

# Real-Time Cooling Power Attribution for Co-Located Data Center Rooms with Distinct Temperatures

Rongrong Wang  
Computer Science & Engineering  
Nanyang Technological University  
Singapore

Duc Van Le  
Computer Science & Engineering  
Nanyang Technological University  
Singapore

Rui Tan  
Computer Science & Engineering  
Nanyang Technological University  
Singapore

Yew-Wah Wong  
Energy Research Institute  
Nanyang Technological University  
Singapore

Yonggang Wen  
Computer Science & Engineering  
Nanyang Technological University  
Singapore

## ABSTRACT

At present, a co-location data center often applies an identical and low temperature setpoint for its all server rooms. Although increasing the temperature setpoint is a rule-of-thumb approach to reducing the cooling energy usage, the tenants may have different mentalities and technical constraints in accepting higher temperature setpoints. Thus, supporting distinct temperature setpoints is desirable for a co-location data center in pursuing higher energy efficiency. This support calls for a new cooling power attribution scheme to address the inter-room heat transfers that can be up to 9% of server load as shown in our real experiments. This paper describes our approaches to estimating the inter-room heat transfers, using the estimates to rectify the metered power usages of the rooms' air handling units, and fairly attributing the power usage of the shared cooling infrastructure (i.e., chiller and cooling tower) to server rooms by following the Shapley value principle. Extensive numeric experiments based on a widely accepted cooling system model are conducted to evaluate the effectiveness of the proposed cooling power attribution scheme.

## CCS CONCEPTS

• **Hardware** → **Impact on the environment; Enterprise level and data centers power issues; Power and energy.**

## KEYWORDS

Co-located data centers, cooling power attribution, heat transfer

### ACM Reference Format:

Rongrong Wang, Duc Van Le, Rui Tan, Yew-Wah Wong, and Yonggang Wen. 2020. Real-Time Cooling Power Attribution for Co-Located Data Center Rooms with Distinct Temperatures. In *The 7th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation (BuildSys '20)*, November 18–20, 2020, Virtual Event, Japan. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3408308.3427607>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

*BuildSys '20*, November 18–20, 2020, Virtual Event, Japan

© 2020 Association for Computing Machinery.

ACM ISBN 978-1-4503-8061-4/20/11...\$15.00

<https://doi.org/10.1145/3408308.3427607>

## 1 INTRODUCTION

The advance of Internet of Things (IoT) and the deployments of 5G networks place proliferated demands on data processing, storage, and analytics in the back end. This calls for ever growing capacity of data centers (DCs) as the back-end computing infrastructure. Since DC construction and operations require extensive expertise and significant investment, various computing users rent server rooms from co-location DC operators to host their information technology (IT) equipment. In a co-location DC, the centralized operations and management of shared facilities including the cooling systems reduce the operating expenses of the tenants. As such, co-location DC has become a major form in the DC industry. By 2019, there are more than 4,000 co-location DCs operating worldwide [1].

DCs consume lots of energy and use a big portion for cooling. In 2014, the electricity used by DCs in U.S. accounted for 1.8% of the country's electricity consumption [20]. In tropics such as Singapore, this ratio was up to 7% [3]. On average, about 40% DC energy is used for cooling [20]. Thus, reducing cooling energy is an important mission. Currently, a co-location DC often applies a low temperature setpoint (e.g., 21°C) for all server rooms. Research has shown that increasing the server room temperature setpoint is a rule-of-thumb approach to reducing cooling energy [15, 17]. Specifically, a 1°C increment can lead to about 4% cooling energy saving [5]. To prompt higher temperature setpoints, the American Society of Heating, Refrigeration and Air-Conditioning Engineers (ASHRAE) has been working on extending the recommended allowable temperature range of IT equipment [18]. For instance, the servers compliant with the A3 requirement [4] can operate continuously with inlet temperature up to 40°C. Many latest servers (e.g., Dell's gen14 and HPE's DLx gen9 servers) are A3-compliant.

However, the tenants in a co-location DC may have different mentalities and technical constraints in accepting higher temperature setpoints. For instance, a tenant running a server room with cold aisle air containment can easily accept higher room temperature setpoints if the supply air is still cold. In the absence of air containment, the tenants running IT equipment that requires low temperatures can hardly accept higher room temperature setpoints. As such, supporting distinct temperature setpoints for server rooms is desirable for a co-location DC in pursuing higher energy efficiency. To this end, two important new issues need to be studied to make sense the encouragement of adopting distinctly higher

temperature setpoints subject to the tenants' technical constraints. First, as these rooms share parts of the co-location DC's cooling infrastructure (i.e., chiller and cooling tower), how to fairly attribute the power usage of the shared cooling infrastructure to the rooms adopting distinct temperature setpoints? We need to re-examine the applicability of the prevailing attribution policies (e.g., proportional splitting based on rooms' IT loads). Second, as the server rooms of distinct temperature setpoints may have inter-room heat transfers, the impact of the heat transfers on the attribution of the cooling power to the rooms needs to be understood and addressed properly.

This paper investigates the above two issues based on a two-stage cooling system model that captures the essence of the cooling system designs in co-location DCs. The first stage consists of the air handling units (AHUs) in the individual server rooms. An AHU transfers the heat carried by the return hot air from the IT equipment to the influx cold water. It controls the water flow rate via a valve and air flow rate via internal fans to maintain the return air temperature at the setpoint. Thus, the AHU power usage is mainly due to its internal fans. The second stage is a cooling infrastructure shared by all server rooms, which consists of a *chiller* and a *cooling tower*. The chiller uses a refrigeration cycle to transfer the heat from the AHUs to a second water cycle with a higher temperature. The cooling tower further dissipates the heat carried by the second water cycle to the ambient air. The power usage by the second stage is mainly due to the water pumps in the two water cycles, the compressors in the chiller plant, and the fans in the cooling tower.

Based on the above model, our analysis for distinct room temperatures gives the following two properties. First, the second-stage cooling power (i.e., the shared part) only depends on the total IT load of all rooms. Second, the inter-room heat transfers affect individual AHU powers, and do not affect the second-stage cooling power. Based on the above two properties, we propose a cooling power attribution scheme that computes the power share of each room by two components. The first component is the sum of the metered power usage of the considered room's AHU and a *rectification* that addresses the AHU's extra power usage due to the heat transfers with the neighbor rooms. The second component is a fair share of the second-stage cooling power based on all rooms' IT loads. This paper aims to achieve real-time cooling power attribution based on the readings of the relevant meters in real time (e.g., every minute). The fine time granularity of the power attribution improves the accuracy of energy accounting and charging. However, the real-time computations of the above two components face respective challenges as discussed below.

Significant heat can dissipate from a high-temperature room to a low-temperature room via the separation (e.g., walls and floors). From our experiment conducted in a real server room, when its temperature is 12°C higher than the building ambient, the heat dissipated through its enclosure is 9% of its IT load. Due to the inter-room heat transfers, the AHU of a room that receives net influx heat transfer uses more power to maintain the temperature. This is because the AHU's internal fans need to rotate faster to transfer more heat from the return air to the chilled water. Thus, the currently prevailing AHU cooling power charging scheme merely based on the metered AHU power usage will be biased. A rectification is needed to address the increment/decrement of the AHU

power caused by the heat transfer with each neighbor room. However, under a general setting of  $n$  rooms, estimating  $\binom{n}{2}$  inter-room heat transfers from  $n$  equations each formulating a room's net influx/outflow heat transfer based on its measured IT load and heat removed by its AHU is an *underdetermined* problem. To address this challenge, we exploit the first principle that the heat transfer is proportional to the temperature difference and then integrate sufficient historical measurements with varied temperatures of the rooms into an *overdetermined* equation system with  $\binom{n}{2}$  unknown heat transfer coefficients. With the estimated coefficients, we can estimate the real-time inter-room heat transfer based on the temperature difference of any two rooms and use that for rectification.

Attributing the second-stage cooling power to the server rooms also faces challenges. From analysis, we cannot divide the power into portions, each determined by an individual room's IT load only. For this scenario, the principle of Shapley value [19] can be applied to achieve certain fairness axioms. However, although the Shapley value is a well accepted conceptual device, it incurs high compute overhead due to its complexity of  $O(n \cdot 2^n)$ . To reduce the compute overhead, we propose two approaches to approximate the Shapley power attribution function. The first uses a multilayer perceptron (MLP) trained using data generated by feeding the Shapley power attribution function with random IT loads of the rooms. As MLP inference is fast, the MLP-based power attribution can be executed in real time. In nature, this approach offloads the intensive Shapley value computation to the offline training data generation process. The second approach uses a heuristic algorithm to compute the power attribution with  $O(n)$  complexity. Evaluation shows that the MLP approximation achieves 2.4% mean relative error (MRE) but requires lengthy training data generation, whereas the heuristic algorithm is lightweight and scalable, but gives higher MRE of 6%.

