Asuncion	1	2A	0.35
Gabon	Libreville	2	0A	0.37	India	New Dehli	17	1B	0.38
Morocco	Rabat	2	3A	0.33	United Arab Emirates	Abu Dhabi	4	0B	0.4

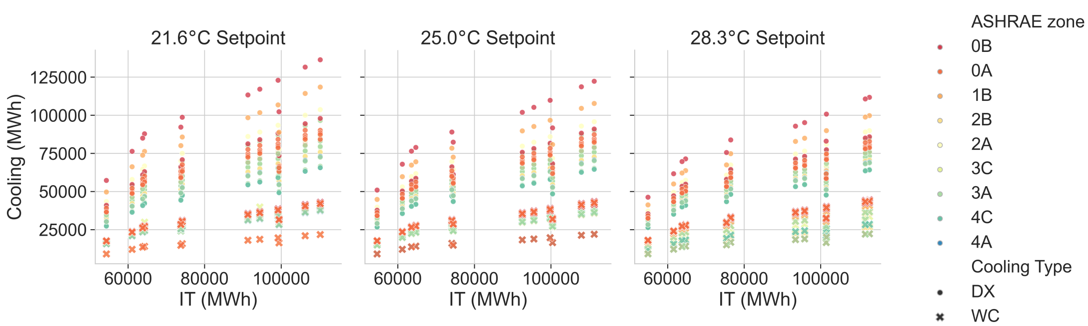

Fig. 12. Modelled cooling energy for DCs comprising the range of servers at different IT loads in each of the chosen locations.

7 Conclusion

We have profiled a number of typical enterprise rack servers while in varying environmental conditions within the manufacturer's guidelines and found:

1) regardless of generation, the server's efficiency is impacted by the environment
2) the compute performance remained unaffected
3) temperature had a much larger effect than pressurization, but both effects are measurable.

We analyzed how this behavior would impact DCs using a case study, to determine whether the increase in server power would offset gains from increasing DC temperatures. In general, cooling energy decreases with higher inlet temperatures and water-based chillers are more efficient but even in our idealized example, interactions between many factors influenced the DC-level power consumption, including utilization rate, cooling system type, server specification and local climate. Overall there were no universal, guaranteed savings from increased inlet temperatures, increases in IT energy can offset any beneficial decreases in cooling energy and therefore raising temperatures is potentially a false economy. This is particularly true with older IT hardware, but also depends on the local climate and cooling technologies, where we found different responses between DX and WC cooling in hotter climates. Our findings indicate that relationships and interactions between the different systems in the DC are complex and warrant further investigation including more profiling and real-world experiments. While there are intuitive guidelines for increasing efficiency, each

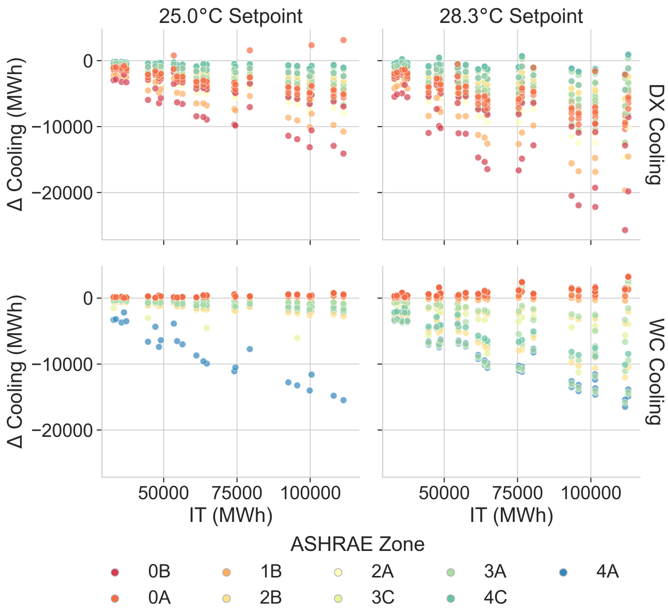

Fig. 13. The expected change in annual energy for each cooling system across the range of servers when increasing the set-point from 21.6°C.

