Impacts of Increasing Temperature and Relative Humidity in Air-Cooled Tropical Data Centers

Duc Van Le, Jing Zhou, Rongrong Wang, Rui Tan, and Fei Duan

Abstract—Data centers (DCs) are power-intensive facilities which use a significant amount of energy for cooling the servers. Increasing the temperature and relative humidity (RH) setpoints is a rule-of-thumb approach to reducing the DC energy usage. However, the high temperature and RH may undermine the server's reliability. Before we can choose the proper temperature and RH settings, it is essential to understand how the temperature and RH setpoints affect the DC power usage and server's reliability. To this end, we constructed and experimented with an air-cooled DC testbed in Singapore, which consists of a direct expansion cooling system and 521 servers running real-world application workloads. This paper presents the key measurement results and observations from our 11-month experiments. Our results suggest that by operating at a supply air temperature setpoints of 29℃, our testbed achieves substantial cooling power saving with little impact on the server's reliability. Furthermore, we present a total cost of ownership (TCO) analysis framework which guides settings of the temperature and RH for a DC. Our observations and TCO analysis framework will be useful to future efforts in building and operating air-cooled DCs in tropics and beyond.

Index Terms—Data centers, temperature, humidity, energy usage, total cost of ownership analysis.

1 Introduction

Data centers (DCs) are power-intensive facilities which use up to 50% of their total energy usage for cooling the information technology (IT) equipment [1]. The main task of the cooling system is to maintain the supply air temperature and relative humidity (RH) conditions below certain setpoints. Increasing these setpoints is the rule-of-thumb approach to reducing the DC cooling energy usage [2], [3], [4]. However, many DCs are still operating at low temperature and RH setpoints in the ranges of [21°C, 27°C] and [45%, 70%], respectively. This is because the high temperature and RH may undermine the IT equipment's thermal safety, reliability and computing performance [5], [6], [7]. Moreover, the server power in general increases with the temperature because the server's internal fans rotate at higher speeds when the air temperature at server's inlet is higher [3]. The server's leakage power also increases with temperature [8]. As a result, adopting too high temperatures may not always lead to reduction of the total DC energy usage.

In this paper, we consider such a question: *what are the proper temperature and RH setpoints that balance the trade-off between the positive and negative impacts of raising temperature and RH setpoints on the total DC energy usage and IT equipment's reliability, respectively, in the tropical climate?* We focus on the tropical DCs because in the tropics with year-round high ambient temperature and RH levels, the DCs use more energy for cooling, compared with other locations in the world. For instance, in Singapore, the DC industry sector accounted for 7% of the country's total electricity usage in 2012, which is expected to reach 12% by 2030 due to the growth of digital economy [9]. In U.S., this ratio is 1.8% only in 2014 [10].

To answer the above question, it is essential to investigate how the temperature and RH setpoints varied in wide ranges affect the IT and cooling power usages and server reliability. Previous works [2], [3], [4], [5], [6], [7] conducted various studies to achieve this goal. However, these existing works mostly relied on the data collected from production DCs, in which the temperature and RH conditions are often maintained in narrow ranges. Moreover, they studied the DCs in the different areas with various climate conditions. Thus, their results may not cover high temperature and RH setpoints which may provide further opportunities to improve the DC energy efficiency with insignificant impact on the server reliability in the tropics. To this end, we constructed and experimented with a real DC testbed in Singapore with instantaneous ambient temperature up to 37°C and typical RH from 70% to 80%.

Our DC testbed consists of 521 servers deployed in 10 IT racks with a planned per-rack peak power rating of 8 kW. The deployed servers run real-world business applications managed by our industry partner. The testbed adopts a raised-floor DC design which employs a direct expansion (DX) cooling system consisting of two computer room air conditioning (CRAC) units with a maximum total cooling capability of about 80 kW. In the testbed space, multiple sensors in various sensing modalities are installed to monitor the real-time environmental condition of the IT racks and room as well as the power usages of the servers and cooling system. We conducted extensive experiments on the testbed to investigate the impact of the different supply air temperature and RH setpoints on the power usage of IT and cooling systems, server's reliability and safety. In particular, the temperature setpoint is varied from 20°C to 32°C with

- *Duc Van Le, Rongrong Wang, and Rui Tan are School of Computer Science and Engineering, Nanyang Technological University, Singapore 639798. Emails: {vdle,rrwang,tanrui}@ntu.edu.sg.*
- *Jing Zhou and Fei Duan are School of Mechanical and Aerospace Engineering, Nanyang Technological University, Singapore 639798. Emails: {jing_zhou,feiduan}@ntu.edu.sg.*

a step size of 1°C. In addition, we also study the impact of the other environmental factors including the air flow rate, outside temperature, and air quality (i.e., corrosion rate) based on the collected experimental results.

In this paper, we present the key measurement results and findings from our completed 11-month experiments. Specifically, our key observations show that the cooling and total DC power usages decrease by about 37% and about 16%, respectively, when the air supply temperature setpoint increases from 20°C to 32°C. The DC testbed achieves the minimum total energy usage with the supply air temperature setpoint of 29°C which leads to the server's annualized failure rate of about 2.3%. Moreover, when the supply temperature setpoint is higher than 23°C, without RH control, the supply air RH can be still maintained below 70% that is the optimal RH limit recommended by the American Society of Heating, Refrigeration and Air-Conditioning Engineers (ASHRAE) [11]. However, the DC operators and their tenants may have different mentalities and technical constraints in accepting high temperatures and RH levels because their servers may require different supply air conditions and costs for repairing/replacing the failed components. Thus, in this paper, we also propose a total cost of ownership (TCO) analysis framework to jointly consider all relevant DC operating costs affected by the temperature and RH conditions. Then, multiple cost-effective operating levels are defined to guide the temperature and RH settings in the DCs. These levels help facilitate the system implementation and the communication among DC operators and tenants. Finally, we apply the proposed framework to analyze the TCO costs of the testbed based on the collected measurement results.

As a systematic attempt of building and experimenting with a real DC testbed, our work has generated practical insights that the DC operators can adopt to improve the energy efficiency of their air-cooled DCs, especially in the tropical areas, like Singapore. In particular, our proposed TCO analysis framework has been adopted by a new Singapore Standard [12] on the operation of tropical DCs. While our work mainly focuses on improving the energy efficiency of the tropical DCs, our experimental results and observations can be also applicable to the DCs in other locations. For instance, the cooling energy saving obtained by raising the temperature setpoint in our work can be considered as the lower bound of the cooling energy saving obtained in other cooler locations. This is because our measurements show that the DC cooling power usage increases with the outside air temperature.

The main contributions of our work can be summarized as follows:

- We construct a real-world air-cooled DC testbed and conduct extensive experiments to study the impacts of various supply air conditions on the DC power usage and server reliability under the tropical climate conditions. Our experimental results and observations provide scientific evidence to promote the air-cooled high-temperature DCs with improved energy efficiency, especially in the tropical areas.
- We build multiple DC energy usage and server reliability models based on the data traces collected from the testbed. These models are useful to the DC research community which often lacks real-data-driven models due to the high cost of constructing and operating the real-world DC testbeds.
- We propose a TCO analysis framework which defines multiple cost-effective levels to guide the temperature and RH settings in the practical DCs.
- We release a dataset [13] of collected power usage and environment condition traces for public access. To the best of our knowledge, the existing publicly available DC-related datasets (e.g., Alibaba's [14] and Google's [15]) are merely about server workloads.

The remainder of this paper is organized as follows. §2 reviews the DC infrastructure background and the related work. §3 describes the design and experimental plan of the testbed. §4 discusses the experimental results and observations. §5 presents the proposed TCO analysis framework. §6 concludes this paper.

2 Background & Related Work

This section introduces the background on DC infrastructures in §2.1. Then, it reviews related work in §2.2.

2.1 DC Infrastructures

A DC is a complex cyber-physical system which consists of two main systems: the IT and cooling systems. The IT system hosts multiple servers, storage and networking devices that generate heat to be removed by the cooling system. For cooling, the DCs often adopt one of two mature cooling solutions [16]: *air cooling* and *direct liquid cooling*. An Uptime Institute's recent survey [17] shows that 92% of the 402 global respondents implemented the air cooling in their DCs. With the main objective of generating insights that can be useful to most DCs, our work focuses on investigating a real DC testbed with an air cooling system. Moreover, the air cooling can be implemented in three major forms [18]: *chilled water (CW)*, *direct expansion (DX)*, and *free-cooling* systems. Specifically, the CW and DX cooling systems often use a significant amount of energy since their cooling components (e.g., the CRACs, water chillers, cooling towers, condensers and compressors) are power-intensive. In general, the total cooling power usage comprising of the power usages of these components decreases with the supply air temperature setpoint [3]. Moreover, a high RH setpoint eliminates the need of dehumidification capability and its relevant costs.

In both the CW and DX cooling systems, the air is circulated through a closed loop inside the DC building, and not mixed with the outside air. Differently, the free-cooling system uses fans to directly blow the outside air into the server room, and then moves the hot air carrying the heat generated by the servers out of the DC. The free-cooling system is energy-efficient because it does not use the power-intensive cooling devices. In our previous work [19], we constructed and experimented with a real air free-cooled DC testbed in Singapore. Our 18-month experiments showed that with the free-cooling system, increasing the supply air temperature from 25°C to 37°C can lead to the cooling energy saving of about 38%. However, implementing an air free-cooled DC in tropics is challenging due to the high outside ambient air RH level and airborne contaminants.

TABLE 1
ASHRAE's recommended and allowable limits for inlet temperature and RH in DC environments.