The solutions introduced above form a real-time cooling power attribution scheme for co-location DC adopting distinct room temperature setpoints. It can be used to encourage the tenants to increase their temperature setpoints without causing controversies. With our proposed scheme implemented, a co-location DC can take a more advantageous position on the market since its tenants keen to hotter server rooms can enjoy lower cooling costs in return.

The contributions of this paper are summarized as follows.

- Based on a representative model of co-location DC cooling systems, we analyze the impacts of the server rooms' distinct temperatures and IT loads on the cooling power usage, while considering heat transfers among server rooms.
- We design an approach to estimating inter-room heat transfer coefficients. Based on that, we propose an AHU power usage rectification approach aiming at eliminating the impact of inter-room heat transfers.
- We approximate the Shapley function for the second-stage cooling power using an MLP or a heuristic algorithm. Thus, the power attribution can be performed in real time due to the low compute overhead of the approximations.

Paper organization: §2 reviews related work. §3 presents preliminaries. §4 studies the impacts of room temperatures and IT loads on cooling power. §5 overviews our scheme. §6 and §7 expatiate AHU power rectification and second-stage cooling power attribution, respectively. §8 presents evaluation results. §9 concludes this paper.

## 2 BACKGROUND AND RELATED WORK

There are two broad categories of cooling cost attribution policies used in co-location DCs. The first category charges fixed costs for cooling to tenants disregarding the actual usages. For instance, the primitive *equal division* policy equally distributes the cooling cost to all server rooms regardless of their IT loads. It is seldom adopted due to its clear bias [12]. The *space-based* policy [10] applies a fixed per-square-foot or per-rack-space rate to compute the cost for each server room based on the room area or rack space. The second category of policies charges based on the server rooms' actual power/energy usages. The *load-proportional division* (LPD) policy [10] attributes the instantaneous cooling power to the server rooms proportionally according to their instantaneous IT loads. This policy only requires the total cooling power usage and each room's IT load. Its simplicity and fairness at the first glance promote its wide adoption. However, it is also biased since the power usage of the cooling system is analytically indivisible with respect to the IT loads of individual server rooms (cf. §3.2). In addition, it does not consider the inter-room heat transfers, if the server rooms adopt distinct temperature setpoints. In the performance evaluation of this paper (cf. §8), LPD is employed as a baseline approach.

The issue of inter-room heat transfer has been considered in the context of centralized heat provision in multi-apartment residential buildings, because the residents may have different preferences on room temperatures. The studies [6, 25] focus on modeling the relationship between the amount of transferred heat and various affecting factors including the temperature difference and the properties of walls. The studies [14, 21] estimate the amount of transferred heat between adjacent apartments and then reallocate the heating costs calculated based on the metered heat delivered to the apartments. However, the approaches developed in these existing studies [6, 14, 21, 25] that concern heating provision cannot be readily applied to cooling provision in co-location DCs. For instance, as shown in this paper, the inter-room heat transfers affect the power usage of the first-stage cooling in the server rooms. The apartment heating systems do not have this feedback effect because the heat exchangers in the apartments do not use power. In addition, different from these existing studies that build models of heat transfer from detailed parameters such as building layout, apartments' 3D structures, and wall material properties, which are tedious processes, we apply data analytics to estimate inter-room heat transfer coefficients without resorting to detailed modeling.

A recent study [11] considers fair attribution of cooling cost to the server racks that reside in the same server room and belong to different tenants. Thus, these racks share the same room environment and are cooled by the same AHUs. The study [11] applies the Shapley value principle to attribute the room's cooling cost to the racks. Differently, we consider the more common scenario in which each room is used exclusively by a tenant. Thus, the in-room cooling cost attribution considered in [11] is not applicable. The second-stage cooling power attribution is not addressed in [11].

Energy apportionment, i.e., to estimate residents' energy footprints, has been studied in the context of commercial buildings [23, 24]. Its core problem is to associate residents' positions and activities with the building's real-time power usage, different from our problem of precisely attributing cooling power to server rooms.

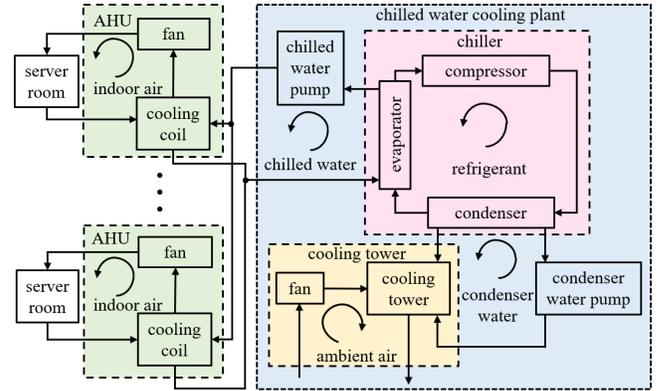


Figure 1: Two-stage cooling system model in co-location DCs

## 3 PRELIMINARIES

This section presents the analytical models characterizing the heat processes and cooling power usage. These models are used to study the impacts of temperature and IT load on cooling power in §4, and drive the evaluation in §8. This paper considers a typical two-stage cooling system adopted by co-location DCs, which is illustrated in Fig. 1. The first stage consists of the AHUs deployed in the server rooms. It transfers the heat carried by the return air from the servers to the water flowing in the cooling coil. The second stage is an infrastructure shared by all rooms. It consists of three cycles, i.e., chilled water cycle, refrigerant cycle, and condenser cycle. Lastly, it dissipates the heat through the cooling tower to the ambient.

The supply cold air from the AHU is often conducted to the server inlets. To improve cooling efficiency, many DCs build air containment for the cold aisle. The return hot air from the servers is in general conducted to the server room ambient. This paper aims to support distinct setpoints for the server rooms' return air temperatures (i.e., their ambient temperatures). From [22], 60% DCs adopt return air temperature as the main condition of server room. In current DCs, the setpoints are often around 21°C.

In what follows, we analyze the heat processes and the cooling power. Table 1 summarizes the notation used in this paper.

Table 1: List of notation

Notation	Definition
$T_{as}, T_{ar}$	Supply and return air temperatures
$T_{chws}, T_{cws}$	Chilled and condenser water temperatures
$T_o$	Outdoor air wet-bulb temperature
$\dot{m}_{ai}, \dot{m}_{chw_i}$	Room $i$ 's air and chilled water mass flow rates
$\dot{m}_{chw}$	Chilled water mass flow rate of chiller
$\dot{Q}_{tr_{ij}}$	Heat transfer rate between room $i$ and $j$
$\dot{Q}_{r_i}, \dot{Q}_{tr_i}$	Room $i$ 's net heat and heat transfer rates
$\dot{Q}_{ch}, \dot{Q}_{ct}$	Heat rates of chiller and cooling tower
$P_{IT_i}, P_{AHU_i}$	Room $i$ 's IT load and AHU power
$P_{ch}, P_{ct}, P_{chp}, P_{cp}$	Powers of chiller, cooling tower, two pumps
$P_{cw}, P$	Chilled water and DC cooling power
$\alpha_{ij}$	Heat transfer coefficient between room $i$ and $j$
$c_a, c_w$	Heat capacities of air and water

### 3.1 Heat Process Model

Consider a co-location DC of  $n$  server rooms. Denote by  $T_{ar_i}$  the temperature setpoint of room  $i$ . If two neighboring rooms  $i$  and  $j$  have different temperature setpoints  $T_{ar_i}$  and  $T_{ar_j}$ , the static heat transfer rate from room  $i$  to  $j$ , denoted by  $\dot{Q}_{tr_{ij}}$ , is  $\dot{Q}_{tr_{ij}} = \alpha_{ij}(T_{ar_i} - T_{ar_j})$ , where  $\alpha_{ij}$  is the heat transfer coefficient that depends on the material property and the area of the shared separation between the two rooms. A negative  $\dot{Q}_{tr_{ij}}$  means that the heat is transferred from room  $j$  to room  $i$ . A DC building is often built to have good thermal insulation from the atmospheric ambient for cooling efficiency [2]. Thus, in this paper, we ignore the heat transfer between any server room and the ambient.