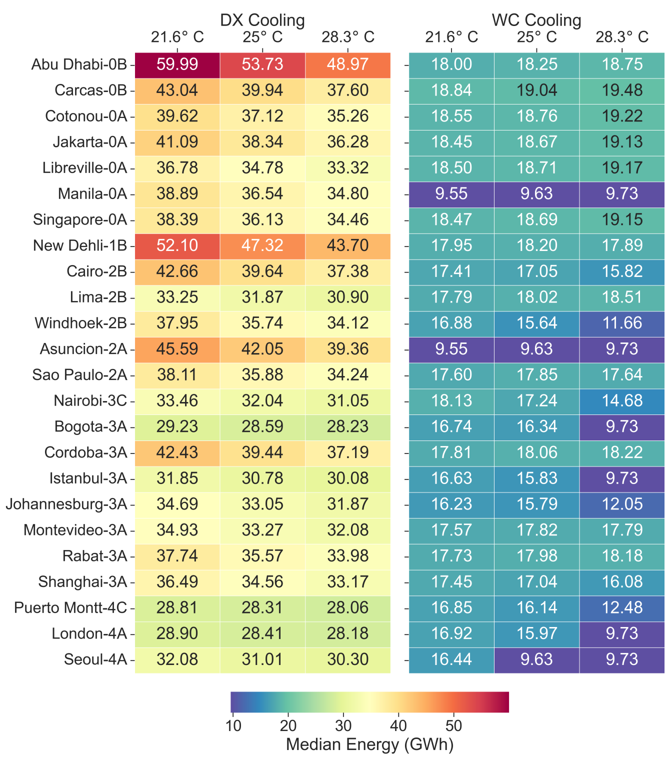

	DX Cooling 21.6° C	DX Cooling 25° C	DX Cooling 28.3° C	WC Cooling 21.6° C	WC Cooling 25° C	WC Cooling 28.3° C
Abu Dhabi-0B	59.99	53.73	48.97	18.00	18.25	18.75
Carcas-0B	43.04	39.94	37.60	18.84	19.04	19.48
Cotonou-0A	39.62	37.12	35.26	18.55	18.76	19.22
Jakarta-0A	41.09	38.34	36.28	18.45	18.67	19.13
Libreville-0A	36.78	34.78	33.32	18.50	18.71	19.17
Manila-0A	38.89	36.54	34.80	9.55	9.63	9.73
Singapore-0A	38.39	36.13	34.46	18.47	18.69	19.15
New Dehli-1B	52.10	47.32	43.70	17.95	18.20	17.89
Cairo-2B	42.66	39.64	37.38	17.41	17.05	15.82
Lima-2B	33.25	31.87	30.90	17.79	18.02	18.51
Windhoek-2B	37.95	35.74	34.12	16.88	15.64	11.66
Asuncion-2A	45.59	42.05	39.36	9.55	9.63	9.73
Sao Paulo-2A	38.11	35.88	34.24	17.60	17.85	17.64
Nairobi-3C	33.46	32.04	31.05	18.13	17.24	14.68
Bogota-3A	29.23	28.59	28.23	16.74	16.34	9.73
Cordoba-3A	42.43	39.44	37.19	17.81	18.06	18.22
Istanbul-3A	31.85	30.78	30.08	16.63	15.83	9.73
Johannesburg-3A	34.69	33.05	31.87	16.23	15.79	12.05
Montevideo-3A	34.93	33.27	32.08	17.57	17.82	17.79
Rabat-3A	37.74	35.57	33.98	17.73	17.98	18.18
Shanghai-3A	36.49	34.56	33.17	17.45	17.04	16.08
Puerto Montt-4C	28.81	28.31	28.06	16.85	16.14	12.48
London-4A	28.90	28.41	28.18	16.92	15.97	9.73
Seoul-4A	32.08	31.01	30.30	16.44	9.63	9.73

Fig. 14. Summary table of median total annual energy for the six servers at 50% utilization, comparing DX and WC cooling systems at varying operating temperatures by city. The shared color map indicates that despite different responses between the two cooling technologies, WC cooling uses less energy regardless of the local climate.