Class		Max. Temperature			Max. RH			
		2004	2008	2011, 15, 21	2004	2008	2011	2015, 21
Recommend	All	25°C	27°C	27°C	55%	60%	60%	70%
Allowable	A1	NA	32°C	32°C	NA	80%	80%	80%
	A2	NA	35°C	35°C	NA	80%	80%	80%
	A3	NA	NA	40°C	NA	NA	85%	85%
	A4	NA	NA	45°C	NA	NA	80%	90%

From 2004 to 2021, ASHRAE has published five editions of its IT equipment thermal guidelines to suggest the temperature and RH settings in DCs.

Specifically, our analysis showed that the co-presence of the high RH, corrosive gases and particulate contaminations contained in the ambient air is the main reason of corrosion occurrence on the server's hardware, which leads to a server failure rate of about 30%. Moreover, our previous work also suggested that the conventional DCs that circulate the clear air in enclosed buildings can increase their temperature setpoints for better energy efficiency without degrading server reliability and computing performance. Meanwhile, in our Singapore's tropical area, the usage of the free-cooling system is limited, and most DCs still adopt a CW/DX cooling system. Thus, in this work, we construct a real DC testbed with a DX system such that the representative results can be obtained and widely adopted by the DC industry to pursue the high-temperature DCs with improved energy efficiency in our tropical areas.

2.2 Related Work

ASHRAE's Guidelines: To promote the high-temperature-RH DCs with improved energy efficiency, the ASHRAE [20] has been working over years to enlarge the allowable temperature and RH ranges of the air supplied to the IT equipment. Specifically, the ASHRAE defines four classes of IT equipment (i.e., A1, A2, A3, and A4) and introduces the recommended and allowable inlet temperature and RH limits for all classes, as shown in Table 1. In practical DC operations, such ASHRAE's temperature and RH setting suggestions are often considered as a recommendation. Determining the optimal temperature and RH settings for a particular DC needs to consider all combined impacts of the temperature and RH on the DC energy usage, server reliability and computing performance.

Impacts on DC Energy Usage: The amount of the energy saving obtained by increasing the temperature and RH setpoints depends on the type/configuration of the cooling system and the DC ambient air condition. Meanwhile, the server power in general increases with the temperature. The reason is that the IT equipment manufacturers deploy various algorithms to control the server's internal fan speed based on the temperature measured at server's inlet [21]. Such algorithms generally increase the server's fan speed when the inlet temperature is higher. Thus, with higher temperatures, the server uses more powers for rotating its internal fans at higher speeds. For instance, El-Sayed *et al.* [5] found that the server power consumption increases by 50% when the temperature increases from 30°C to 45°C. Similarly, increasing the RH setpoint can also reduce the energy usage of the dehumidifiers.

Impacts on Server Reliability: Too high temperature and RH may undermine the server reliability [2], [3], [4], [5], [6], [7], [22]. Specifically, server hardware components can be damaged due to too high instantaneous supply air temperatures. For instance, Patterson [2] reported that increasing the temperature by 10°C could reduce the capacitor's lifetime by half. Based on the data collected from seven different DCs of Google over two years, El-Sayed *et al.* [5] observed that the HDD errors and failures increase with the temperature slower than the Arrhenius model's prediction. In addition, ASHRAE has provided the relative failure rates of IT equipment called the *x-factors* for a specific temperature range to guide the choice of DC temperature setpoint [11]. In this work, we adopt the x-factor to interpret the annualized failure rate (AFR) of the servers at any temperature based on the measured AFR at a certain temperature. The details are presented in §4.2.

Furthermore, from the general intuition, the high RH level can lead to occurrence of the moisture condensation which potentially promotes corrosion of the server hardware materials, thus negatively affecting the server reliability. In 2016, Manousakis *et al.* [7] found that the RH has dominant impact on the server component failures in the air free-cooled DCs. For instance, the HDD failure rate is 9x more correlated with RH than temperature. The studies in [23], [24] reported that in the presence of the air corrosive gases (e.g., NO_2, SO_2, and H_2S) and particulate contaminations, the high RH can result in the increased server hardware failure rate. However, the existing research [23], [24] has shown that with clean air, the RH has insignificant impact on the server failure rate. Thus, ASHRAE allows a RH level up to 90% for modern servers compliant with A4 requirement in typical DCs in which the air is filtered and recirculated within the enclosed buildings.

Impacts on Server Computing Performance: Several studies [5], [22], [19] have set up testbeds using thermal chambers to evaluate the impact of the high temperature on servers' performance. For instance, El-Sayed *et al.* [5] observed that the throughput of all tested seven different HDD models does not drop when the ambient temperature is less than 40°C. Our previous work [19] observed that the servers can operate without computing performance degradation under the supply temperature up to 37°C. In summary, existing studies consistently show that temperature has little impact on the server's computing performance.

The above existing works have shown heterogeneous observations on the impacts of the temperature and RH on the DC power usage and server reliability because they relied on different data sources and/or DC designs and configurations. The results of these works cannot be directly applied for determining the optimal temperature and RH settings for the DCs in a specific region such as our tropical area. In this work, to achieve the goal, we design and experiment with a real DC testbed in which the temperature and RH setpoints can be varied in wide ranges. From the experimental results on the testbed, we build various models which characterize the relationship between the temperature and RH with the DC power usage and server reliability. Then, we propose a TCO analysis framework to guide the temperature and RH settings.

Sensor Networks for DC Monitoring and Manage-

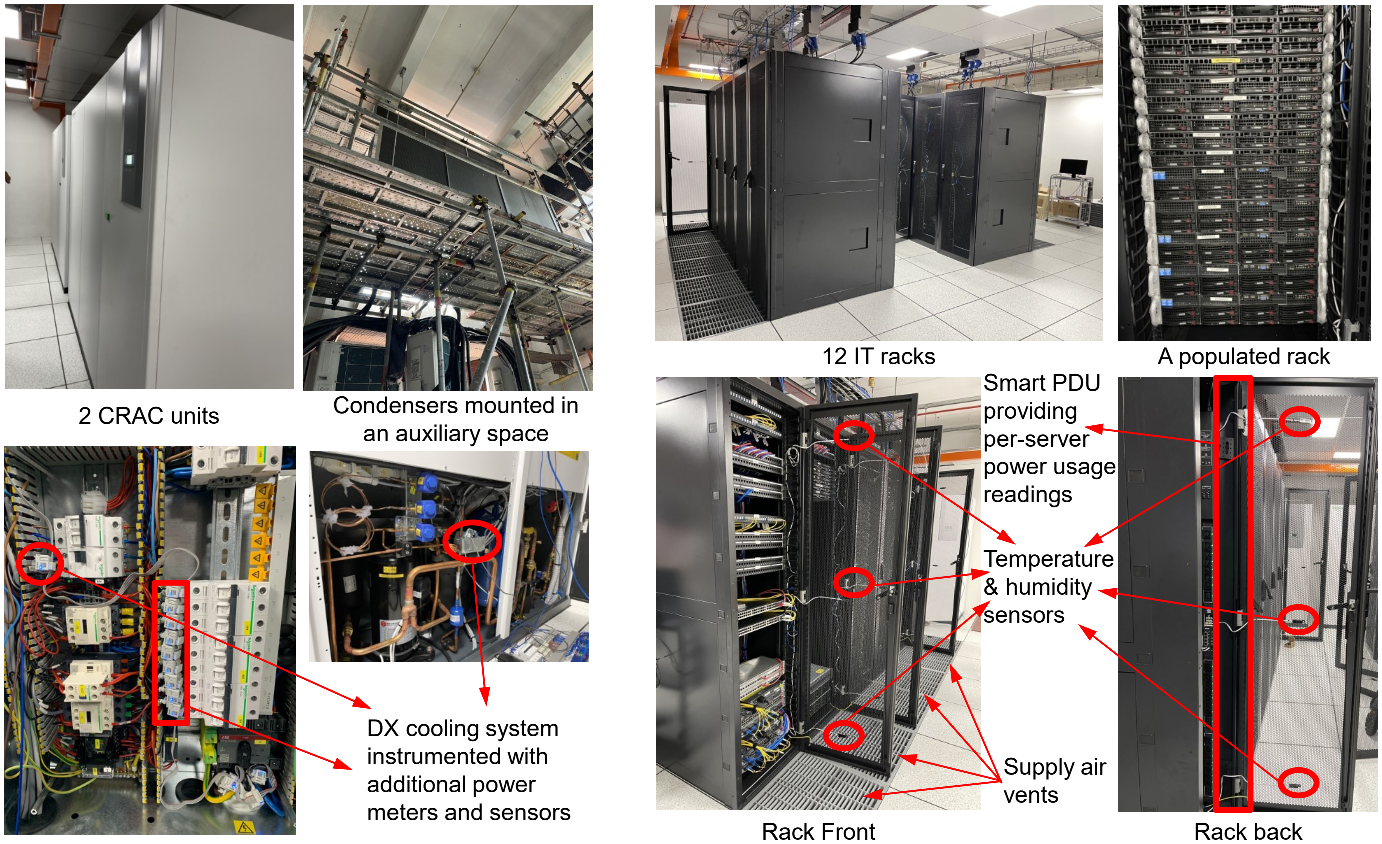

(a) Cooling System. (b) IT equipment and instrumented sensors.