In a server room, the AHU removes heat generated by the IT equipment and transferred from the neighboring rooms. Denoting by  $M_i$  the set of room  $i$ 's neighboring rooms, the overall heat rate of room  $i$ , denoted by  $\dot{Q}_{r_i}$ , is  $\dot{Q}_{r_i} = P_{IT_i} + \dot{Q}_{tr_i} = P_{IT_i} + \sum_{k \in M_i} \dot{Q}_{tr_{ki}}$ , where  $P_{IT_i}$  is room  $i$ 's IT load and  $\dot{Q}_{tr_i}$  is the net heat transfer rate that room  $i$  takes from  $M_i$ . In this paper, we assume that all the electrical power used by the IT equipment is converted to heat. Thus, the IT load is identical to the room's heat generation rate.

Now, we model the heat exchange in the AHU. The chiller supplies chilled water with a temperature of  $T_{chws}$ . From the law of conservation of energy, the air and chilled water mass flow rates (denoted by  $\dot{m}_{a_i}$  and  $\dot{m}_{chw_i}$ ), the AHU's supply cold air temperature (denoted by  $T_{as}$ ), and the temperature of the chilled water leaving the AHU (denoted by  $T_{chwr}$ ) satisfy  $\dot{Q}_{r_i} = c_a \dot{m}_{a_i} (T_{ar_i} - T_{as}) = c_w \dot{m}_{chw_i} (T_{chwr} - T_{chws})$ , where  $c_a$  and  $c_w$  are the heat capacities of air and water, respectively. Note that the AHU controls its internal fans and cooling coil to maintain the return and supply air temperatures at  $T_{ar_i}$  and  $T_{as}$ . The resulted  $\dot{m}_{a_i}$  due to the fan control ensures the heat removal rate of  $\dot{Q}_{r_i}$ . An internal valve of the AHU controls the  $\dot{m}_{chw_i}$  to maintain the temperature of the chilled water leaving the AHU at  $T_{chwr}$ . As a result, the total mass flow rate of the chilled water is  $\dot{m}_{chw} = \sum_{i=1}^n \dot{m}_{chw_i}$ , which needs to be maintained by a pump in the chilled water cycle. In this paper, we follow [7] to set  $T_{chws} = 7^\circ\text{C}$ ,  $T_{as} = 17^\circ\text{C}$ , and  $T_{chwr} = 12^\circ\text{C}$ .

The chiller uses a compressor to lower the temperature of the returned water and then transfers the heat to the cooling tower by a condenser. Due to the operation of the compressor, the chiller consumes a power of  $P_{ch}$  and converts it to heat, as well as removes the heat at a rate of  $\dot{Q}_{ch} = \sum_{i=1}^n \dot{Q}_{r_i}$ . They compose the removed heat rate of the cooling tower, which is denoted by  $\dot{Q}_{ct}$ . Specifically,  $\dot{Q}_{ct} = \dot{Q}_{ch} + P_{ch} = \sum_{i=1}^n \dot{Q}_{r_i} + P_{ch} = c_w \dot{m}_{chw} (T_{cwr} - T_{cws})$ , where  $\dot{m}_{chw}$  is the mass flow rate of condensed water,  $T_{cwr}$  and  $T_{cws}$  are the temperature setpoints of the water entering and leaving the cooling tower. Note that these two temperature setpoints are implemented by the condenser and the cooling tower, respectively. In this paper, we set  $T_{cws} = 20^\circ\text{C}$  and  $T_{cwr} = 27^\circ\text{C}$ .

### 3.2 Cooling Power Model

Let  $P_{cw}$  and  $P_{AHU_i}$  denote the power usages of the chilled water system and room  $i$ 's AHU, respectively. The total power usage of the cooling system, denoted by  $P$ , is given by  $P = \sum_{i=1}^n P_{AHU_i} + P_{cw}$ . The  $P_{cw} = P_{ch} + P_{ct} + P_{chp} + P_{cp}$ , where  $P_{ch}$ ,  $P_{chp}$ ,  $P_{cp}$ ,  $P_{ct}$  are the power usages of the water chiller, chilled water pump, condensed water pump and cooling tower, respectively. The  $P_{AHU_i}$ ,

$P_{ch}$ ,  $P_{chp}$ ,  $P_{cp}$ , and  $P_{ct}$  can be modeled as  $P_{AHU_i} = f_1(\dot{m}_{a_i})$ ,  $P_{ch} = f_2(T_{chws}, T_{cws}, \dot{Q}_{ch})$ ,  $P_{chp} = f_3(\dot{m}_{chw})$ ,  $P_{cp} = f_4(\dot{m}_{cw})$ , and  $P_{ct} = f_5(T_{cws}, T_{cwr}, T_o, \dot{Q}_{ct})$ , respectively, where  $T_o$  is the ambient air temperature and the functions  $f_k$  ( $k = 1, \dots, 5$ ) specify the factors that affect the cooling devices' power usages. Their detailed forms depend on the device specifications. In general, they are non-linear. Note that the  $f_1$  models the power usage of the AHU's internal fans; the AHU's cooling coils do not consume power, because they just passively transfer heat from the return air to the chilled water.

From the above modeling, the total cooling power is affected by IT loads ( $P_{IT_i}$ ), return air temperature setpoints ( $T_{ar_i}$ ), and the ambient air temperature ( $T_o$ ). In this paper, we focus on investigating the impacts of IT loads and temperature setpoints on the total cooling power. Thus, we view  $T_o$  as a constant. In this paper, we set  $T_o = 16^\circ\text{C}$ . Therefore, the total cooling power can be modeled as a function of the rooms' IT loads and temperature setpoints. Specifically,  $P = F(P_{IT_1}, P_{IT_2}, \dots, P_{IT_n}, T_{ar_1}, T_{ar_2}, \dots, T_{ar_n})$ . In general, this function is non-linear and analytically indivisible. Say, it cannot be written as  $P = \sum_{i=1}^n F_i(P_{IT_i}, T_{ar_i})$  for straightforward power attribution. Thus, the power attribution is a non-trivial problem.

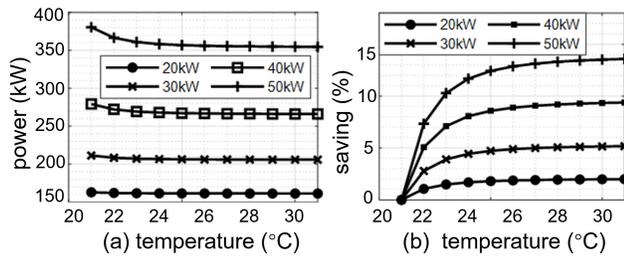
To apply the cooling power attribution scheme proposed in this paper, the co-location DC operator needs two power models: AHU power model  $f_1$  and the second-stage cooling system's composite power model  $f_2 + f_3 + f_4 + f_5$ . Such models are in general available to the operator that has detailed information of the facilities in its premise. In addition, the operator can also use historical data to build data-driven models. In the numeric experiments of this paper, we adopt the instantiated models in the ASHRAE's manual [7] for  $f_k$  ( $k = 1, \dots, 5$ ). Due to space limitation, we omit their details here.

## 4 IMPACTS OF TEMPERATURE AND IT LOAD ON HEAT PROCESSES

In this section, we conduct a set of numeric experiments to study the impacts of temperature setpoints and IT loads on the cooling power usage and heat transfer rate in a co-location DC. The results provide insights to guide the design of power attribution. In the numeric experiments, we vary the temperature setpoint from  $21^\circ\text{C}$  to  $33^\circ\text{C}$ . These temperatures are within the allowed ranges of ASHRAE A2, A3 and A4 equipments [4]. Note that most off-the-shelf data center IT devices meet A2 and A3 requirements.

### 4.1 Impact of Temperature on Cooling Power

We perform numeric experiments based on the two-stage cooling system model in §3 to study the cooling power savings achieved by raising the temperature setpoint. We consider a co-location DC consisting of ten server rooms, each of which has the same temperature setpoint and IT load. In the experiments, we vary the temperature setpoint of each room from  $21^\circ\text{C}$  to  $31^\circ\text{C}$  with a step size of  $1^\circ\text{C}$ . The room's IT load is varied from 20 kW to 50 kW with a step size of 10 kW. We define  $\sigma_t$  to be the relative cooling power saving achieved by raising the temperature setpoint from  $21^\circ\text{C}$  to a certain temperature of  $t^\circ\text{C}$ . Specifically,  $\sigma_t$  is calculated as  $\sigma_t = \frac{P_{21} - P_t}{P_{21}} \times 100\%$ , where  $P_t$  is total cooling power of the ten server rooms with the same temperature setpoint of  $t^\circ\text{C}$ . Note that DCs typically adopt a temperature setpoint between  $20^\circ\text{C}$  and  $22^\circ\text{C}$ .



