

CoolEdge: Hotspot-Relievable Warm Water Cooling for Energy-Efficient Edge Datacenters

Qiangyu Pei

Huazhong University of Science and
Technology
Wuhan, Hubei, China
peiqliangyu@hust.edu.cn

Shutong Chen

Huazhong University of Science and
Technology
Wuhan, Hubei, China
shutongcs@hust.edu.cn

Qixia Zhang

Huazhong University of Science and
Technology
Wuhan, Hubei, China
zhangqixia427@hust.edu.cn

Xinhui Zhu

Huazhong University of Science and
Technology
Wuhan, Hubei, China

Fangming Liu*

Huazhong University of Science and
Technology
Wuhan, Hubei, China
fmliu@hust.edu.cn

Ziyang Jia

Huazhong University of Science and
Technology
Wuhan, Hubei, China

Yishuo Wang

Huazhong University of Science and
Technology
Wuhan, Hubei, China

Yongjie Yuan

Huazhong University of Science and
Technology
Wuhan, Hubei, China

ABSTRACT

As the computing frontier drifts to the edge, edge datacenters play a crucial role in supporting various real-time applications. Different from cloud datacenters, the requirements of *proximity to end-users*, *high density*, and *heterogeneity*, present new challenges to cool the edge datacenters efficiently. Although warm water cooling has become a promising cooling technique for this infrastructure, the one-size-fits-all cooling control would lower the cooling efficiency considerably because of the severe thermal imbalance across servers, hardware, and even inside one hardware component in an edge datacenter. In this work, we propose CoolEdge, a hotspot-relievable warm water cooling system for improving the cooling efficiency and saving costs of edge datacenters. Specifically, through the elaborate design of water circulations, CoolEdge can dynamically adjust the water temperature and flow rate for each heterogeneous hardware component to eliminate the hardware-level hotspots. By redesigning cold plates, CoolEdge can quickly disperse the chip-level hotspots without manual intervention. We further quantify the power saving achieved by the warm water

cooling theoretically, and propose a custom-designed cooling solution to decide an appropriate water temperature and flow rate periodically. Based on a hardware prototype and real-world traces from SURFsara, the evaluation results show that CoolEdge reduces the cooling energy by 81.81% and 71.92%, respectively, compared with conventional and state-of-the-art water cooling systems.

CCS CONCEPTS

• **Hardware** → Thermal issues; Enterprise level and data centers power issues; • **Computer systems organization** → Architectures.

KEYWORDS

edge datacenter energy, warm water cooling, hotspot relieving, vapor chamber

ACM Reference Format:

Qiangyu Pei, Shutong Chen, Qixia Zhang, Xinhui Zhu, Fangming Liu, Ziyang Jia, Yishuo Wang, and Yongjie Yuan. 2022. CoolEdge: Hotspot-Relievable Warm Water Cooling for Energy-Efficient Edge Datacenters. In *Proceedings of the 27th ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS '22)*, February 28 – March 4, 2022, Lausanne, Switzerland. ACM, New York, NY, USA, 16 pages. <https://doi.org/10.1145/3503222.3507713>

1 INTRODUCTION

Edge datacenters are emerging as a critical infrastructure for edge computing. To provide real-time services in close proximity to end-users, edge datacenters are widely distributed from commercial buildings to industrial complexes in forms of micro datacenters or server clusters. Gartner predicts around 75% of enterprise-generated data will be created and processed at the edge by 2025, though the value is only 10% in 2018 [30], bringing explosive growth in the number of edge datacenters. According to a report for edge computing [50], the edge datacenters will cost as high as \$100 billion

*Corresponding author: Fangming Liu (fmliu@hust.edu.cn). Q. Pei, S. Chen, Q. Zhang, X. Zhu, F. Liu, Z. Jia, Y. Wang, and Y. Yuan are with the National Engineering Research Center for Big Data Technology and System, the Services Computing Technology and System Lab, Cluster and Grid Computing Lab in the School of Computer Science and Technology, Huazhong University of Science and Technology, 1037 Luoyu Road, Wuhan 430074, China.

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.

ASPLOS '22, February 28 – March 4, 2022, Lausanne, Switzerland

© 2022 Association for Computing Machinery.