Fig. 1. Cooling and IT systems of the testbed. In (b), the IT system consists of 10 racks hosting 521 servers and 2 backup racks.

ment: Existing studies [25], [26], [27], [28] have utilized sensor networks to achieve fine-grained monitoring of DC's operating conditions, which helps improve the reliability and efficiency of various DC management functions, such as the cooling control, power capping and preventative maintenance prediction. For instance, Liu and Terzis [25] implemented a sensing system using multiple sensor nodes called Genomotes to monitor the temperature and humidity conditions in a Microsoft's production DC. Moreover, Saifullah *et al.* [26] designed and deployed a wireless sensor network, called CapNet, which employs an event-driven control approach for real-time power capping in DCs. Similar to these existing studies, we also deploy multiple sensors to monitor the real-time operating conditions in our DC testbed. Moreover, these studies mainly leverage the sensor networks to achieve better DC management while our work analyzes the collected sensing data to study the impact of various environmental conditions on the DC power usage and server reliability.

TABLE 2
Description of deployed sensors in the testbed.

Sensor	Qty	Location
Temperature	36	6 for each rack*
RH	36	6 for each rack*
Temperature	1	outdoor
RH	1	outdoor
Temperature	2	CRAC's inlet and outlet
RH	2	CRAC's inlet and outlet
Air velocity	2	supply air vents
Corrosion sensor	1	cold aisle
Corrosion coupons	2	cold aisle
Power	3	DX cooling system
Power	24	2 for each IT rack

*Temperature and RH sensors are installed at three heights on the front and back sides of each rack.

3 Design and Experimental Plan of Testbed

In this section, we describe the design and configuration of the testbed. Then, we present our experimental plan.

3.1 Design of Testbed

The main design objective of the testbed is that we can maintain the temperature and RH of the air supplied to the IT equipment at certain setpoints for a period of time (e.g., from several days to several years). Moreover, the testbed can adjust the supply air temperature and RH setpoints in wide ranges, such that we can evaluate the IT and cooling energy usages, and the reliability under various supply air conditions. To meet these objectives, we design a DC testbed consisting of a server room sized $9.48m \times 8.07m \times 3.40m$ in the premise of a commercial DC operator. Fig. 1 shows the cooling and IT systems of the testbed. Specifically, the

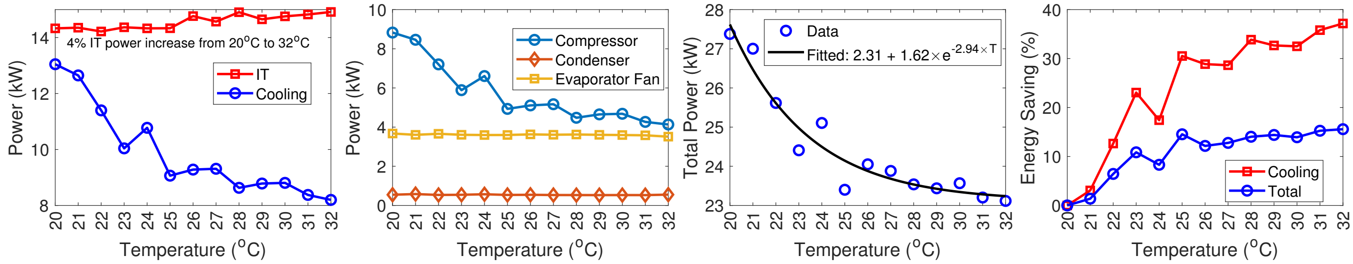

(a) IT and cooling powers. (b) Cooling power breakdown. (c) Total power. (d) Energy savings.

Fig. 2. Energy profile of the testbed with the temperature control only under different supply air temperature setpoints. In (a), (b) and (c), each data point is the average power usage over a period of 2 days. In (c), the fitted total power usage curve is an exponential function of the supply air temperature setpoint, denoted by T.

TABLE 3
Experimental plan on the testbed.

Test (Duration)	Controlled Parameters		
	Temperature	RH	Airflow rate (m^3/h)
Energy (94 days)	[20°C, 32°C]	[50%, 75%]	13600, 16500
Reliability (1 year)	29°C	No control	13600

testbed employs a DX cooling system consisting of two CRAC units with a total cooling capacity of about 80 kW. The mostly dynamic stage of the testbed is the cooling system. In particular, the stability of the cooling stage depends on the control algorithm which is used to control the operation of the DX cooling system to maintain the desired supply air temperature setpoints. We worked with a cooling equipment vendor to investigate various commercial DX cooling system products. The objective is to select a DX system which is capable of precisely maintaining every supply air temperature setpoint in the planned range of [20°C, 32°C] with a minimum control error. Finally, the DX system with a control error within 1°C is selected for the testbed.

In the testbed, we installed 12 45U IT racks which are aligned in two rows without hot/cold aisle containment as shown Fig. 1(b). A total of 521 server nodes from four different IT manufacturers are deployed on 10 racks. The planned per-rack peak power rating is 8 kW. The rest of 2 racks hosted 25 network devices (i.e., the switches and routers). During our experiments, the deployed servers run real business applications managed by our industry partner. In the testbed, we deployed multiple sensors and power meters with the quantity and locations as summarized in Table 2. Specifically, the sampling interval of these environmental sensors are set to 30 seconds. The collected sensor measurements are streamed into a server which has a web interface for real-time data visualization and downloading

3.2 Experimental Plan

Table 3 summarizes our experimental main test plan on the testbed. Specifically, we plan to conduct the following two groups of experiments which are *energy* and *reliability* tests. The energy test focuses on profiling the energy usages of the IT and cooling systems under various conditions of the supply air. During an energy test unit, the supply air temperature, RH and air volume flow rate setpoints are maintained at a certain level for a specific time period. Specifically, we

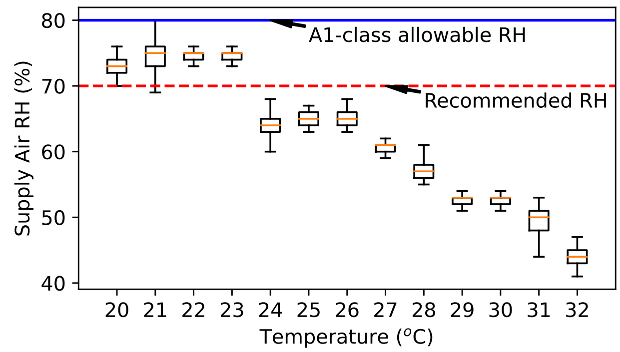

Fig. 3. Supply air RH distribution at different temperature setpoints. The box, line, upper and lower whiskers represent the middle 50%, average, ranges for the bottom 25% and the top 25% of the RH samples, respectively.

define the temperature setpoint as the temperature of the air that the CRACs supply to the server room. The CRACs use the temperature values measured by a temperature sensor at their inlet as environment feedbacks for their control algorithm to maintain a desired temperature setpoint. The RH setpoint is defined as the maximum supply air RH level that the CRACs can maintain during the test. The air volume flow rate setpoint is the volume speed of the air that the CRACs blow to the server room.

The energy test consists of multiple unit tests with combinations of the temperature, RH and air flow rate each sweeping the respective range as shown in Table 3. So far, we have completed all planned energy test units in 94 days. Moreover, the goal of the reliability test is to measure the server failure rate at the certain supply air condition over a time period of 1 year.

4 Experimental Results

In this section, we present the energy usages of the IT and cooling systems under various supply air conditions. Specifically, each power data result is the average power which is calculated by dividing the measured total energy usage by the test period. Then, we present the server reliability results. Specifically, we present the detailed forms of multiple DC energy usage and server reliability models that are built based on the collected data traces. As such, these models can be easily adopted for other DC studies.

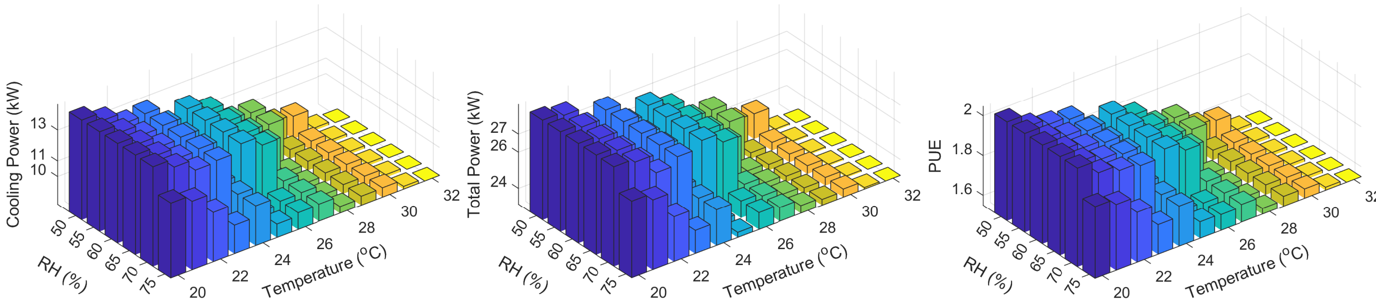

(a) Cooling power usage. (b) Total power usage. (c) PUE.

Fig. 4. Energy profile and PUE of the testbed with both the temperature and RH control.

4.1 Energy Test

4.1.1 Temperature control

We conducted the first experiment set to understand the energy profile of the testbed with the temperature control only. Fig. 2 shows the power usage of the IT and cooling systems, total DC power usage, and energy savings when the supply air temperature setpoint varies from 20°C to 32°C during a 26-day experiment. Each temperature setpoint is maintained for 48 hours. From Fig. 2(a), when the temperature setpoint increases from 20°C to 32°C, the IT power usage is mostly stable and then increases when the temperature is higher than 25°C. For instance, the IT power usage increases by about 4% when the temperature increases from 20°C to 32°C. The main reason is that the server's internal fans rotate faster when the inlet temperature is higher. Note that the tested servers run the real-world applications that may have different computation workloads over time. Thus, the IT power usage has variance at the close temperature setpoints which are tested at different time durations.