**Figure 2: Impact of temperature setpoint and IT load on (a) the cooling power and (b) relative cooling power saving.**

Thus, we investigate the relative cooling power saving with respect to the baseline of 21°C.

Fig. 2 shows the total cooling powers and corresponding relative power savings versus temperature setpoint under various IT loads. The cooling power decreases with the temperature setpoint. Specifically, given a certain per-room IT load between 20 kW and 50 kW, the relative power saving  $\sigma_t$  increases sharply when the temperature setpoint increases from 21°C to 25°C. Then, the relative power savings flatten out when the temperature setpoint is greater than 25°C. Note that the temperature value of 25°C for such an observation depends on the detailed settings. However, it is intuitive that the increment of 1°C in temperature setpoint when the setpoint is low brings more relative power savings than that when the setpoint is high. Overall, under a certain IT load, the cooling power saving has a non-linear relationship with the temperature setpoint. The reason is as follows. From the model of AHU, i.e.,  $\dot{Q}_{ri} = c_a \dot{m}_{a_i} (T_{ar_i} - T_{as})$ , to remove a certain amount of heat  $\dot{Q}_{ri}$ , the mass air flow rate  $\dot{m}_{a_i}$  is lower if the temperature setpoint  $T_{ar_i}$  is higher. Thus, the AHU's internal fans can run at lower speeds and use less power.

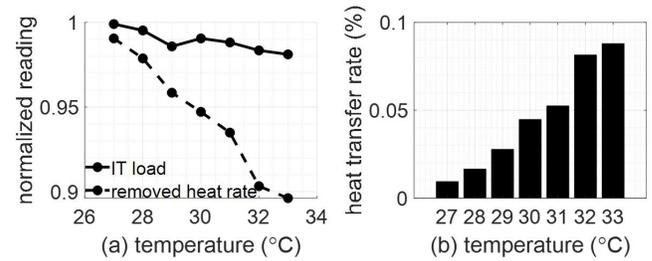
From Fig. 2(a), under the same temperature setpoint, the cooling power greatly increases with the IT load. This is because the cooling system needs to operate at a higher cooling capability to remove more heat generated by the IT equipment. In addition, from Fig. 2(b), the increased IT load results in a larger relative cooling power saving under the same temperature setpoint. For instance, with the temperature setpoint of 31°C, the relative power saving increases from 2% to 14% when the IT load increases from 20 kW to 50 kW. The reason is that the second-stage cooling system in general has higher power efficiency in moving more heat.

**Observation 1:** The relative cooling power saving increases with the temperature setpoint and IT load.

## 4.2 Impact of Heat Transfer on Cooling Power

In this subsection, we conduct a set of experiments with a real server room to quantify the amount of heat transfer in reality. Then, we study the impact of the inter-room heat transfers on the DC's cooling power usage based on the cooling system model in §3.

**4.2.1 Heat transfer from a server room.** In this experiment, we operate a server room hosting a number of IT racks. The server room has cement separations from its ambient. An AHU is deployed in this room to move the heat generated by the IT equipment to the



**Figure 3: Heat transfer from a real server room. (a) IT load and removed heat rate normalized with respect to a constant; (b) Relative heat transfer rate with respect to the corresponding IT load.**

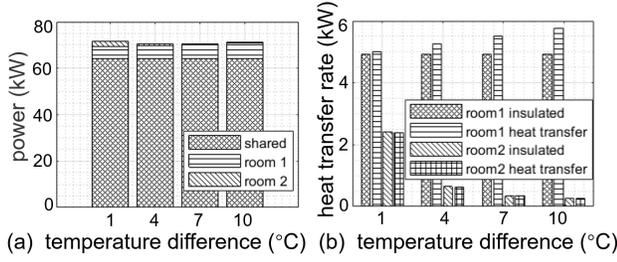
second-stage cooling system. The AHU is equipped with meters to measure the rate of the heat moved from the room and the return air temperature. The power distribution unit of the room provides real-time IT load readings. We run experiments in which the temperature setpoint varies from 27°C to 33°C. Under a certain temperature setpoint, the experiment lasts for 12 hours. The server room is located in a building that has a constant ambient temperature lower than the lowest temperature setpoint in our server room (i.e., 27°C).

Fig. 3(a) shows the IT load and the rate of heat removed by the AHU, both normalized with a constant, versus the temperature setpoint. Each point is based on the average value of the measurements over 12 hours. In Fig. 3(a), the slight decrease of IT load with temperature during the course of the experiment is due to server workload change. We can see that the removed heat rate is always less than the IT load. This suggests that a portion of the heat generated by the IT equipment is transferred out through the enclosure of the server room. We plot the relative heat transfer rate with respect to the corresponding IT load in Fig. 3(b). The relative heat transfer rate exhibits a linearly increasing trend with the room temperature. This verifies the heat transfer model in §3.1. When the room temperature is 33°C, the heat transfer rate is about 9% of IT load. A cooling system expert also investigated the area and the material of the server room's enclosure. Our experiment result and the expert's roughly calculated heat transfer rate match.

**Observation 2:** The inter-room heat transfer can be a significant factor of the heat processes, should the server rooms adopt distinct temperature setpoints.

**4.2.2 Impact of heat transfer on cooling power.** We run numeric experiments based on the cooling system model in §3. We consider a co-location DC consisting of two neighboring server rooms, i.e., room 1 and room 2. Room 1's temperature setpoint is fixed at 21°C. Room 2's temperature setpoint varies from 22°C to 31°C. As a result, a portion of the heat generated in room 2 is transferred to room 1. In the experiments, we set the heat transfer coefficient between the two rooms as  $\alpha_{12} = 0.27$  kW/°C.

Fig. 4(a) shows the total cooling power and its breakdown to the shared second-stage cooling system and the two rooms' AHUs under various differences between the two rooms' temperatures. The total cooling power changes slightly with the temperature difference. The power usage of the second-stage cooling system



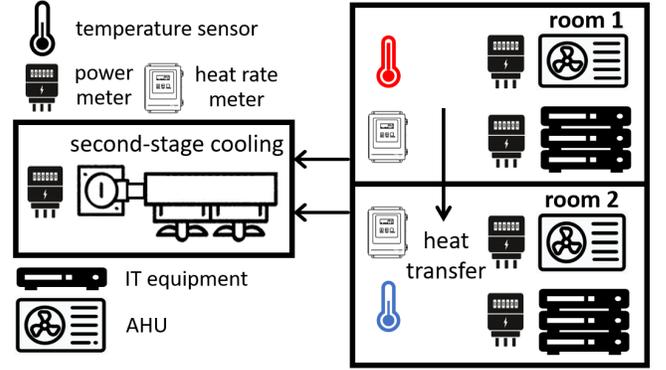
**Figure 4: Impact of heat transfer on cooling power. (a) Cooling power and breakdown; (b) AHU's cooling power.**

remains the same across various temperature differences. This implies that the heat transfer does not affect the power usage of the second-stage cooling. This can be formally stated as follows.

**PROPOSITION 4.1.** *Under the cooling system model described in §3, the second-stage cooling system's power usage depends on the total IT load and is not affected by the inter-room heat transfers.*

**PROOF.** From the conservation of heat, we have  $\sum_{i=1}^n Q_{tr_i} = 0$ . Thus, the rate of heat removed by the chiller from the chilled water cycle  $\dot{Q}_{ch} = \sum_{i=1}^n \dot{Q}_{r_i} = \sum_{i=1}^n P_{IT_i} + \dot{Q}_{tr_i} = \sum_{i=1}^n P_{IT_i}$ . Moreover, the mass flow rate of the chilled water is  $\dot{m}_{chw} = \sum_{i=1}^n \dot{m}_{chw_i} = \sum_{i=1}^n \frac{\dot{Q}_{r_i}}{c_w(T_{chwr} - T_{chws})} = \frac{\sum_{i=1}^n P_{IT_i}}{c_w(T_{chwr} - T_{chws})}$ , where  $c_w$ ,  $T_{chwr}$ , and  $T_{chws}$  are constants. Thus, the chiller plant's operating status that is characterized by  $\dot{Q}_{ch}$  and  $\dot{m}_{chw}$  only depends on  $\sum_{i=1}^n P_{IT_i}$ . Hence, the second-stage cooling power depends on  $\sum_{i=1}^n P_{IT_i}$  only.  $\square$