ACM ISBN 978-1-4503-9205-1/22/02...\$15.00

<https://doi.org/10.1145/3503222.3507713>

in information technology (IT) equipment capital expenditures in 2028. Although the power rating of an edge datacenter is only 10's to 100's of kW that is three orders of magnitude smaller than a cloud datacenter, such a growing number of edge datacenters will inevitably bring heavy energy burden. By 2028, the energy demand of edge datacenters will reach the same order of magnitude as that of the global datacenters in 2020 [42, 51].

Despite the small power capacity, the power density of an edge datacenter is generally much higher than that of a cloud datacenter due to the area restriction. Inspur proposes an edge server NE5260M5 whose depth is 65% of the standard depth in Open Compute Project [65]. For one thing, such a short depth makes the implementation more flexible and space-saving, so that the server can be mounted on a short rack or even on the wall. For another, the short rack and compact-aisle arrangement further increase the power density. For instance, a well-designed edge datacenter can work at 2.1 kW per square foot [94], which is one magnitude higher than the power density of a cloud datacenter.

Recently, Tencent Cloud has opened its first edge datacenter to provide real-time services of video processing, cloud gaming, smart healthcare, and so on [17]. Although traditional lightweight workloads like Web services are suitable to be scheduled on a central processing unit (CPU), the emerging computational edge workloads like deep learning inference, rely heavily on accelerators, such as graphics processing units (GPUs), tensor processing units, field-programmable gate arrays, smart network interface cards, etc. Hence, to support these diverse performance-critical edge applications, the edge server needs to comprise various high-powered heterogeneous hardware, leading to a high power provisioning to edge servers [12, 51, 108].

The specific requirements of edge datacenters, including *proximity to end-users*, *high density*, and *heterogeneity*, make existing cooling techniques inefficient or even impracticable. The free cooling technique requires specific low-temperature locations with free cooling sources like the cold outdoor air, contradicting the edge's demand for proximity to end-users. In addition, high density and heterogeneity further increase the difficulty of efficient cooling in edge datacenters. As the power density grows dramatically, air cooling would be no more suitable and even unsafe for the edge datacenters [9], especially when dealing with thermal imbalance of heterogeneous hardware [10, 95]. Hence, we argue that water cooling is a promising technique for edge datacenters and it also has great potential for saving energy.

Currently, warm water cooling (e.g., 40°C~50°C) emerges to reduce cooling costs by avoiding over-cooling servers running at a low utilization and allowing less or even no use of the chiller [41, 118]. Recent studies indicate that the cooling costs can be reduced by 40% through raising the water temperature [29], which is significant for the edge datacenters located in populated areas with a high electricity pricing. By employing warm water cooling, less cooling demand also enables a chiller with a smaller size and lower cooling capacity, saving about \$100~\$300 per kW of the capacity [28].

In spite of these promising advantages, state-of-the-art coarse-grained warm water cooling would be highly inefficient for edge datacenters due to the severe hotspot issue [41] at multiple levels. On the one hand, the imbalanced hardware utilization as well as different thermal specifications of heterogeneous hardware leads to




Figure 1: Water cooling architecture in an edge datacenter.

thermal imbalance across servers and hardware. To cool down even a small portion of hotspot hardware components for their safety, the inlet water for every hardware component should be chilled to a very low temperature synchronously. This over-provisioning strategy is exceedingly inefficient since the centralized chiller needs to consume extra energy to provide cold water to other non-hotspots at the same time [41, 100, 118]. One possible solution might be installing distributed chillers or pumps for each hardware component, but the high costs, additional space demands, and other technical problems make this infeasible [19, 41]. On the other hand, the thermal imbalance also exists inside a hardware component because of different thermal specifications and imbalanced utilization of internal units. As shown in later Fig. 2 and Sec. 2.2, the temperature difference inside a CPU core can be over 20°C, and the value exceeds 30°C inside a GPU. Concerning the temperature measurement granularity, undetected localized hotspots will not only result in performance degradation and a shorter lifespan, but also increase the cooling costs by over-cooling other non-hotspot units [20, 69, 71]. Thus, if the above multiple-level hotspots can be dispersed effortlessly, the cooling efficiency can be further improved while ensuring the safety.