Meanwhile, as shown in Fig. 2(a), the cooling power usage has variance but a decreasing trend when the temperature increases. Specifically, the cooling power usage is a sum of power usages of the three main components of the DX cooling system including the compressor, condenser and evaporator. From Fig. 2(b), increasing the supply air temperature leads to reduction in the power usage of the compressor that drives the refrigerant through the closed-loop DX system as described in §2.1. The main reason is that a lower supply air temperature setpoint requires a higher flow rate of the refrigerant driving through the evaporator to absorb more heat from the server room. As a result, the higher temperature setpoint requires the lower refrigerant flow rate, which reduces the demand for the electrical motor load of the compressor and its energy usage. Thus, the cooling power usage in generally decreases with the temperature setpoint. However, among the close temperature setpoints, the cooling system may consume more power to maintain a higher temperature. This is because the compressor's power usage also increases with the outside air temperature. We will explain the reason in §4.1.4. Meanwhile, the outside air temperatures may be different over time. For instance, our measurement shows that the average outside temperature over 2-day periods when testing the temperatures of 23°C and 24°C are 30.64°C and 32.1°C, respectively. As a result, the average cooling power of 24°C is about 0.775kW higher that of 23°C as shown in Fig. 2(a).

In summary, although the IT power usage increases, the total DC power usage (i.e., sum of the IT and cooling power usages) has an exponential decreasing relationship with the supply air temperature as shown in the fitted curve of the total power usage in Fig. 2(c). Moreover, from Fig. 2(d), we can see that the testbed achieves the cooling and total energy savings of about 37% and about 16%, respectively, when the temperature setpoint increases from 20°C to 32°C. This result indicates that increasing the temperature setpoint can help improve the overall energy efficiency of the DCs.

Observation 1: When the air supply temperature setpoint increases from 20°C to 32°C, the IT power usage increases by about 4%, while the cooling and total power usages decrease by about 37% and about 16%, respectively. The total power usage has an exponential decreasing relationship with the supply air temperature setpoint within a range of [20°C, 32°C].

Fig. 3 shows the box plots for the distribution of the supply air RH measured by an RH sensor with a sampling interval of 30 seconds over 2 days when the temperature setpoint varies from 20°C to 32°C. From Fig. 3, we can see that the supply air RH is varied in a narrow range under each temperature setpoint. This is because in the above temperature control experiments, we do not perform the RH control. In other words, the cooling system does not introduce/remove the water vapor to/from the air supplied to the servers. Moreover, the air is circulated within the enclosed building and not mixed with the outside air. As a result, the testbed's air mostly contains the similar amount of water vapor during the experiments. The RH is the ratio of the amount of moisture contained in the air at a certain temperature to the maximum amount of moisture that the air can hold at the same temperature. Thus, under a certain temperature setpoint, the RH remains stable due to the similar amount of water vapor contained in the air. Moreover, as shown in Fig. 3, the RH in general decreases with the temperature. The reason is that the hot air can hold more moisture than the cold air. In addition, from Fig. 3, the supply air RH is always less than the ASHRAE's A1-A2 class allowable RH level of 80% [11]. The A3 and A4 classes have less stringent allowable RH levels of 85% and 90%, respectively. As mentioned earlier, these allowable RH levels are the maximum RH levels under which the servers are expected to function properly. Moreover, when the supply air temperature setpoint is greater than 23°C, the supply air RH is always less than 70% that is the optimal RH level

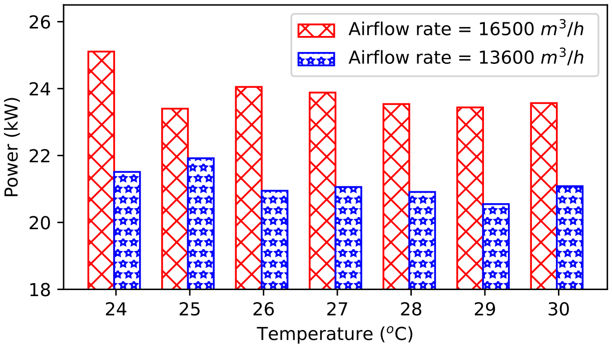

Fig. 5. Total power usage with the air volume flow rate control.

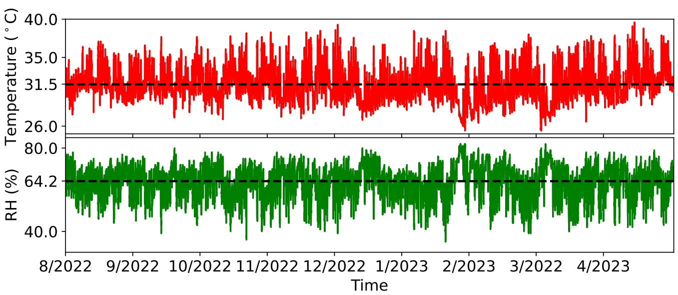

Fig. 6. Outside air condition during August 2022 to April 2023. In top and bottom sub-figures, the dotted lines represent the average outside air temperature and RH of 31.5℃ and 64.2%, respectively.

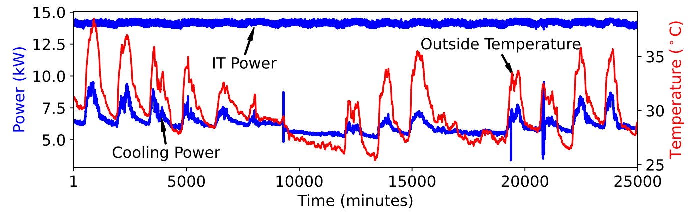

Fig. 7. Traces of the outside air temperatures, IT and cooling power usages over 25,000 minutes.

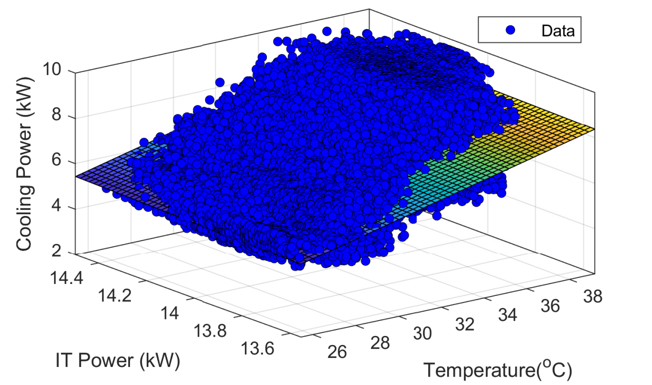

Fig. 8. Fitted cooling power function of $P_c(\text{kW}) = -3.441 + 0.228T_{\text{out}} + 0.2143P_{\text{IT}}$, where T_{out} and P_{IT} are the outside temperature and IT power, respectively. The fitting root mean square error (RMSE) is 0.809 (kW) over 269,095 input-output samples collected during 103 days.

recommended by ASHRAE.

Observation 2: The supply air RH can be maintained below the ASHRAE's recommended RH level of 70% without the RH control when the supply temperature setpoint is greater than 23°C. This indicates that increasing the supply air temperature setpoint not only reduces the energy usage but also eliminates the need of the RH control and its relevant costs.

4.1.2 *RH control*

As mentioned earlier, we deploy an extra dehumidifier to control the supply air RH condition in the testbed. Specifically, with each combination setting of the temperature and RH setpoint, if the DX cooling system without the RH control cannot maintain the supply air RH below the required RH setpoint, the dehumidifier is activated to remove the moisture from the air. As a result, the RH level can be reduced.

Fig. 4 shows the cooling and total power usages, and PUE of the testbed when the supply air temperature and RH setpoints vary in ranges of [20°C, to 32°C] and [50%, 75%], respectively. In these tests, the cooling power usage is a total power usage of the DX cooling system and dehumidifier. From Fig. 4(a), under a certain RH setpoint, the cooling power usage decreases with the temperature setpoint. Moreover, under each temperature setpoint, the cooling power usage remains at groups of the close RH setpoints and has decreasing trend with the RH setpoint. This is because when the less stringent RH level is required, the dehumidifier is not needed. Moreover, under a certain temperature setpoint, the IT power usage mostly remains the same when the RH setpoint varies from 50% to 75%. Thus, the total power usage and PUE have the same trend with the cooling power usage as shown in Figs. 4(b) and (c).

4.1.3 *Air volume flow rate control*

In the above temperature and RH control tests, we fixed the supply air volume flow rate at 16500 m^3/h. This setpoint is recommended by our cooling equipment vendor and is 85% of the maximum volume air flow rate that the testbed's DX cooling system can provide. To further explore energy saving opportunities, we conducted additional temperature control experiments with a reduced air volume flow rate of 13600 m^3/h (i.e., 70% of the cooling system's maximum flow rate). Moreover, to reduce the impact of the IT workload variation on the test results, we run each test unit for a longer time period of six days.

Fig. 5 shows the total power usage of the testbed over variations of the supply air temperature and volume flow rate setpoints. From Fig. 5, decreasing the air volume flow rate helps achieve significant energy saving. For instance, with the temperature setpoint of 29°C, reducing the air volume flow rate from 16500 m^3/h to 13600 m^3/h (i.e., a reduction of 2900 m^3/h) can lead to a power usage reduction of 2.89 kW i.e., an energy saving of 12.3%. Moreover, among the tested temperature setpoints from 24°C to 30°C, the 29°C is the optimal temperature setpoint with which the testbed can achieve the minimum total energy usage. Moreover, with the temperature setpoint of 29°C, the testbed can maintain an average supply air RH of about 55% (cf. Fig. 3) without the RH control.