Differently, as shown in Fig. 4(b), the power usage of each room's AHU changes with the temperature difference. To further investigate the impact of the heat transfer on the AHU power usage, we run additional numeric experiments in which the two rooms are thermo-insulated (i.e., no heat transfer). Fig. 4(b) presents the power usages of the two rooms' AHUs in the presence and absence of thermal insulation. In the presence of thermal insulation, the power usage of room 1's AHU remains the same, since the heat generation rate in room 1 (i.e., the IT load only) is constant. In the absence of thermal insulation, the power usage of room 1's AHU increases with the temperature difference. This is because more heat is transferred from room 2 to room 1 when the temperature setpoint of room 2 increases. As a result, room 1's AHU needs to operate its internal fans at higher rotation speeds to remove more heat to maintain the temperature at the setpoint. In the presence and absence of thermal insulation, the power usage of room 2's AHU decreases with the temperature difference. The reason is that a higher temperature setpoint allows the AHU to operate its internal fans at lower speeds. The power usage of room 2's AHU in the presence of thermal insulation is higher than that in the absence of thermal insulation. However, room 2's AHU power usage reduction caused by the heat transfer is less significant in comparison with room 1's AHU power usage increase. This is because AHU power usage is non-linear with the heat removal rate. In this two-room example, the heat transfer results in higher total power usage of the two rooms' AHUs, in comparison with the case with thermal



**Figure 5: A two-room example. Arrow represents heat flow.**

insulation. This suggests that the impact of heat transfers on the cooling power usage should be considered.

**Observation 3:** The heat transfers among server rooms with distinct temperatures affect the AHU power usages. They do not affect the power usage of the second-stage cooling system.

## 5 PROBLEM AND APPROACH OVERVIEW

### 5.1 Problem Statement

We consider a co-location DC consisting of multiple server rooms with distinct return air temperature setpoints. The DC uses the two-stage cooling system described in §3. Each server room has at least four meters to measure return air temperature, IT load, the rate of heat removed by the AHU, and AHU power usage  $P_{AHU_i}$ . The DC operator also deploys a meter to measure the second-stage cooling system's power usage  $P_{c_w}$ . Fig. 5 shows a minimal example of a co-location DC with two server rooms, illustrating the required meters. In this paper, we study the problem of attributing the total cooling power  $P = \sum_{i=1}^n P_{AHU_i} + P_{c_w}$  to the server rooms. As discussed in §3.2, we assume that the DC operator has models of  $P_{AHU_i}$  (i.e.,  $f_1$ ) and  $P_{c_w}$  (i.e., the composite of  $f_2 + f_3 + f_4 + f_5$ ).

The attribution is challenging due to the following reasons. First, due to the inter-room heat transfers, the measured  $P_{AHU_i}$  may be different from the power that the AHU<sub>*i*</sub> is supposed to use to remove the heat generated by the room *i*'s IT equipment only. The  $P_{AHU_i}$  can include power usage for removing heat transferred from the neighbours with higher temperatures. Also, the  $P_{AHU_i}$  can be lower than the supposed AHU power usage if the room *i* transfers heat to its neighbors. Therefore, the  $P_{AHU_i}$  measurement cannot be directly attributed to room *i* as its AHU power usage. Second, as discussed in §3, as  $P$  is non-linear and analytically indivisible, there is no straightforward attribution to server rooms.

The load-proportional division (LPD) policy is a prevailing power attribution mechanism. Its simplest form, which is widely adopted, charges the tenant by: electricity tariff (\$/kWh)  $\times$  IT electricity usage (kWh)  $\times$  a constant that factors in the power usage effectiveness (PUE) of the whole DC infrastructure. When server rooms' adopt distinct temperatures, the above simplest LPD policy falls short of considering the impacts of room temperature and inter-room heat transfers on AHU power usage. It can be improved by charging the tenant for AHU power based on meter reading and

second-stage cooling power based on LPD separately. It still does not address inter-room heat transfers. In addition, applying LPD on the analytically indivisible power usage of the second-stage cooling system is a solution offering no profound fairness explanations.

## 5.2 Approach Overview

From our measurements and analysis in §4, the server rooms' temperature setpoints and the inter-room heat transfers affect the power usages of the rooms' AHUs, and do not affect the power usage of the second-stage cooling system. The latter only depends on the total IT load of the rooms. Therefore, in our proposed power attribution scheme, we address the following two sub problems in §6 and §7 respectively: (1) AHU power rectification and (2) second-stage cooling power attribution. The AHU power rectification aims at rectifying the metered AHU power usage  $P_{AHU_i}$  of each room  $i$ , such that room  $i$  is attributed with the supposed AHU power usage that removes the heat generated by room  $i$ 's IT equipment only. To this end, we perform two steps. First, we develop a data-driven approach to estimate the heat transfer coefficient between any two server rooms based on the historical meter measurements. Then, the real-time inter-room heat transfers are estimated and the extra AHU powers that are used to remove the incoming heat from neighboring rooms are determined and attributed back to these neighbors. The second-stage cooling power attribution adopts the Shapley value method based on the server rooms' IT loads. The attribution by the Shapley value method meets several fairness axioms. Our main objective in this paper is to avoid the intensive and long-lasting computation of the Shapley value at run time, because the instantaneous power attribution needs to be performed with short periods (e.g., every one minute). To this end, we design an MLP and train it with sufficient power attribution samples computed by the Shapley value method offline. At run time, by feeding the MLP with the server rooms' real-time IT loads, the MLP inference gives the attribution with low latency. However, the offline data generation still incurs overhead. Thus, we also develop a fast heuristic attribution algorithm based on observations on Shapley attribution, achieving good accuracy and scalability with  $n$ .

## 6 AHU POWER RECTIFICATION

Inter-room heat transfers are needed to perform the rectification. However, under a general setting, we cannot estimate the  $\binom{n}{2}$  inter-room heat transfer values from a system of  $n$  equations instrumented with the meter measurements, because it is an underdetermined problem. Although the number of unknown heat transfers can be reduced by considering building topology (i.e., not every two rooms are adjacent), it may not turn the problem determined. This issue motivates us to resort to estimating all inter-room heat transfer coefficients, because with these, we can estimate the instantaneous heat transfers based on the rooms' instantaneous return air temperatures. §6.1 presents a data-driven approach to estimate the heat transfer coefficients. §6.2 presents the rectification approach.

### 6.1 Estimation of Heat Transfer Coefficients

The heat transfer coefficient depends on the material and area of the separation structure. The estimation approaches described in existing studies [9, 21] require detailed information such as the

separation material properties. These approaches incur tedious processes of modeling all room separations. This paper proposes a data-driven approach that estimates the heat transfer coefficients merely based on the historical measurements of server room temperatures, heat removal rates, and IT loads over multiple time steps.

Our analysis uses the following notations. The  $\alpha_{ij}$  denotes the unknown heat transfer coefficient between rooms  $i$  and  $j$ . For the  $k$ th time step:  $\dot{Q}_{r_i}[k]$  denotes the measured heat rate at which the AHU removes heat from room  $i$ ;  $P_{IT_i}[k]$  denotes the measured IT load of room  $i$ ;  $\dot{Q}_{tr_i}[k] = \dot{Q}_{r_i}[k] - P_{IT_i}[k]$  denotes room  $i$ 's net influx heat transfer rate;  $T_{ij}[k] = T_i[k] - T_j[k]$  denotes the difference between the measured return air temperatures of rooms  $i$  and  $j$ . From the first principle governing the heat transfer as presented in §3.1, we have  $\dot{Q}_{tr_i}[k] = \sum_{p,q \in [1,N], p < q, q=i} T_{pq}[k] \cdot \alpha_{pq} - \sum_{p,q \in [1,N], p < q, p=i} T_{pq}[k] \cdot \alpha_{pq}$ , where the first sum is room  $i$ 's total influx heat transfer rate and the second sum is room  $i$ 's total outflow heat transfer rate. The above equality can be vectorized as

$$\begin{pmatrix} \dot{Q}_{tr_1}[k] \\ \dot{Q}_{tr_2}[k] \\ \vdots \\ \dot{Q}_{tr_n}[k] \end{pmatrix} = \begin{pmatrix} -T_{12}[k] & -T_{13}[k] & \cdots & 0 \\ T_{12}[k] & 0 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & T_{(n-1)n}[k] \end{pmatrix} \begin{pmatrix} \alpha_{12} \\ \alpha_{13} \\ \vdots \\ \alpha_{(n-1)n} \end{pmatrix}.$$

We write the above equation as  $\dot{Q}_{tr}[k] = \mathbf{T}[k]\boldsymbol{\alpha}$ , where  $\dot{Q}_{tr}[k] \in \mathbb{R}^{n \times 1}$ ,  $\mathbf{T}[k] \in \mathbb{R}^{n \times \binom{n}{2}}$ , and  $\boldsymbol{\alpha} \in \mathbb{R}^{\binom{n}{2} \times 1}$ . It is underdetermined. If the return air temperatures vary over  $K$  time steps, we can integrate the equations into a single equation. Specifically, by defining matrices  $\dot{Q}_{tr} = (\dot{Q}_{tr}[1]; \dot{Q}_{tr}[2]; \dots; \dot{Q}_{tr}[K]) \in \mathbb{R}^{Kn \times 1}$  and  $\mathbf{T} = (\mathbf{T}[1]; \mathbf{T}[2]; \dots; \mathbf{T}[K]) \in \mathbb{R}^{Kn \times \binom{n}{2}}$ , the integrated equation is  $\dot{Q}_{tr} = \mathbf{T}\boldsymbol{\alpha}$ . A necessary condition for this equation to be determined is  $K \geq \frac{n-1}{2}$ . However, if the return air temperature variations over time are small, the matrix  $\mathbf{T}$  may be ill-conditioned. Thus, it is beneficial to integrate many time steps more than the necessary condition to ensure that the integrated equation is overdetermined and the least squares approach can be applied to solve  $\boldsymbol{\alpha}$ .