In summary, conventional cold water cooling wastes vast amounts of unnecessary energy in cooling many low-utilization servers, while coarse-grained warm water cooling raises the hotspot issue. To make the best of their advantages while avoiding these negative impacts, we propose CoolEdge, a fine-grained warm water cooling system for relieving multiple-level hotspots and saving cooling energy of high-density and heterogeneous edge datacenters. Specifically, we make the following contributions:

- We summarize the new challenges to efficiently cool the high-density and heterogeneous edge datacenters, and argue that a solution to the hotspot issue is extremely urgent due to the rapid growth of edge computing.
- We put forward a hotspot-relievable warm water cooling architecture with two major innovations. Specifically, through fine-grained cooling control under well-designed water circulations, hardware-level hotspots can be eliminated with high efficiency; by means of our newly developed cold plates with vapor chambers, chip-level hotspots can be dispersed without manual intervention.

Table 1: Thermal specifications of some IT hardware components

Hardware type	Intel Xeon E5-2680 v4 CPU [35]	Nvidia GeForce RTX 2080 Ti GPU [75, 88]	Nvidia A100 80GB PCIe GPU [64, 76]	DRAM [10]	Samsung 983 DCT SSD [78]
MOT (°C)	86	89	100	85	70
TDP (W)	120	250	300	Typically ≤ 10	Read: 8.7, Write: 10.6

- To the best of our knowledge, we are the first to theoretically quantify energy savings achieved by warm water cooling. Based on the quantification, we propose a custom-designed cooling solution using the Heat Dissipation Oriented (HDO) or Chiller Power Oriented (CPO) strategy to provide cooling setting adaptation to heterogeneous hardware while improving the cooling efficiency.
- We build a hardware prototype to validate the practicability of CoolEdge, and conduct datacenter-level simulations to show the remarkable performance of CoolEdge in balancing multiple-level hotspots and saving cooling costs. The evaluation results demonstrate that compared with two baselines, CoolEdge reduces 81.81% and 71.92% of the cooling energy, respectively. A cost saving analysis estimates that CoolEdge can save up to \$3,513,600 yearly in a city.

2 BACKGROUND AND MOTIVATION

In this section, we first recap existing cooling techniques and discuss their limitations of achieving high cooling efficiency in edge datacenters. Then, we disclose the hotspot issue from both the hardware level and the chip level, and explain our motivation to design a new cooling architecture in the end.

2.1 State-of-the-Art Cooling Techniques VS. Demands of Edge Datacenters

First of all, we present a brief synopsis of existing cooling techniques, and analyze their inefficiency from the perspective of three requirements of edge datacenters: *proximity to end-users*, *high density*, and *heterogeneity*.

Free cooling vs. proximity to end-users: Free cooling, a way of directly cooling servers by free cooling sources, e.g., the outdoor air and lake water, has been well studied in recent years [58]. A growing number of cloud datacenters are built in cold and dry areas, or near the sea or lake, with free coolers like the dry cooler for access to cold air or water. For example, Iceland has become one of the world's most cost-effective destinations for datacenters owing to its ideal weather [27], and Microsoft even built its datacenter under the sea [74]. However, in order to provide low-latency services, edge datacenters should be sited in proximity to end-users and widely distributed in cities, which can hardly meet the strict requirements of free cooling.

Air cooling vs. high density: The contradiction between the increasing demand for edge datacenters and the shortage of urban land forces servers to be stacked up at a higher density. This dramatically increases the cooling demand since it becomes trickier to

timely take away the heat, especially when the servers run at 100% utilization for sustained periods. To ensure high performance of edge applications, air cooling struggles to satisfy the strict cooling demand at such a high power density [73], because of its low heat conduction capacity and the difficulty in managing the airflow efficiently when the rack and aisle are increasingly compact. Although there are some techniques using elaborate engineering on the air-aisle arrangement, like the circular pattern of racks designed by Vapor IO [93] and the high-density cooling proposed by Intel [62], these techniques show poor performance in efficiency, scalability, and/or adaptivity for edge datacenters.