Observation 3: The DC testbed can achieve the minimum total energy usage with the supply air temperature

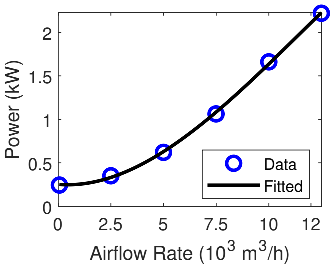

(a) Fan power in free-cooling.

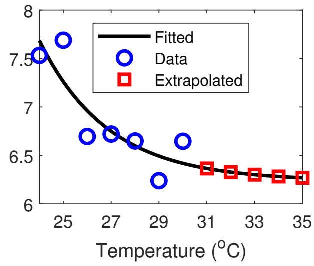

(b) DX cooling power.

Fig. 9. Fitted functions of the fan (i.e., free-cooling system) and DX power usages. In (a), the fitted fan power is $P_{\text{fan}}(\text{kW}) = 0.4769 - 1.265 \times 10^{-5} F + 3.05 \times 10^{-8} F^2 - 6.139 \times 10^{-13} F^3$ where $F(\text{m}^3/\text{h})$ is the supply air flow rate. In (b), the DX cooling power is $P_{\text{DX}}(\text{kW}) = 6.237 + 5.925 \exp^{-3.465T}$ where T (℃) is the supply air temperature setpoint and the air flow rate is set to 13600 m^3/h.

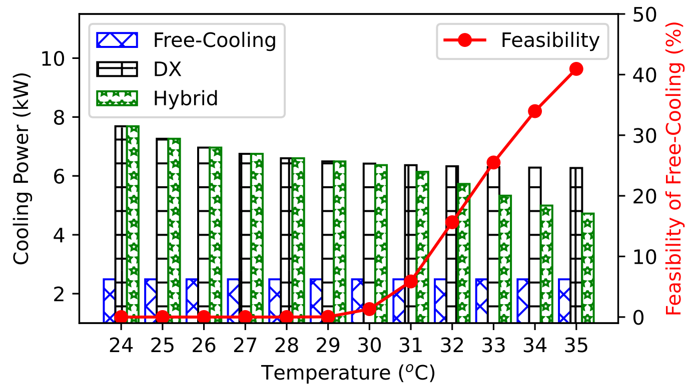

Fig. 10. Power usages of DX and hybrid (DX + free-cooling) systems under different supply air temperature setpoints and a fixed air volume flow setpoint of 13600 m^3/h.

and volume flow rate setpoints of 29°C and 13600 m^3/h, respectively. With these optimal setpoints, the supply air RH is always less than the ASHRAE's recommended RH level of 70%.

4.1.4 *Impact of outside air condition in tropics*

This section studies how the outside air condition of our tropical area affects the power usage of the testbed's DX cooling system. Fig. 6 shows the outside air temperature and RH traces over 8 months from August 2022 to March 2023 in the testbed area. We can see that the outside air temperature and RH fluctuate around the average values of 31.5°C and 64.2%, respectively, with the instantaneous levels up to 39°C and 80%, respectively.

Fig. 7 represents traces of the outside temperature, IT and cooling power usages over a time period of 25,000 minutes (i.e., about 14 days) when the supply air temperature is maintained at 29°C. From Fig. 7, the cooling power usage mostly increases with the outside air temperature. This is because in order for the condenser to dissipate the heat from the refrigerant to the ambient environment, the temperature of the refrigerant is required to be higher than the temperature of the outside air surrounding the condenser coils [29]. Higher temperature difference between the outside air and refrigerant temperatures makes the condenser reject the heat more efficiently. Thus, with a higher outside air temperature, the compressor has to use more powers to raise the temperature of the refrigerant for efficient heat dissipation. Fig. 8 shows the DX cooling power function of the IT power and outside air temperature, which is fitted using the 103 days' data traces collected from the testbed with a fixed supply air temperature setpoint of 29°C. We can see that the power usage of the DX system has a linear relationship with the outside air temperature.

The above results indicate that in the tropics with the year-round high ambient temperature and RH levels, the DCs with a DX system use more cooling energy to maintain the same temperature setpoint, compared with other cold locations in the world. In addition, the tropical climate also reduces the opportunity of using the energy-efficient air free-cooling system which utilizes the natural outside air to cool the server. As mentioned earlier, our previous work [19] constructed and experimented with a real air free-cooled tropical DC testbed which uses a fan system to supply the outside air into the server room to take away the heat generated by the servers. Fig. 9(a) presents the fan power usage model as a function of the supply air volume rate, which is fitted using the data traces collected from our testbed. We can see that the fan (i.e., free-cooling) system uses much lower energy than the DX system shown in Fig. 9(b). However, the usage of the free-cooling in tropics is limited, especially when the desired supply air temperature and RH setpoints are low. Thus, the DCs often adopt a hybrid scheme of the free-cooling and the traditional cooling (e.g., the DX or CW) systems [18]. When the outside air temperature and RH are higher than the desired temperature and RH setpoints, the traditional cooling system is activated to cool the servers. Otherwise, the free-cooling system is used to directly blow the outside air into the server room.

Fig. 10 shows the power usages of the free-cooling, DX and hybrid systems as well as the feasibility of the free-cooling in our tropical area when the supply air temperature setpoint varies from 24°C to 35°C. The supply air flow rate and RH setpoints are set to 13600 m^3/h and 65%, respectively. Specifically, our previous investigation [19] suggested that under the free-cooling system without the high-class filters, the particulate contaminants containing in the outside air will absorb the moisture and cause corrosion on the server hardware materials when the supply air RH is higher than the contaminants' deliquescent RHs. Thus, in this evaluation, we set the RH setpoint to 65% which is the deliquescent RH of many contaminants according to the ASHRAE's guidelines [30] on the air quality management for the DCs. Moreover, it is assumed that the free-cooling system adopts the fan system of our previous free-cooled DC testbed while the DX system of this work is considered the traditional cooling system. The hybrid system iteratively uses the free-cooling or DX system depending on the outside air condition. Moreover, the feasibility of the free-cooling is calculated as a percentage of samples in which the outside temperature and RH levels do not exceed the desired temperature and RH setpoints, respectively, over the 9-month outside air condition traces as shown in Fig. 6.

From Fig. 10, we can see that in our Singapore's tropical area, the opportunity of using the free-cooling system is limited. For example, the free-cooling feasibility is around

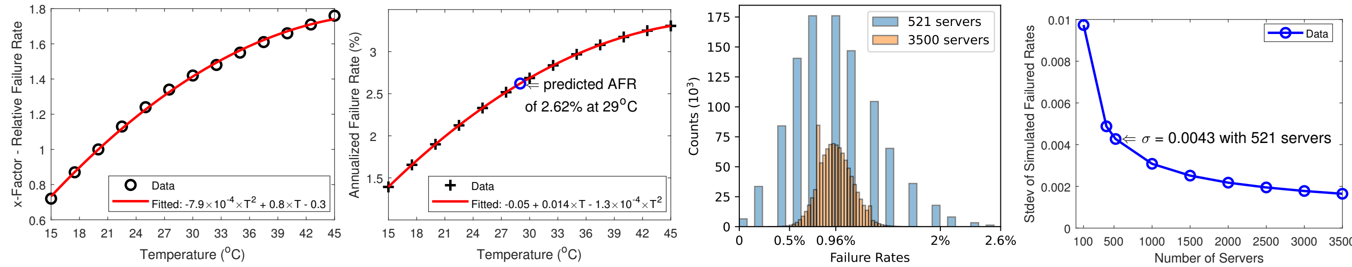

(a) x-factors. (b) Predicted AFRs. (c) Distribution of simulated FR_Υ. (d) Standard deviation σ.

Fig. 11. Server reliability results. In (a) and (b), the fitted x-factor and AFR curves are quadratic functions of the temperature setpoint of T.

zero when the temperature is less than 31°C. This result implies that in Singapore, the energy-efficient free-cooling system cannot be used if the DCs adopt the ASHRAE's recommended maximum temperature setpoint of 27°C. This is because the free-cooling system cannot satisfy the supply air temperature and RH requirements. Meanwhile, the DX and hybrid cooling systems can maintain the supply air temperature and RH below desired setpoints at the cost of higher power usages, compared with the free-cooling system. Moreover, with the temperature less than 31°C, the power usages of the DX and hybrid systems are the same due to the infeasibility of the free-cooling.

Observation 4: The tropical climate increases the power usage of the DX cooling system and limits the opportunity of using the free-cooling scheme. In Singapore, the free-cooling scheme is infeasible at all time when the supply air temperature setpoint is less than 31°C while the supply air RH setpoint is set to 65%.

4.2 Server Reliability Test

4.2.1 *Experiment settings*

In the above energy tests, we did not observe any server overheating, shutdown, and failure, especially when the supply air temperature setpoint increases up to 32°C. In general, the servers will automatically shutdown when the temperature measured in their inlet exceeds a certain safety threshold (e.g., 45°C) to avoid the overheating and hardware damages. Thus, this observation means that all tested servers can safely work with the temperature setpoint up to 32°C. Thus, we decided to maintain the sweet-spot supply air temperature and air flow rate setpoints of 29°C and 13600 m^3/h, respectively, in the reliability test for a period 1 year. We measure the server reliability in terms of the annualized failure rate (AFR) during this reliability test.