DC must collect data and evaluate these changes individually, as it is easy for changes made with the intention of saving energy to have the opposite effect overall.

Acknowledgments

This work was supported by a UKRI Future Leaders Fellowship [grant number MR/T04389X/1]

References

[1] IEA, "Data Centres and Data Transmission Networks," IEA, Tech. Rep., 2021. [Online]. Available: https://www.iea.org/reports/data-centres-and-data-transmission-networks

[2] A. Shehabi, S. J. Smith, D. A. Sartor, R. E. Brown, M. Herrlin, J. G. Koomey, E. R. Masanet, N. Horner, I. L. Azevedo, and W. Lintner, "United states data center energy usage report," *Lawrence Berkeley National Laboratory*, 2016.

[3] A. Korte, V. Tiberius, and A. Brem, "Internet of things (IoT) Technology Research in Business and Management Literature: Results From a Co-citation Analysis," *J. Theor. Appl. Electron. Commerce Res.*, vol. 16, no. 6, pp. 2073–2090, 2021.

[4] B. N. Silva, M. Khan, and K. Han, "Towards sustainable smart cities: A review of trends, architectures, components, and open challenges in smart cities," *Sustain. Cities Soc.*, vol. 38, pp. 697–713, 2018.

[5] B. Whitehead, D. Andrews, A. Shah, and G. Maidment, "Assessing the environmental impact of data centres part 1: Background, energy use and metrics," *Building and Environ.*, vol. 82, pp. 151–159, 12 2014.

[6] G. A. Brady, N. Kapur, J. L. Summers, and H. M. Thompson, "A case study and critical assessment in calculating power usage effectiveness for a data centre," *Energy Convers. Manage.*, vol. 76, pp. 155–161, 2013.

[7] *Power usage effectiveness*, ISO/IEC Std. 30 134-2:2016.

[8] Google, "Google environmental report 2021," Google, Tech. Rep., 2021. [Online]. Available: https://www.gstatic.com/gumdrop/sustainability/google-2021-environmental-report.pdf

[9] Facebook, "Sustainability report 2020," Facebook, Tech. Rep., 2020. [Online]. Available: https://sustainability.fb.com/wp-content/uploads/2021/06/2020_FB_Sustainability-Report.pdf

[10] S. R. Department, "Data center average annual pue worldwide 2007-2021," Statista, Tech. Rep., 2021. [Online]. Available: https://www.statista.com/statistics/1229367/data-center-average-annual-pue-worldwide/

[11] H. Singh, T. Reuters, D. Azevedo, S. D. Ibarra, R. Newmark, E. O. Steve, and J. Pflueger, "Data center maturity model," The Green Grid, Tech. Rep., 2011.

[12] M. Acton, P. Bertoldi, and J. Booth, "2022 best practice guidelines for the eu code of conduct on data centre energy efficiency," European Commission Joint Research Centre (JRC), Tech. Rep., 2022.

[13] A. Capozzoli and G. Primiceri, "Cooling systems in data centers: State of art and emerging technologies," *Energy Procedia*, vol. 83, pp. 484–493, 2015.

[14] C. Jin, X. Bai, and C. Yang, "Effects of airflow on the thermal environment and energy efficiency in raised-floor data centers: A review," *Sci. Total Environ.*, vol. 695, 2019.

[15] J. Sarkinen, R. Brännvall, J. Gustafsson, and J. Summers, "Experimental analysis of server fan control strategies for improved data center air-based thermal management," in *19th IEEE Intersoc. Conf. Therm. Thermomech. Phenom. Electron. Syst. (ITherm)*. IEEE, 2020, pp. 341–349.

[16] K. Ebrahimi, G. F. Jones, and A. S. Fleischer, "A review of data center cooling technology, operating conditions and the corresponding low-grade waste heat recovery opportunities," *Renewable Sustain. Energy Rev.*, vol. 31, pp. 622–638, 2014.

[17] E. Frachtenberg, D. Lee, M. Magarelli, V. Mulay, and J. Park, "Thermal design in the open compute datacenter," in *Proc. 13th IEEE Intersoc. Conf. Therm. Thermomech. Phenom. Electron. Syst. (ITherm)*. IEEE, 2012, pp. 530–538.