Water cooling vs. heterogeneity: Water cooling emerges as an energy-efficient paradigm for datacenters. As water has a higher density and greater thermal capacity per unit volume than air, water cooling supports a higher power density [79]. Currently, cloud datacenters mainly use direct-to-chip cooling [95]. As shown in Fig. 1, a cold plate with water flowing inside is directly pressed on the surface of a hardware component to absorb the heat. After absorbing heat from different hardware components on different branches, the water gathers together and then is cooled to some temperature by the centralized chiller (also by the cooling tower in some large-scale cloud datacenters). For the safety of hardware components, this temperature is usually set to be low enough so as to cool down some high-utilization hardware components with a high temperature. In this coarse-grained water cooling system, however, since different hardware components share the same inlet water temperature and flow rate without considering the specific cooling demand of each individual hardware component, a lot of cooling energy is inevitably wasted, showing excessively low efficiency.

Nowadays, some cloud providers propose the warm water cooling technique to reduce cooling costs further. However, coarse-grained warm water cooling suffers from a severe hotspot issue, which we will expatiate in the following subsection.

2.2 The Hotspot Issue in Edge Datacenters

Compared with cloud datacenters, the hotspot issue becomes more severe in edge datacenters due to the requirements of high density and heterogeneity, along with skewed hardware utilization of edge workloads [108] and non-ideal ambient conditions of the edge. Typically, there are two kinds of thermal imbalance in edge datacenters, i.e., at the hardware level and the chip level.

The hotspot issue at the hardware level: Previous works [41, 55] have shown the hotspot issue exists among homogeneous hardware (e.g., CPUs or dynamic random-access memories (DRAMs))

**Figure 2: Temperature distribution inside a CPU core.**




Figure 3: Temperature variation of heterogeneous hardware components.

of the same type). For heterogeneous hardware, the thermal imbalance becomes more significant due to their divergent thermal specifications and dynamic characteristics. Table 1 illustrates their thermal specifications of Maximum Operating Temperature (MOT) and Thermal Design Power (TDP) [106]. As we can see, there are vast differences in both MOT and TDP from one hardware type to another, especially between the computing hardware and the memory or storage hardware. We also evaluate the dynamic characteristics of heterogeneous hardware components when changing their load levels. As plotted in Fig. 3¹, these hardware components show different operating temperatures and temperature variation rates in the same status. Usually, the operating temperature of computing hardware is above 40°C, while the operating temperature of memory hardware is lower than 40°C. When changing their statuses, the temperature of computing hardware goes up/down significantly faster than that of the memory, and reaches a stable level more quickly. Hence, it is essential to design a custom-designed solution for heterogeneous hardware, which will be introduced in Sec. 4.4.

The hotspot issue at the chip level: Considering the hardware type and workload characteristics, different internal units inside the hardware component may be at different utilization and power levels, which brings the hotspot issue at the chip level. We investigate this hotspot issue under four cases: among CPU cores, inside a CPU core, inside a GPU, and inside a DRAM, respectively.

(1) Hotspots among CPU cores: A CPU usually contains many processing units, i.e., cores. An existing work shows their average temperature imbalance can be over 7°C [69]. Our later simulation result in Fig. 7 also illustrates these hotspots.

(2) Hotspots inside a CPU core: A micro CPU core contains several parts from low-powered cache units to high-powered computing units. We use the HotSpot simulator [33] to acquire temperature distribution inside a CPU core, as presented in Fig. 2. In particular, the temperature difference between computing units and cache units can be over 20°C. When running integer workloads, there exist several hotspots especially in the integer register marked as IntReg in Fig. 2.

(3) Hotspots inside a GPU: A GPU consists of multiple units including computing units, memory units, etc. According to the measurement result of an AMD GPU for a stress test, inside the GPU,

the hotspots can reach above 100°C and the maximum temperature difference is over 30°C.

(4) Hotspots inside a DRAM: A DRAM is mainly composed of several DRAM chips and one buffer chip which have different rated temperatures [55]. According to the previous research, the temperature difference among DRAM chips is over 15°C while the value between DRAM chips and the buffer chip can reach more than 30°C [55, 117].