4.2.2 *An approach to predict the AFRs*

The server reliability test for a certain temperature often requires a long test period (e.g., 1 year) to obtain the statistically significant result. Running such long-term reliability test for temperatures in a wide range is challenging due to the high DC operating cost. Thus, in this work, we run the reliability test and measure the AFR of tested servers for a temperature of 29°C only. Then, we develop an approach that can predict the AFRs for other temperatures based on the measured AFR and the ASHRAE's x-factors [11] as follows. Specifically, the ASHRAE defines the x-factor for a certain temperature as a relative failure rate of IT equipment at this temperature, compared with the failure rate at the baseline temperature of 20°C. Fig. 11(a) shows the fitted function of the x-factors for the temperature from 15°C to 45°C. Let X_i and AFR_i denote the x-factor and AFR at a temperature of T_i, respectively. Since the x-factor X_i is the relative failure rate of IT equipment at T_i, compared with the failure rate at temperature of 20°C, we have $\mathrm{AFR}_i = X_i \times \mathrm{AFR}_{20}$. Given a measured AFR_m at a temperature denoted by T_m, the AFR_{20} can be calculated as $\mathrm{AFR}_{20} = \mathrm{AFR}_m / X_m$. Then, the AFR_i at a temperature of T_i can be calculated as follows: $\mathrm{AFR}_i = X_i \times \frac{\mathrm{AFR}_m}{X_m}$. From our industry partner, the historical AFR of the tested servers in our testbed is 2% at the temperature of 24°C, i.e., $\mathrm{AFR}_m = 2\%$. Fig. 11(b) presents the AFRs for the temperatures in a range of [15°C, 45°C], which are predicted using our developed approach based on the $\mathrm{AFR}_m = 2\%$ and x-factors. We can see that the predicted AFR_{29} for the temperature of 29°C tested in our reliability test is 2.3%. This result indicates that increasing the temperature setpoint from 24°C to 29°C leads to about 0.3% AFR increase only.

4.2.3 *Impact of number of servers on accuracy of measured server failure rate*

In this section, we study impacts of the number of tested servers on the accuracy of the measured server failure rate at 29°C in our reliability test. Let λ and P denote the server failure probability and mean time between failures (MTBF) at 29°C, respectively. Moreover, we define FR_Υ as the server failure rate (i.e., ratio of the number of failed servers to the total number of tested servers) over a certain test period denoted by Υ. Then, the relationship between the P, λ, and FR_Υ can be represented as [31]:

$$P = \frac{1}{\lambda} = \frac{\Upsilon}{\ln \frac{1}{1-\mathrm{FR}_\Upsilon}}. \tag{1}$$

In a simulation, given a certain number of servers, denoted by N, we draw N failure time samples from the exponential distribution of $\mathrm{Exp}(\lambda\tau)$ where τ is the failure time. Then, we consider the ratio of the total number of samples with the failure time less than Υ to the number of servers N as the simulated FR_Υ at 29°C.

So far, we have run our reliability test for 5 months, i.e., $\Upsilon = 5$ months. Given the predicted $\mathrm{AFR}_{29} = 2.3\%$ (i.e., the failure rate over 12 months), the failure rate over $\Upsilon = 5$ months can be calculated as $\mathrm{FR}_\Upsilon = \frac{5\mathrm{AFR}_{29}}{12} \approx 0.96\%$. Fig. 11(c) shows the distributions of simulated FR_Υ values with a mean value of 0.96% over 10^6 repeated simulations

(a) Corrosion coupons. (b) Real-time corrosion sensor.

Fig. 12. Deployment of corrosion coupons and sensor in the testbed.

TABLE 4
Our corrosion rate measurement and ISA's classification.

Copper Corrosion Thickness over 30 Days				
Cu_2O	CuO	Cu_2S	Total	ISA* Class
97Å	9Å	19Å	**125Å**	G1 Mild (<300Å)
Sliver Corrosion Thickness over 30 Days				
AgCl	Ag_2S	Unknown	Total	ISA Class
10Å	38Å	0	**48Å**	G1 Mild (<200Å)

*Instrument Society of Automation (ISA) standard.

when the number of servers N are 512 and 3500. We can see that the FR_Υ values concentrate at 0.96%. Fig. 11(d) presents the standard deviation (Stdev) of the simulated FR_Υ values, denoted by σ when the number of servers varies from 100 to 3500 servers. The σ decreases with the number of servers. Specifically, with $N = 521$ servers in our testbed, the σ of the FR_Υ is 0.0043, i.e., about 45% of $FR_\Upsilon = 0.96\%$. From the 3σ rule-of-thumb, the FR_Υ in our testbed will be within [0%, 2.25%] with a high probability. Meanwhile, in our reliability test, we observed no server failure over the completed test period of $\Upsilon = 5$ months, i.e., the measured $FR_\Upsilon = 0$ lies in [0%, 2.25%]. This result implies the ground-truth FR_Υ is lower than its predicted value of 0.96%. Similarly, the ground-truth AFR_{29} is not higher than 2.3%.

4.2.4 Corrosion rate measurement

During the reliability test, we deployed corrosion coupons and real-time corrosion sensor to measure the corrosion rate in the testbed as shown in Fig. 12. Specifically, a set of corrosion coupons consisting of the copper and sliver metal trips was deployed on the top of an IT rack for a period of 30 days. Then, the coupons were sent to the corrosion sensor vendor for the lab analysis to determine the corrosion rate inside the testbed's space. Table 4 shows the vendor's analysis results. From Table 4, we can see that the corrosive films on the copper and sliver strips have a total thickness of 125 Å and 28 Å, respectively. Since the exposure period of the coupons to the testbed's environment is 30 days, the measured thickness results are equivalent to the copper and sliver corrosion rates of 125 Å/month and 28 Å/month, respectively. These results indicate that the testbed's environment condition can be classified as the corrosivity category of G1 Mild which requires the copper and sliver corrosion rates to be lower than 300 Å/month and 200 Å/month, respectively, according to the Instrument Society of Automation (ISA) standard ANSI/ISA-71.04-2013 [32]. As shown in Fig. 12 (b), the real-time corrosion sensor also classifies the testbed's environment as the G1 Mild category. From the vendor's analysis report, the G1 Mild environment may have the gas concentrations of H_2S<3ppb, SO_2<10ppb, NO_x<50ppb, Cl_2<1ppb, and NH_3<500ppb. Under the G1 Mild environment condition, the corrosion is not a factor in determining the equipment reliability. In other words, the environment condition of our testbed is well controlled and the corrosion is not a reason of the server's hardware failures. This controlled environment allows us to study the sole impact of the air temperature and RH conditions on the reliability of tested servers, which is a primary goal of our long-term reliability test.

4.3 Implications, Limitations, & Future Works

In this section, we summarize the key implications drawn from the above experimental measurement results. Then, we clarify the limitations of our work and discuss the potential future works to address these limitations.

4.3.1 Implications

As an unique opportunity that has full access and control of a real DC in the tropics, our work has generated useful information that the DC operators can use to improve the energy efficiency of their air-cooled DCs. Specifically, our results suggest that the air-cooled DCs operated in enclosed buildings can increase the supply air temperature setpoints up to 29°C for improving their DC energy efficiency with little impact on the server's reliability. By maintaining the supply air temperature setpoint greater than 23°C, the supply air RH can be maintained below the ASHRAE's recommended RH limit of 70% without the need of RH control. Although we only investigate the impacts of increasing the temperature setpoint on the DC energy efficiency and IT equipment reliability in a tropical area (i.e., Singapore), our experimental results and observations can be also applied to the DCs in other locations. Specifically, the outside ambient air conditions of the DC location do not affect the IT power usage and equipment reliability results because the air used for cooling is recirculated within the enclosed building without mixing with the outside air. Moreover, as shown in §4.1.4, the DC cooling power usage increases with the outside air temperature. Thus, our work presents the lower bound of the cooling energy savings that are obtained by raising the temperature setpoint in other cooler locations.

4.3.2 Limitations and Future Works

In this work, we only focus on investigating the energy efficiency of a DX cooling system which is often implemented in the small-scale DCs. Meanwhile, the large-scale DCs often adopt a CW cooling system. However, constructing a real DC testbed with a CW system requires much higher equipment costs. As future work, it is interesting to extend the measurement results of this work to study the energy efficiency of a CW-cooled DC. For instance, our collected real IT power and weather data (i.e., outside air temperature and RH) traces can be integrated into a DC simulator (e.g., EnergyPlus [33]) embedded with a numerical CW cooling model [34] to build a virtual DC testbed for evaluating the

TABLE 5
DC Operational Levels.

Level$^+$	Annualized TCO (ATCO) Range
Level 1	$(\text{ATCO}_{\max} - 0.25\Delta^*, \text{ATCO}_{\max}]$
Level 2	$(\text{ATCO}_{\max} - 0.5\Delta, \text{ATCO}_{\max} - 0.25\Delta]$
Level 3	$(\text{ATCO}_{\max} - 0.75\Delta, \text{ATCO}_{\max} - 0.5\Delta]$
Level 4	$[\text{ATCO}_{\min}, \text{ATCO}_{\max} - 0.75\Delta]$

$^+$ The cost-effective level increases with Level 1 to 4.
$^*\Delta = \text{ATCO}_{\max} - \text{ATCO}_{\min}$.