For a new co-location DC before commission, the DC operator can perform controlled experiments to vary the server rooms' temperatures and collect data for the heat transfer coefficient estimation. For a commissioned co-location DC, since the server rooms' IT loads vary over time and the AHUs have control dynamics, the room temperatures may deviate from their setpoints. The DC operator may select data in  $K$  time steps that are unnecessarily continuous from years' operation history, such that  $\mathbf{T}$  has a good condition number that can be assessed by  $\|\mathbf{T}\|_F \|\mathbf{T}^+\|_F$  [16], where  $\mathbf{T}^+$  is the pseudoinverse of  $\mathbf{T}$  and  $\|\cdot\|_F$  represents Frobenius norm.

### 6.2 Rectification of AHU Power Usages

We face two major challenges. The first is caused by the non-linear relationship between AHU's power usage and its rate of heat removal. Thus, the AHU may use different powers to move the same amount of heat when it operates on different conditions including the return hot air temperature. Considering the simplest two-room case, a room's AHU power increment and the other's decrement caused by the heat transfer is in general different. Thus, in the absence and presence of thermal insulation, the total AHU power usages of the two rooms are different. This has been observed in §4.2.2. Therefore, the hypothetical case of ideally thermo-insulated

rooms is not an ideal target of the rectification. In our proposed approach, we follow a principle of using the heat transfer-induced AHU power usage increment of the lower-temperature room to rectify the power usages of the two involved rooms. Under this principle, the sum of all rooms' rectifications is zero. The rationale of choosing the lower-temperature room is from the observation in Fig. 4(b) that the heat transfer has greater impact on the lower-temperature room's AHU power.

The second challenge is that, if a room's influx heat transfer is from multiple rooms, the room's power usage increment is indivisible. We address this as follows. Considering room  $i$  with metered PMU power  $P_{AHU_i}$ , we first determine the variation of the air mass flow rates caused by the heat rates transferred to room  $i$  from the rooms with higher temperatures. For instance, considering a higher-temperature room  $k$  with metered PMU power  $P_{AHU_k}$ , the transferred heat rate from room  $k$  to room  $i$  is  $\dot{Q}_{trki} = \alpha_{ik}(T_k - T_i)$ , where  $\alpha_{ik}$  is obtained in §6.1. Therefore, the increment of room  $i$ 's AHU air mass flow rate due to  $\dot{Q}_{trki}$ , denoted by  $\Delta\dot{m}_{aki}$ , is  $\Delta\dot{m}_{aki} = \frac{\dot{Q}_{trki}}{c_a(T_i - T_{as})}$ . We adopt the linear approximation to estimate the increment of the AHU power usage (denoted by  $\Delta P_{AHU_{ki}}$ ) caused by  $\Delta\dot{m}_{aki}$  as  $\Delta P_{AHU_{ki}} = \nabla f_1(\dot{m}_{ai})\Delta\dot{m}_{aki}$ , where  $\dot{m}_{ai}$  is the current air mass flow rate of room  $i$ 's AHU and  $\nabla f_1(\cdot)$  represents the first derivative of the function  $f_1(\cdot)$  defined in §3.2. By following the aforementioned principle of choosing the lower-temperature room as the common basis of rectification, we subtract  $\Delta P_{AHU_{ki}}$  from  $P_{AHU_i}$  and add it to  $P_{AHU_k}$ . We follow above to rectify the AHU powers of any two rooms having heat transfer.

The rectification process can be modeled by a directed graph, in which the nodes represent server rooms and a directed edge represents the directional heat transfer. Fig. 6 illustrates the graph for a 3-room case, in which  $T_1 < T_2 < T_3$ . Each node is associated with the room temperature and metered AHU power. For the edge from node  $k$  to node  $i$ , we compute the heat transfer rate  $\dot{Q}_{trki}$ , increment of air mass flow rate  $\Delta\dot{m}_{aki}$ , and increment of power usage of the end node's AHU  $\Delta P_{AHU_{ki}}$ . The  $\Delta P_{AHU_{ki}}$  is the edge cost. Accordingly, the rectified AHU power (denoted by  $\hat{P}_{AHU_i}$ ) is the original AHU power subtracted with all incoming edge costs and added with all outgoing edge costs. For instance, for rooms 1 and 2 in Fig. 6,  $\hat{P}_{AHU_1} = P_{AHU_1} - \Delta P_{AHU_{21}} - \Delta P_{AHU_{31}}$  and  $\hat{P}_{AHU_2} = P_{AHU_2} - \Delta P_{AHU_{32}} + \Delta P_{AHU_{21}}$ .

## 7 SECOND-STAGE COOLING ATTRIBUTION

This section formulates the second-stage cooling power attribution problem and presents our MLP-based and fast heuristic approaches.

### 7.1 Fairness Objective and Challenge

Denote by  $N$  the set of all  $n$  server rooms, by  $P_{c_w}(N)$  the power usage of the second-stage cooling to remove the heat generated by all server rooms, by  $P_{c_{wr}}(i)$  the attribution of  $P_{c_w}(N)$  to room  $i$ . The attribution aims to meet the following three fairness axioms:

**Efficiency:**  $\sum_{i=1}^n P_{c_{wr}}(i) = P_{c_w}(N)$ .

**Symmetry:** If rooms  $i$  and  $j$  contribute equally to the second-stage cooling power, their attributions are same, i.e., if  $P_{c_w}(M \cup \{i\}) = P_{c_w}(M \cup \{j\})$ ,  $\forall M \subseteq N \setminus \{i, j\}$ , then  $P_{c_{wr}}(i) = P_{c_{wr}}(j)$ .

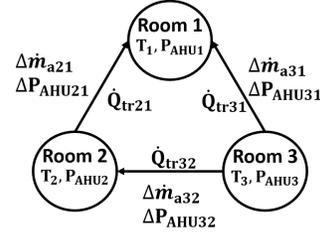


Figure 6: AHU power rectification graph for three rooms.

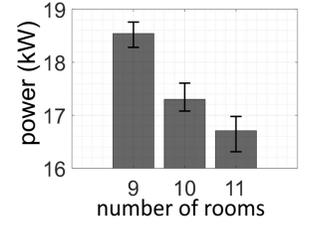


Figure 7: Room 1's Shapley attributions in three cases.

**Dummy player:** If the second-stage cooling power does not change in the presence or absence of room  $i$ , room  $i$  has zero attribution, i.e., if  $P_{c_w}(M \cup \{i\}) = P_{c_w}(M)$ ,  $\forall M \subseteq N \setminus \{i\}$ , then  $P_{c_{wr}}(i) = 0$ .

The widely adopted LPD policy does not guarantee the symmetry axiom. As the second-stage cooling power has a non-linear relationship with the rooms' IT loads, two rooms with different IT loads may have the same contribution to the second-stage cooling power. However, the LPD policy will assign different attributions to the two rooms. The Shapley value approach [19] that is a fair value attribution game method has been proven to meet all the three fairness axioms [11]. We take it as the objective as it is considered the only fair method for cost sharing game [8]. Applying the Shapley value approach in our cooling power attribution context, the  $P_{c_w}(M)$  and  $P_{c_w}(N)$  can be considered as the characteristic cost function of a coalition consisting of a subset  $M$  of  $m$  server rooms and the cost of the entire server room set  $N$ , respectively. Under the Shapley value approach,

$$P_{c_{wr}}(i) = \sum_{M \subseteq N, i \notin M} \frac{(|N| - |M| - 1)! |M|!}{|N|!} (P_{c_w}(M \cup \{i\}) - P_{c_w}(M)). \quad (1)$$

However, Eq. (1) is computationally intensive. Specifically, the complexity of computing  $P_{c_{wr}}(i)$  is  $O(2^n)$ . Thus, the compute complexity of the second-stage cooling power attribution is  $O(n \cdot 2^n)$ . Such a high compute complexity does not allow the real-time power attribution when  $n$  is beyond a certain value.