2.3 Why We Need a New Cooling Architecture to Address the Hotspot Issue for Edge Datacenters?

Many software-based solutions can be implemented to relieve hotspots in a cloud datacenter, including power throttling [44, 52, 71, 119], workload deferral [4], and workload balancing [10, 31, 32, 55, 69, 71]. However, it is usually necessary to consider the tradeoff between the performance guarantee and hotspot relieving. For example, avoiding hotspots by lowering hardware frequency is likely to degrade hardware performance. Also, since some mission-critical edge applications, such as smart traffic management [81], would have no deferrable workloads, the hotspots may unavoidably emerge constantly. As a result, it is necessary to propose a workload-agnostic solution to the hotspot issue for general cases at the edge. Since these software-based solutions require no hardware-level support and have access to higher level information, such as workload deadlines, they can be combined with this workload-agnostic solution to further reduce energy consumption through relieving hotspots.

Recently, Jiang et al. [41] propose a thermoelectric cooler-based (TEC-based) solution to address the hotspot issue in a homogeneous cloud datacenter with only CPUs. Specifically, the authors apply warm water to cool CPUs equally and integrate each CPU with a TEC to provide extra cooling capacity for hotspots. However, the TEC-based solution fails to meet the specific requirements of edge datacenters, i.e., high density and heterogeneity. On the one hand, it requires considerable modifications to the internal structure of servers, which is somewhat impractical for already high-density servers. In particular, given the high energy demand of the TEC, it will be disabled when the CPU becomes a non-hotspot. In this case, a copper plate of double the CPU's size and an additional cold plate are necessary for transferring heat. On the other hand, it cannot be extended to support heterogeneous hardware due to the following reasons. First, the required copper plate cannot cater to the physical layouts of various hardware types. Taking the GPU as an example, as several large capacitors are scattered around the computing units, there is no room for this copper plate. Second, the TECs cannot be applied to high-powered hardware due to their limited cooling capacity. The maximum heat load that an economical TEC can transfer effectively is usually less than 150 W [1], hardly meeting the cooling demand of high-powered hardware like GPUs whose TDP can reach 400 W [64].

In short, prior works are inefficient for latency-sensitive edge workloads or lack support for the heterogeneity of edge datacenters. Differing from them, we propose a hotspot-relievable water cooling system tailored to owner-operated edge datacenters with full consideration for heterogeneous hardware.

¹The details of hardware components are presented in Section 5.




Figure 4: Hotspot-relievable warm water cooling system.

3 SYSTEM ARCHITECTURE

In this section, we formally propose CoolEdge, a hotspot-relievable warm water cooling system for edge datacenters. We begin with the system overview and then elaborate on the design details.

3.1 System Overview

As shown in Fig. 4, each server consists of a series of heterogeneous hardware components, including CPU, GPU, DRAM, etc. There are three major parts in our proposed cooling system: Inner-and-Out Loop, Mini Loop, and controllers.

(1) Inner-and-Out Loop contains two water circulations, i.e., Inner Loop and Outer Loop, to cool the hardware components. In particular, Inner Loop is a hot water circulation directly recycling the “used” water after cooling the hardware; Outer Loop is a cold water circulation where the hot water is pumped to the chiller and turned into “refreshed” chilled water again. Unlike the conventional water cooling system, we use a valve to provide an appropriately customized inlet water temperature and flow rate for each hardware component, by mixing a certain amount of hot water from Inner Loop and cold water from Outer Loop.

(2) Mini Loop is a small vapor-fluid circulation inside a two-phase vapor chamber, which is creatively deployed on the cold plate to enhance thermal conductivity and reduce local hotspots inside the hardware component.

(3) Controllers include a Centralized Controller and multiple subControllers, i.e., one subController for each server. Specifically, based on the information collected in real time (e.g., hardware utilization and temperature) and the specific cooling strategy (introduced in Sec. 4), the Centralized Controller periodically decides on the best cooling setting, i.e., the inlet water temperature and flow rate for each hardware component. Then it sends the control command to each subController to adjust water temperature and flow rate accordingly.




Figure 5: CPU temperature under different cold plates, utilization, water temperature, and flow rate.