[18] "Open compute project." [Online]. Available: https://www.opencompute.org/

[19] A. Lawrence, "Data center pues flat since 2013," Uptime Institute, Apr 2020. [Online]. Available: https://journal.uptimeinstitute.com/data-center-pues-flat-since-2013/

[20] ASHRAE TC 9.9, "2011 thermal guidelines for data processing environments: Expanded data center classes and usage guidance," ASHRAE, Tech. Rep., 2011.

[21] S. Strutt, C. Kelley, H. Singh, and V. Smith, "Data center efficiency and it equipment reliability at wider operating temperature and humidity ranges," The Green Grid, Tech. Rep., 2012.

[22] D. Atwood and J. G. Miner, "Reducing data center cost with an air economizer," Intel Corporation, Tech. Rep.

[23] S. Krishnapura, T. Tang, V. Lal, and S. Achuthan, "Green computing at scale," Intel Corporation, Tech. Rep., Aug 2021. [Online]. Available: https://www.intel.com/content/dam/www/central-libraries/us/en/documents/intel-it-green-computing-at-scale-paper.pdf

[24] E. Pinheiro, W.-D. Weber, and L. A. Barroso, "Failure trends in a large disk drive population," 2007.

[25] N. El-Sayed, I. A. Stefanovici, G. Amvrosiadis, A. A. Hwang, and B. Schroeder, "Temperature management in data centers: Why some (might) like it hot," vol. 40, no. 1, 2012.

[26] X. Gao, Z. Xu, H. Wang, L. Li, and X. Wang, "Why "some" like it hot too: Thermal attack on data centers," *SIGMETRICS Perform. Eval. Rev.*, vol. 45, no. 1, p. 23–24, jun 2017.

[27] S. Alkharabsheh, U. L. Puvvadi, B. Ramakrishnan, K. Ghose, and B. Sammakia, "Failure analysis of direct liquid cooling system in data centers," *J. Electron. Packag.*, vol. 140, 6 2018.

[28] R. R. Schmidt and B. D. Notohardjono, "High-end server low-temperature cooling," *IBM J. Res. Develop.*, vol. 46, no. 6, pp. 739–751, 2002.

[29] N. Hardavellas, M. Ferdman, B. Falsafi, and A. Ailamaki, "Toward dark silicon in servers," *IEEE Micro*, vol. 31, no. 4, pp. 6–15, 2011.

[30] Y. Wang, D. Nörtershäuser, S. L. Masson, and J.-M. Menaud, "Potential effects on server power," *Wireless Netw.*, 11 2018.

[31] Y. Wang, D. Nörtershäuser, S. Le Masson, and J.-M. Menaud, "An Empirical Study of Power Characterization Approaches for Servers," in *Proc. 9th Int. Conf. Smart Grids, Green Commun. IT Energy-aware Technol. (ENERGY 2019*, Athens, Greece, Jun. 2019, pp. 1–6.

[32] ASHRAE, *Thermal Guidelines for Data Processing Environments, 2nd Edition*. ASHRAE, 2008.

[33] D. S. Burdett, "Analysis of Thermal and Compute Performance of Data Centre Servers," Ph.D. dissertation, University of Leeds, 2018.

[34] J. H. Bell and R. D. Mehta, "Contraction Design for Small Low-Speed Wind Tunnels," NASA, Tech. Rep., 1988.

[35] C. L. Wharton, "The influence of the proximity of a wall to the test section exit of a wind tunnel," Ph.D. dissertation, Georgia Institute of Technology.

[36] B. Åžahin, A. J. Ward-Smith, and D. Lane, "The Pressure Drop and Flow Characteristics of Wide-Angle Screened Diffusers of Large Area Ratio," *J. Wind Eng. Ind. Aerodyn.*, vol. 58, no. 1-2, pp. 33–50, 1995.

[37] D. Barratt and T. Kim, "A Banked Wide-Angle Diffuser with Application to Electrostatic Precipitators," *Proc. Inst. Mech. Eng., Part A: J. Power and Energy*, vol. 229, no. 1, pp. 88–98, 2015.