### 7.2 Power Attribution Approaches

This section presents two approaches to achieving real-time attribution of the second-stage cooling power that approximates the Shapley attribution discussed in §7.1. The first approach trains an MLP using training samples of Shapley attribution generated offline and forwards the MLP for real-time attribution. The second approach uses a lightweight heuristic algorithm developed based on a key observation of Shapley cooling attribution. The MLP-based approach achieves high approximation accuracy but requires intensive offline computation. The heuristic approach is training-free, lightweight, but less accurate. They will be evaluated in §8.

**7.2.1 MLP-based attribution approach.** The MLP takes the server rooms' real-time IT loads  $P_{IT_1}, \dots, P_{IT_n}$  as inputs to predict the attributions  $P_{c_{wr}}(1), \dots, P_{c_{wr}}(n)$ . The MLP is trained offline using the data samples of  $P_{c_{wr}}(i)$  ( $i = 1, \dots, n$ ) which are generated by feeding the Shapley power attribution function in Eq. (1) with random IT loads of the rooms. Our current parallel implementation

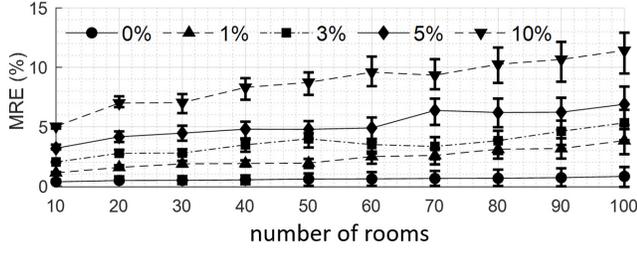


Figure 8: Accuracy of heat transfer coefficient estimation.

of the data generation on a workstation equipped with two 12-core Intel Xeon processors can handle up to 18 server rooms. To handle more server rooms, cloud computing can be used. Note that the training data generation is a one-time effort. The trained MLP is used online to predict the rooms' second-stage cooling power usages based on the rooms' instantaneous IT loads in real time.

**7.2.2 Fast Heuristic attribution approach.** When  $n$  increases, the offline training data generation for the MLP becomes harder. For the cases of large  $n$ , we develop a fast heuristic attribution approach based on a key observation as follows: under the second-stage cooling system model described in §3, given  $n$ , the Shapley attribution for a certain IT load has small variations when the other IT loads vary. We now illustrate this using a numeric example. We consider three cases with  $n = 9$ ,  $n = 10$ , and  $n = 11$ , respectively. We fix the IT load of room 1 to be 30 kW. We conduct 1,000 random experiments as follows. In each experiment, we randomly generate the IT loads of the rooms (except room 1) under the three cases. Each room's IT load is within  $[0, 50]$  kW. Moreover, in each experiment, the total load of all rooms in the 10-room case is same as that of the 11-room case, and different from that of the 9-room case. Fig. 7 shows room 1's Shapley attribution  $P_{Cwr}(1)$  under the three cases. The error bars represent the maximum and minimum during the 1,000 experiments. We can see that (i) the  $n$  has a major impact on  $P_{Cwr}(1)$  and (ii) the  $P_{Cwr}(1)$  has maximal fluctuations of 0.47 kW, 0.52 kW, and 0.66 kW for the three cases, respectively, which are just 2.5%, 3%, 3.9% of the respective  $P_{Cwr}(1)$  averages.

Based on above, we compute room  $i$ 's approximated Shapley attribution as  $P_{Cwr}(i) = \frac{P_{Cw}(n \times P_{IT_i})}{n}$ . It is based on a hypothetical case in which each of the  $n$  rooms has identical IT load of  $P_{IT_i}$ ; thus, the Shapley attribution for each room is the second-stage cooling power  $P_{Cw}(n \times P_{IT_i})$  divided by  $n$ . The compute complexity is  $O(n)$ .

## 8 PERFORMANCE EVALUATION

This section evaluates the proposed power attribution scheme by numeric experiments driven by the models in §3. We adopt settings recommended by ASHRAE [7], e.g.,  $T_{as} = 17^\circ\text{C}$ ,  $T_{chws} = 7^\circ\text{C}$ , and  $T_{chwr} = 12^\circ\text{C}$ . In this section, an error bar represents the maximum and minimum over 1,000 repeated experiments.

### 8.1 Performance of AHU Power Rectification

**8.1.1 Accuracy of heat transfer coefficient estimation.** We conduct experiments with  $n$  increasing from 10 to 100. We generate the adjacent relationship between any two rooms by sampling a Bernoulli distribution with an adjacency probability of 0.6. The ground-truth

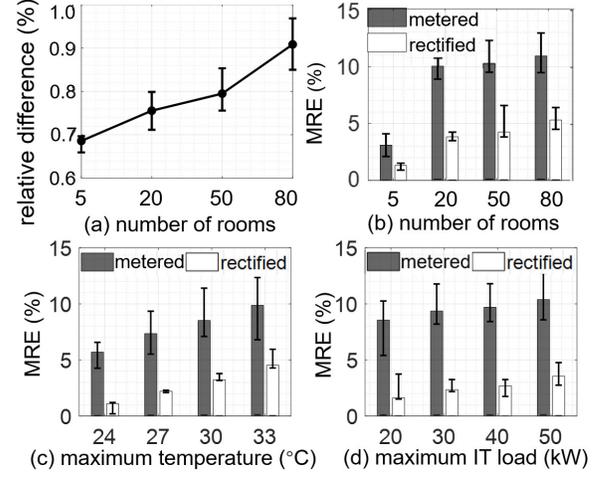


Figure 9: Performance of AHU power rectification.

heat transfer coefficient between any two adjacent rooms  $i$  and  $j$  is randomly and uniformly sampled from  $[0.03, 0.05]$  kW/ $^\circ\text{C}$ . All rooms' IT loads are 30 kW. We generate historical data for  $K = 200n$  time steps, which is much more than the necessary condition of  $K \geq \frac{n-1}{2}$ . In each time step  $k$ , the return air temperature  $T_{ar_i}$  of room  $i$  is randomly selected from  $[21, 30]^\circ\text{C}$ . The ground-truth removed heat rates  $Q_{tr_i}[k]$  ( $i = 1, \dots, n$ ) are calculated from the IT loads  $P_{IT_i}$  and the ground-truth inter-room heat transfer rates. We account for measurement noises by adding random noises to the measured net influx heat transfer rates. Specifically, the simulated measurement is  $\hat{Q}_{tr_i}[k] = Q_{tr_i}[k] + \epsilon$ , where  $Q_{tr_i}[k]$  is the ground truth,  $\epsilon$  is a random noise drawn uniformly from  $[-\delta Q_{tr_i}[k], \delta Q_{tr_i}[k]]$ , and  $\delta$  controls the noise level. In our experiments,  $\delta$  varies from 0 to 10% that is the maximum error of real heat meters [13]. By defining relative error as  $\frac{|E-T|}{T}$ , where  $E$  and  $T$  are the estimated and true values, we use the mean relative error (MRE) over all heat transfer coefficients as the accuracy metric. Fig. 8 shows the MREs under various settings of  $n$  and  $\delta$ . The MRE increases with  $n$  and  $\delta$ , which is consistent with intuition. When the measurements are noiseless (i.e.,  $\delta = 0$ ), the MRE is at most 1.7% when  $n$  is up to 100. When  $\delta = 10\%$ , the MRE is from 5% to 11.4%, similar to  $\delta$ . These results suggest that our estimation approach scales well with  $n$  and  $\delta$ .