CW cooling energy usage. Moreover, the public weather data traces [35], [36] can be integrated into this DC simulator for investigating the energy efficiency of the air-cooled DCs in other locations worldwide. Similarly, the public server workload datasets [14], [15] can be also used to study the large-scale DCs with more high-performance servers. In addition, our proposed server reliability prediction approach only considers impact of the temperatures on the server failure rate. The current research has shown that the age of the IT equipment may also affect its reliability. Thus, our future work also plans to consider the impact of the server age to improve the accuracy of our prediction approach.

5 TCO Analysis Framework

This section presents a total cost of ownership (TCO) analysis framework to calculate DC operating costs under specific temperature and RH setpoints. Then, the proposed framework is applied to perform the TCO analysis for the testbed as a case study.

5.1 Approach Overview

Our experimental results showed that increasing the supply air temperature and RH setpoints can help improve the DC energy efficiency. However, the DC operators and their tenants may have different mentalities in accepting high temperatures and RH levels because their servers may require different supply air conditions and costs for repairing/replacing the failed components. Thus, we propose a TCO analysis framework that aims to quantify all DC operating costs under specific ranges of the temperature and RH setpoints. Then, the DC operator defines multiple cost-effective levels, each of which has certain ranges of the temperature and RH setpoints and operating cost. These levels help facilitate the DC system implementations and the communications among DC operators and tenants. For instance, in the colocation DCs that rent physical spaces out to multiple tenants for housing their own IT equipment, the operator can divide its DC space into multiple operating zones with a certain setting of the temperature and RH setpoints and renting price. Then, the operator introduces these operating zones to its tenants. Each tenant can select one zone based on the air working condition requirement of its IT equipment. The tenant with the IT equipment that can work with the high temperature and RH levels can choose the high-cost-effective level with a cheaper renting price.

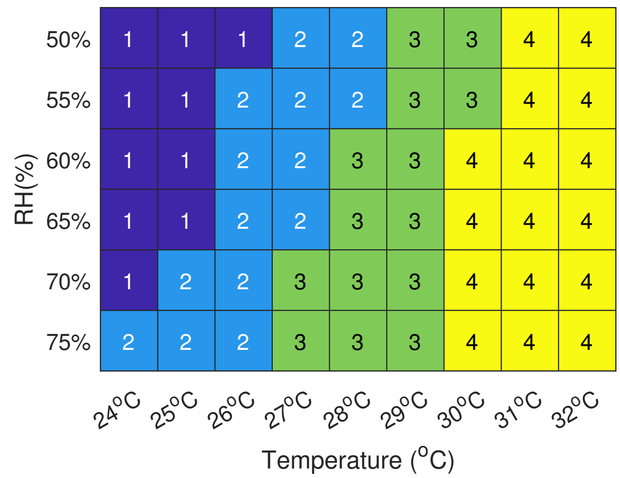

Fig. 13. Four cost-effective levels of the testbed. The number in a cell indicates the level index.

5.2 TCO Analysis

The proposed TCO analysis framework consists of the following three main steps. The first step is to evaluate the total DC power usage including the IT and cooling power usages under specific supply air temperature and RH ranges, denoted by $[T_{\min}, T_{\max}]$ and $[RH_{\min}, RH_{\max}]$, respectively. To obtain the accurate results, this energy usage evaluation step can be performed by the real measurements like the experiments conducted in our DC testbed. However, this measurement method requires high costs for constructing and operating the real-world DC testbeds. Alternatively, the DC operator can perform simulations using the existing DC energy usage and server reliability models/tools [1]. For the cooling power usage, the cooling equipment vendors often have the simulation tools to study the energy profile of their cooling system products under various working conditions. The DC operators can cooperate with their vendor to obtain the energy usage of their cooling system under various temperature and RH settings using these tools. Note that in this step, the total power usage is evaluated for a DC with certain system design and configuration. Thus, the evaluated power usage results also include the impact of the DC system design parameters (e.g., the server room size) on the energy usage. For instance, a larger room may lead to the higher DC power usage because more servers can be deployed, thereby increasing the IT power and requiring a higher cooling capacity (i.e., cooling power) to maintain the same temperature setpoint.

The second step is to calculate the total TCO cost of the DC under each combination setting of the temperature and RH setpoints in $[T_{\min}, T_{\max}]$ and $[RH_{\min}, RH_{\max}]$, respectively. Specifically, under a combination setting i, the TCO cost, denoted by ATCO_i, is the sum of all annualized operating expenses (OpEx) costs. Specifically, $\text{ATCO}_i = E_i + C_i + M_i$, where the E_i, C_i, and M_i are the annualized total energy usage, carbon and IT equipment maintenance costs at the setting i's temperature and RH setpoints, respectively. The ATCO_i does not include the capital expenditures (CapEx) costs because these costs in general do not change with the temperature and RH levels. In particular, the E_i is a total of the IT and cooling energy usages per year multiplying with the electricity price of the region where the DC is located. The carbon cost C_i is calculated by multiplying the per-year total energy usage with the CO_2 intensity emission and country's carbon tax rates. The CO_2 intensity emission rate is the amount of CO_2 released to produce a kilowatt hour (kWh) of electricity [37]. It depends on the type of energy sources used to generate the electricity. The carbon tax rate is often charged based on the total amount of CO_2 emissions. Moreover, the IT equipment maintenance cost M_i is the total cost required to replace/repair the failed servers per year. It means $M_i = \text{AFR}_i \times N \times c_r$, where AFR_i is the AFR at setting i, N is the total number of servers in the DC, and c_r is the per-sever replacement/maintenance cost.

In the last step, we divide the evaluated temperature and RH setpoints into the four operating levels based on their annualized TCO cost ATCO_i. Let $\text{ATCO}_{\min}$ and $\text{ATCO}_{\max}$ denote the minimum and maximum values among all ATCO_i values of all combined temperature and RH settings obtained in the second step. Then, the ATCO range for each level is determined using equations defined in Table 5. Specifically, the level's maximum ATCO values evenly divide the ATCO range of [$\text{ATCO}_{\min}$, $\text{ATCO}_{\max}$] with a step size of 25% of $\text{ATCO}_{\max} - \text{ATCO}_{\min}$. Then, the temperature and RH setpoints of a setting i belong to Level j if the ATCO_i falls within the annualized ATCO range of Level j. Such four discrete levels group the temperature and RH settings which has similar ATCO values into the same group. From Table 5, Level 1 is the least cost-effective level, while Level 4 is the most cost-effective level.

5.3 Cost-Effective Operating Levels of Testbed

This section presents a case study to illustrate how the proposed framework can be used to perform the TCO analysis of our real DC testbed. Specifically, we use the measured total power usages shown in Fig. 4(b) at the tested temperature and RH setpoints to calculate the annualized energy usage cost E_i with the electricity price of USD0.24 per kWh (i.e., the electricity price rate in our region). Moreover, we use the CO_2 emission rate of $0.731\,\text{kg/kWh}$ and carbon tax rate of USD3.7 per tone of CO_2 emission (tCO2e) in our regions to calculate the carbon cost C_i. The predicted AFRs shown in Fig. 11(a) and the per-sever replacement cost $c_r = \text{USD}1500$ are used to calculate the IT equipment maintenance cost M_i for our testbed with a total number of servers $N = 521$. Fig. 13 shows the four operating levels of our testbed with specific temperature and RH setpoints in ranges of [24°C, 32°C] and [50%, 75%], respectively.

In the DCs, the different servers or tenants may require distinct maximum temperature and RH thresholds. The DC operator can configure its space into up to four operating zones corresponding to the determined four cost-effective levels. From Fig. 13, we can see that each level may include a set of different temperature and RH setpoints. Thus, each operating zone should be configured with the temperature and RH setpoints that lead to the lowest ATCO value among the level's temperature and RH setpoints. Then, each server should be deployed into the highest cost-effective level from 1 to 4 which has the temperature and RH setpoints below its required thresholds. As a result, the operator can satisfy the distinct requirements of multiple servers with the minimum operating cost.

5.4 Computing Performance Degradation Costs

Our proposed TCO analysis framework does not consider costs incurred by the server's computing performance drop due to the increased temperature and RH setpoints. The reason is that as reviewed in §2.2, the existing studies consistently show that temperature and RH has no/little impact on the server's computing performance. Specifically, in our previous work [19], we conducted with a set of experiments to study how the supply air temperature and RH affects the computing performance of three main server's components, i.e., the CPU, HDD and memory. Our experimental result analysis showed that the main server components can operate without performance degradation under combined impact of the temperature up to 37°C and the RH above 90%. More details on our experiment design and results can found in [38].

6 Conclusion

We describe the design and configuration of an air-cooled DC testbed in a tropical area. Then, we present the key results and observations obtained from the experiments conducted on the testbed, including the power usages of the IT and cooling systems, and server reliability during the experiments. We also propose a TCO analysis framework which consider all relevant operating costs for setting the supply air temperature and RH in the tropical DCs. The observations and TCO analysis framework presented in this paper will be useful to future efforts of building and operating air-cooled DCs in the tropics.

Acknowledgments

This project is supported by the National Research Foundation, Singapore, funded under Energy Research Test-Bed and Industry Partnership Funding Initiative, part of the Energy Grid (EG) 2.0 programme. Part of the objective of this project is to support the development of the Singapore standard (SS 697) [12]. Rui Tan's work is also supported by the Ministry of Education, Singapore, under its Academic Research Fund Tier 1 (RG88/22).