[38] K.-D. Lange, M. G. Tricker, J. A. Arnold, H. Block, and C. Koopmann, "The implementation of the server efficiency rating tool," in *Proc 3rd joint WOSP/SIPEW Int. Conf. Performance Eng. (ICPE)*. ACM Press, 2012, p. 133.

[39] R. Bashroush, N. Rteil, R. Kenny, and A. Wynne, "Optimizing server refresh cycles: The case for circular economy with an aging moore's law," *IEEE Tran. on Sustain. Comput.*, vol. 7, no. 01, pp. 189–200, jan 2022.

[40] Airedale, "London Data Exchange Case study," Airedale, Tech. Rep., 2020. [Online]. Available: https://www.airedale.com/wp-content/uploads/2020/01/LDeX-Case-Study-4.pdf

[41] (2022) London Data Exchange. [Online]. Available: https://www.ldexgroup.co.uk/uk-data-centres/london-datacentre-ldex1/

[42] (2018) Liquid cooling total cost of ownership. The Green Grid. [Online]. Available: https://www.thegreengrid.org/en/resources/library-and-tools/468-Liquid-Cooling-Total-Cost-of-Ownership-Calculation-Tool

[43] A. Lawrence. (2020) Rack density is rising. Uptime Institute. [Online]. Available: https://journal.uptimeinstitute.com/rack-density-is-rising/

[44] *Energy Standard for Buildings Except Low-Rise Residential Buildings*, ANSI/ASHRAE/IES Std. 90.1-2007, 2007.

[45] *Energy Standard for Buildings Except Low-Rise Residential Buildings*, ANSI/ASHRAE/IES Std. 90.1-2016, 2016.

Stephen Clement is a Senior Researcher at Techbuyer. He received BEng and PhD degrees in Computer Science from the University of York, in 2010 and 2015 respectively. Recently he was awarded a UKRI Future Leader Fellowship for holistic approaches to improving data center efficiency. His research interests include distributed simulations, digital twins and sustainable computing, specifically related to data centers.

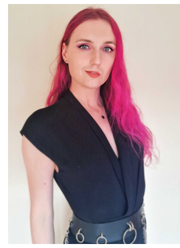

Kat Burdett is a Research and Development Engineer in the Research department at Techbuyer. Her research interests include computational efficiency, thermodynamics, green energy, and sustainable computing. Burdett has a PhD in Mechanical Engineering from the University of Leeds under the tutelage of Dr Jon Summers and has previously collaborated with the Research Institutes of Sweden (RISE).

Nour Rteil is the Lead Application Developer at Interact. She received BEng degree in computer and communications engineering from the American University of Beirut in 2015. In 2019, she joined the University of East London and Techbuyer as a Knowledge Transfer Partnership (KTP) associate to model the energy efficiency of servers. She is now the Lead Developer at Interact. Her research interests include green IT and sustainable computing.

Astrid Wynne is a Chartered Envrionmentalist, the sustainability lead with Techbuyer and head of sustainability at Interact. She has co-authored a number of academic and academic conference papers on the subjects of circular economy and energy efficiency for servers. She represents Techbuyer as a full partner on the Interreg NW funded CEDaCI (Circular Economy in the Data Centre Industry) project, which spans the UK, France, the Netherlands, Germany, Ireland, Belgium and Luxembourg. She regularly submits evidence to academic research projects and governmental bodies in the EU and UK on the subject of sustainable digital practices. She was part of the Consensus Body on the review of the American National Standards Institute revision of its Remanufacturing Standard.

Rich Kenny is the Sustainability and Research Director at Techbuyer and Managing Director at Interact. Named CTO of the Year (IT Hardware) in 2021 his research and development focus are hardware energy efficiency and the environmental and economic impacts of IT hardware during creation, use and disposal. He has contributed to the UK's Industrial Strategy through the Business, Energy and Industrial Strategy (BEIS) Committee and co-authored academic and conference papers on circular economy, energy efficiency and IT Performance. He represents Techbuyer on the Interreg NW funded CEDaCI (Circular Economy in the Data Centre Industry) project and leads research teams funded by both Innovate UK and UK Research Institute (UKRI).