**8.1.2 Performance of AHU power rectification.** We perform evaluation under wide ranges of settings for  $n$ , IT load, and temperature setpoint. We set  $\delta = 1\%$ . **First**, we vary  $n$ . Under each setting of  $n$ , each room's temperature setpoint is randomly drawn from  $[21, 30]^\circ\text{C}$  while the room's IT load is fixed at 30 kW. As discussed in §6.2, the hypothetical case of ideally thermo-insulated rooms is not a gold standard, because the inter-room heat transfers will lead to change of total AHU power. Fig. 9(a) shows the relative difference between the total AHU powers in the absence and presence of thermal insulation versus  $n$ . The relative difference is less than 1%. Thus, the AHU powers in the hypothetical case of ideally thermo-insulated rooms can be used as good targets of the AHU power rectification. Therefore, we evaluate the MREs of the

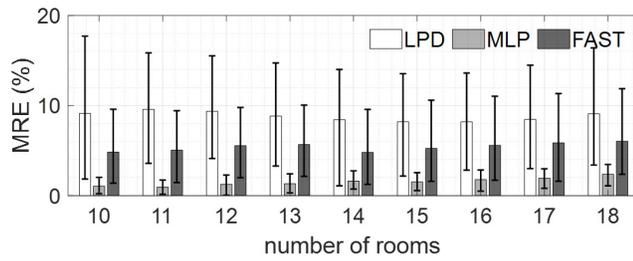


Figure 10: Second-stage cooling power attribution.

metered AHU power and the rectified result with respect to that in the hypothetical case. Fig. 9(b) shows the results. We can see that the metered AHU powers deviate from the thermo-insulated case by up to 13% when  $n = 80$ . Our rectifications reduce MREs by more than half. **Second**, we vary the maximum temperature. Each room’s temperature setpoint is randomly selected from 21°C to the maximum temperature. Each room’s IT load is fixed at 30 kW. Fig. 9(c) shows the results. MRE increases with the maximum temperature. This is because a larger setting of maximum temperature leads to larger amounts of heat transfers. Our rectifications also reduce MREs by more than half compared with the metered AHU powers. **Third**, we vary the maximum IT load. Each room’s IT load is randomly selected from 0 to the maximum IT load. We set  $n = 50$  and maximum temperature to be 30°C. Fig. 9(d) shows the results. Our rectifications reduce MREs by up to 78%.

## 8.2 Performance of MLP & Fast Approaches

For different  $n$  settings, we build different MLPs. An MLP consists of an input layer, multiple hidden layers, and a linear output layer. Rectified linear unit (ReLU) is used as the activation function for the input and hidden layers. Each MLP is trained with 1,000 samples, for which each room’s IT load is randomly sampled from  $[0, 50]$  kW. The Adam optimizer with a learning rate of 0.001 is used for training; training batch size is 128. The test dataset consists of another 1,000 samples. For each  $n$  setting, we conduct extensive evaluation to choose the hyperparameter settings including the numbers of hidden layers and neurons each layer, to minimize the root mean square error (RMSE) between the prediction and the Shapley ground truths of training data. Due to the space constraints, we omit presenting the optimal hyperparameter settings.

We compare the performance of the prevailing LPD policy, and our proposed MLP and fast heuristic (FAST) approaches. We measure the relative error of each room’s attribution with respect to the Shapley ground truth. We use the MRE over all rooms as the accuracy metric. Fig. 10 shows MREs of the three approaches versus  $n$ . The MRE of LPD can be up to 19%. MLP is the most accurate. Its average MRE is 2.38% when  $n$  is 18. FAST gives higher MREs compared with MLP, but lower compared with LPD. FAST’s average MRE is 6.04%. From the above results, MLP can approximate the Shapley attribution function accurately. However, it requires a compute-intensive training data generation. When the computing resources are insufficient for generating training data, FAST is an acceptable alternative.

## 9 CONCLUSION

This paper proposes a real-time cooling power attribution scheme for co-located server rooms with distinct temperatures. First, it rectifies the metered power usages of AHUs to address inter-room heat transfers that are estimated in real time based on the heat transfer coefficients obtained via data analytics. Second, it uses an algorithm to compute the rooms’ approximated Shapley shares of the second-stage cooling system’s power usage. Evaluation shows the effectiveness of the proposed scheme. With it implemented, a co-location DC can encourage the tenants to raise their temperature setpoints subject to their own different constraints such that the DC’s energy efficiency is improved and the tenants to hotter server rooms receive cooling cost savings in return.

## REFERENCES

- [1] [n.d.]. <https://www.datacentermap.com/datacenters.html>.
- [2] [n.d.]. <https://www.datacenter-widnau.com/en/construction.html>.
- [3] 2013. *Singapore Data Center Energy Consumption Study*. Technical Report. Infocomm Development Authority of Singapore (IDA).
- [4] ASHRAE TC 9.9. 2011. Thermal Guidelines for Data Processing Environments - Expanded Data Center Classes and Usage Guidance. *White Paper* (2011).
- [5] Douglas Alger. 2009. *Grow a Greener Data Center*. Cisco Press.
- [6] Afef Gafsi and Gilles Lefebvre. 2003. Stolen heating or cooling energy evaluation in collective buildings using model inversion techniques. *Energy and buildings* 35, 3 (2003), 293–303.
- [7] Supriya Goel, Michael I Rosenberg, and Charles Eley. 2017. *ANSI/ASHRAE/IES Standard 90.1-2016 Performance Rating Method Reference Manual*. Technical Report. Pacific Northwest National Lab (PNNL), Richland, WA (United States).
- [8] Ragavendran Gopalakrishnan, Jason R Marden, and Adam Wierman. 2014. Potential games are necessary to ensure pure nash equilibria in cost sharing games. *Mathematics of Operations Research* 39, 4 (2014), 1252–1296.
- [9] Anti Hamburg, Martin Thalfeldt, Teet-Andrus Kõiv, and Alo Mikola. 2014. Investigation of heat transfer between neighbouring apartments. In *Proceedings of the International Conference on Environmental Engineering (ICEE)*, Vol. 9, 1.
- [10] Mohammad A Islam, Hasan Mahmud, Shaolei Ren, and Xiaorui Wang. 2015. Paying to save: Reducing cost of colocation data center via rewards. In *IEEE 21st International Symposium on High Performance Computer Architecture (HPCA)*.
- [11] Mohammad A. Islam and Shaolei Ren. 2016. A New Perspective on Energy Accounting in Multi-Tenant Data Centers. In *USENIX Workshop on Cool Topics on Sustainable Data Centers (CoolDC 16)*. Santa Clara, CA.
- [12] Weixiang Jiang, Shaolei Ren, Fangming Liu, and Hai Jin. 2018. Non-it energy accounting in virtualized datacenter. In *ICDCS*. IEEE, 300–310.
- [13] AECOM Limited. 2018. *An Investigation into Heat Meter Measurement Errors Final Report*. Technical Report. Department of Energy & Climate Change, UK.
- [14] Paweł Michnikowski. 2017. Allocation of heating costs with consideration to energy transfer from adjacent apartments. *Energy and Buildings* 139 (2017).
- [15] David Moss and John H Bean. 2009. Energy impact of increased server inlet temperature. *APC white paper* 138 (2009).
- [16] Christopher C Paige and Michael A Saunders. 1982. LSQR: An algorithm for sparse linear equations and sparse least squares. *ACM Transactions on Mathematical Software (TOMS)* 8, 1 (1982), 43–71.
- [17] Michael K Patterson. 2008. The effect of data center temperature on energy efficiency. In *2008 11th Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems*. IEEE, 1167–1174.
- [18] John Rath. 2011. Data Centers Can be Even Warmer. <https://bit.ly/2ppu0sq>.
- [19] Lloyd S Shapley. 1953. A value for  $n$ -person games. *Contributions to the Theory of Games* 2, 28 (1953), 307–317.
- [20] Arman Shehabi and et al. 2016. *United states data center energy usage report*. Technical Report. Lawrence Berkeley National Lab (LBNL), Berkeley, CA.
- [21] Simon Siggelsten. 2014. Reallocation of heating costs due to heat transfer between adjacent apartments. *Energy and Buildings* 75 (2014), 256–263.
- [22] Matt Stansberry. 2013. Uptime Institute-Data Center Industry Survey. (2013).
- [23] Shailja Thakur, Manaswi Saha, Amarjeet Singh, and Yuvraj Agarwal. 2014. Wattshare: Detailed energy apportionment in shared living spaces within commercial buildings. In *BuildSys*. 30–39.
- [24] Peter Wei, Xiaoqi Chen, Jordan Vega, Stephen Xia, Rishikanth Chandrasekaran, and Xiaofan Jiang. 2017. ePrints: a real-time and scalable system for fair apportionment and tracking of personal energy footprints in commercial buildings. In *BuildSys*. 1–10.
- [25] Peng Xue, Fan Yang, Ye Zhang, Mengjing Zhao, Jingchao Xie, and Jiaping Liu. 2019. Quantitative study on adjacent room heat transfer: Heating load and influencing factors. *Sustainable Cities and Society* 51 (2019), 101720.