References

[1] M. Dayarathna, Y. Wen, and R. Fan, "Data center energy consumption modeling: A survey," *IEEE Communications Surveys & Tutorials*, vol. 18, no. 1, pp. 732–794, 2015.
[2] M. K. Patterson, "The effect of data center temperature on energy efficiency," in *Proc. 11th Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems*, 2008, pp. 1167–1174.
[3] D. Moss and J. H. Bean, "Energy impact of increased server inlet temperature," *APC white paper*, vol. 138, 2009.
[4] "Google's green data centers: Network pop case study," https://shorturl.at/vBD24, 2011.
[5] 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," *ACM SIGMETRICS Performance Evaluation Review*, vol. 40, no. 1, pp. 163–174, 2012.
[6] S. Sankar, M. Shaw, K. Vaid, and S. Gurumurthi, "Datacenter scale evaluation of the impact of temperature on hard disk drive failures," *ACM Transactions on Storage*, vol. 9, no. 2, p. 6, 2013.
[7] I. Manousakis, S. Sankar, G. McKnight, T. D. Nguyen, and R. Bianchini, "Environmental conditions and disk reliability in free-cooled datacenters," in *Proc. 14th USENIX Conference on File and Storage Technologies (FAST)*, Feb. 2016, pp. 53–65.

[8] P. Arroba García, M. Zapater Sancho, J. L. Ayala, J. M. Moya Fernández, K. Olcoz, and R. Hermida, "On the leakage-power modeling for optimal server operation," 2013.
[9] I. D. A. of Singapore (IDA), "Singapore data center energy consumption study," 2013.
[10] A. Shehabi and et al., "United states data center energy usage report," Lawrence Berkeley National Lab (LBNL), Berkeley, CA, Tech. Rep., 2016.
[11] ASHRAE, "Thermal guidelines for data processing environments," Fifth Edition, 2021.
[12] "Deployment and Operation of Data Centre IT Equipment under Tropical Climate," Singapore Standard 697:2023. [Online]. Available: https://www.singaporestandardseshop.sg/Product/SSPdtDetail/f64b2a31-57a0-49ec-8716-57ccbdc33365
[13] D. V. Le, "An air-cooled tropical data center (TDC2.0) dataset," DR-NTU, https://shorturl.at/bnFLR.
[14] "Alibaba cluster trace program," https://bit.ly/3wvFQBN.
[15] "Google cluster workload traces," https://bit.ly/3vh9YR2.
[16] K. Haghshenas, B. Setz, Y. Blosch, and M. Aiello, "Enough hot air: the role of immersion cooling," *Energy Informatics*, vol. 6, no. 1, p. 14, 2023.
[17] "Uptime institute data center cooling systems survey 2021," *Uptime Institute*, 2021.
[18] A. Habibi Khalaj and S. K. Halgamuge, "A review on efficient thermal management of air- and liquid-cooled data centers: From chip to the cooling system," *Applied Energy*, vol. 205, pp. 1165 – 1188, 2017.
[19] D. Van Le, Y. Liu, R. Wang, R. Tan, and L. H. Ngoh, "Air free-cooled tropical data center: Design, evaluation, and learned lessons," *IEEE Transactions on Sustainable Computing*, vol. 7, no. 3, pp. 579–594, 2022.
[20] R. Schmidt, "A history of ashrae technical committee TC9.9 and its impact on data center design and operation," *Electronics Cooling*, 2013. [Online]. Available: https://bit.ly/3MA6sNi
[21] Z. Wang, C. Bash, N. Tolia, M. Marwah, X. Zhu, and P. Ranganathan, "Optimal fan speed control for thermal management of servers," in *Proc. International Electronic Packaging Technical Conference and Exhibition*, vol. 43604, 2009, pp. 709–719.
[22] Microsoft and Intel, "Server power and performance evaluation in high-temperature environments," *White Paper*, 2012.
[23] P. Singh, L. Klein, D. Agonafer, J. Shah, and K. Pujara, "Effect of relative humidity, temperature and gaseous and particulate contaminations on information technology equipment reliability," in *Proc. International Technical Conference and Exhibition on Packaging and Integration of Electronic and Photonic Microsystems (InterPACK2015)*, 2015, pp. 1–9.
[24] J.-E. Svensson and L.-G. Johansson, "A laboratory study of the effect of ozone, nitrogen dioxide, and sulfur dioxide on the atmospheric corrosion of zinc," *Journal of the Electrochemical Society*, vol. 140, no. 8, pp. 2210–2216, 1993.
[25] J. Liu and A. Terzis, "Sensing data centres for energy efficiency," *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, vol. 370, no. 1958, pp. 136–157, 2012.
[26] A. Saifullah, S. Sankar, J. Liu, C. Lu, R. Chandra, and B. Priyantha, "Capnet: Exploiting wireless sensor networks for data center power capping," *ACM Transactions on Sensor Networks (TOSN)*, vol. 15, no. 1, pp. 1–34, 2018.
[27] M. Jafarizadeh and R. Zheng, "Optimal design of lemonet for environmental monitoring of data centers," *IEEE Transactions on Green Communications and Networking*, vol. 5, no. 4, pp. 1820–1832, 2021.
[28] P. Pinmas, P. Panyadee, M. Kaewmoracharoen, S. Khueankhan, and A. Wannachai, "Towards digital twin data center using building information modeling and real-time data sensing," in *Proc. 37th International Technical Conference on Circuits/Systems, Computers and Communications (ITC-CSCC)*, 2022, pp. 908–911.
[29] "Direct expansion air conditioning systems." [Online]. Available: https://bit.ly/419hkW1
[30] ASHRAE, "Gaseous and particulate contamination guidelines for data centers," *White Paper*, 2011.
[31] W. Li, Y. Yang, and D. Yuan, "Literature review," in *Reliability Assurance of Big Data in the Cloud*, 2015, pp. 9–17.
[32] ISA-71.04-2013, "Environmental conditions process measurement and control systems: Airborne contaminants," *ANSI/ISA - The Instrumentation, System and Automation Society*, 2013.
[33] "Energyplus," https://energyplus.net.
[34] R. Wang, D. Van Le, R. Tan, and Y.-W. Wong, "Real-time cooling power attribution for co-located data center rooms with distinct temperatures and humidities," *ACM Transactions on Cyber-Physical Systems (TCPS)*, vol. 6, no. 1, pp. 1–28, 2022.
[35] "Repository of free climate data for building performance simulation," https://climate.onebuilding.org.
[36] "The MIT supercloud dataset," https://registry.opendata.aws/dcc.
[37] "What is carbon intensity?" [Online]. Available: https://ngrid.com/3HDEyff
[38] D. V. Le, Y. Liu, R. Wang, and R. Tan, "Tropical data centre proof-of-concept," Nanyang Technological University, Tech. Rep., 2019.

Duc Van Le received the BEng degree in electronics and telecommunications engineering from Le Quy Don Technical University, Vietnam, in 2011, and the PhD degree in computer engineering from University of Ulsan, South Korea, in 2016. Currently, he is a Senior Research Fellow at School of Computer and Engineering, Nanyang Technological University, Singapore. Previously, he was a Research Fellow (2016-2018) at Department of Computer Science, National University of Singapore. His research interests include sensor networks, internet of things, and cyber-physical systems. He is the recipient of Best Paper Award from ACM/IEEE ICCPS'23. He is a Senior Member of IEEE.

Jing Zhou received a PhD degree in thermal engineering from the Huazhong University of Science and Technology, China, in 2021. Currently, he is a research fellow at School of Mechanical and Aerospace Engineering, Nanyang Technological University (NTU), Singapore. His research interests include data center cooling, thermodynamic analysis, and system integration and optimization.

Rongrong Wang received the BEng degree in control science and engineering from Northwestern Polytechnical University, China, in 2017, and the MS degree in 2018 in electrical and electronic engineering from Nanyang Technological Univer- sity (NTU), Singapore, where she is currently working toward the PhD degree at School of Computer Science and Engineering (SCSE). She is also working as a research associate at SCSE, NTU. Her research interests include advanced sensing and smart CPS environment.

Rui Tan is an Associate Professor at School of Computer Science and Engineering, Nanyang Technological University, Singapore. Previously, he was a Research Scientist (2012-2015) and a Senior Research Scientist (2015) at Advanced Digital Sciences Center of University of Illinois at Urbana-Champaign, and a postdoctoral Research Associate (2010-2012) at Michigan State University. He received the Ph.D. (2010) degree in computer science from City University of Hong Kong, the B.S. (2004) and M.S. (2007) degrees from Shanghai Jiao Tong University. His research interests include cyber-physical systems and Internet of things. He is the recipient of Best Paper Award, Best Paper Award Runner-Up/Finalist from ICCPS 2022 and 2023, SenSys 2021 and 2022, IPSN 2014 and 2017, and PerCom 2013. He served as Associate Editor of ACM Transactions on Sensor Networks, TPC Co-Chair of e-Energy'23, EWSN'24, SenSys'24, and General Co-Chair of e-Energy'24. He received the Distinguished TPC Member recognition thrice from INFOCOM in 2017, 2020, and 2022. He is a Senior Member of IEEE.

Fei Duan is an associate professor at School of Mechanical and Aerospace Engineering, Nanyang technological University (NTU), Singapore. He graduated with a Ph.D. degree from University of Toronto, Canada. The topics of his research cover droplet wetting and evaporation dynamics, particle self-assembly, enhanced energy storage and transfer, date center thermal management, efficient cogeneration system, etc. So far, he has advised over 25 postdocs or research associates, 19 Ph.D. and 14 master's students. He has published over 170 journal papers, more than 120 conference papers, 5 book chapters, etc. He serves as a Subject Editor for Applied Thermal Engineering (Elsevier), and at Editorial Board of Scientific Reports (Nature Portfolio), and Frontiers in Heat and Mass Transfer (Tech. Science Press).