3.2 Inner-and-Out Loop: Hardware-Level Hotspot Elimination with Mixed Water

As illustrated in Sec. 2, both homogeneous and heterogeneous hardware components have different cooling demands at different times. In order to handle hotspots among different hardware components, we design two water cooling circulations, i.e., Inner Loop and Outer Loop, to achieve fine-grained and flexible cooling control in an edge datacenter. Specifically, we use a pulse-width modulation (PWM) controlled proportional solenoid valve [86] at the inlet of each hardware component. The water temperature and flow rate can be regulated at the desired values by mixing different amounts of hot and cold water based on each hardware component’s real-time cooling demand.

As plotted in Fig. 4, Inner Loop is a hot water circulation, gathering “used” water from the outlet of each hardware component to the water tank and pumping it to the inlet again. Since the hot water from Inner Loop cannot cool down some high-utilization hardware components, Outer Loop pumps hot water from the water tank to the chiller and then sends the chilled water to the inlet. Based on the control command from controllers, different amounts of hot water from Inner Loop and cold water from Outer Loop will be sent to the hardware component. As compared to merely sending the chilled water, the mix of both hot and cold water not only reduces the required amount of chilled water and saves cooling energy at the current time, but also increases natural heat dissipation (discussed in Sec. 4) which in turn saves cooling energy thereafter.

3.3 Mini Loop: Chip-Level Hotspot Dispersion with Two-Phase Vapor Chambers

To relieve hotspots inside a hardware component, we integrate a two-phase vapor chamber into the cold plate and realize vapor-fluid Mini Loop inside the chamber. As shown in Fig. 4, the cold plate is attached to the hardware component to transfer heat into the cooling water. Between the hardware component and the cold plate is the thermal paste, used to eliminate air and thus provide higher thermal conductivity. It is worth noting that the vapor chamber is typically standalone, and attached between a heat source and a cooling component to conduct heat directly. However, we find this is exceedingly inefficient in transferring heat from a hardware




Figure 6: Illustration of Mini Loop.

component to the cooling water inside an intact cold plate, due to the long thermal path and an extra layer of the thermal paste. Therefore, instead of directly attaching the vapor chamber to the bottom of the cold plate, we replace the commonly used cold plate's baseplate with the vapor chamber for achieving higher thermal conductivity as shown in later Fig. 12b.

To verify this observation, we further conduct some experiments to reveal the impact on CPU temperature of the commonly used cold plate (CC) and two cold plates leveraging aforementioned integration approaches. One approach is attaching the vapor chamber to the commonly used cold plate directly (VC-CC) and the other is replacing its baseplate with the vapor chamber (VC, our approach). The results are plotted in Fig. 5², where the horizontal line indicates MOT (i.e., 86°C) and CC, 30, 120 refers to using the commonly used cold plate under the inlet water temperature of 30°C and flow rate of 120 L/h. We can see that VC outperforms the others, especially when the CPU utilization and inlet water temperature get high. For selecting a cost-effective vapor chamber for the integration, we also test vapor chambers with different materials (copper and aluminum) and thicknesses under various load and cooling conditions (e.g., hardware utilization, water temperature, flow rate, etc.). Now we introduce the physical structure, working principle, and attractive characteristics of the vapor chamber as follows.

Physical structure: As shown in Fig. 6a, the vapor chamber consists of a sealed vacuum vessel and an internal wick structure. The outside of the vacuum vessel is typically made of copper or aluminum to achieve high thermal conductivity, while the wick structure contains a small amount of working fluid in equilibrium with its vapor to transfer heat from one side of the chamber to the other [70]. Another performance metric influencing the thermal conductivity is the filling ratio, defined as the ratio of the volume of the working fluid out of the total volume of the vapor chamber [89]. Usually, the value is set at 20%~45% [107].

Working principle: As shown in Fig. 6b, the vapor chamber includes an evaporator and a condenser. The evaporator consists of a wick structure, where the working fluid flows. The fluid absorbs heat from hardware components, and immediately vaporizes and rises to the condenser driven by the pressure difference. Once arriving at the condenser, the vapor condenses again and releases the latent heat of condensation to the cooling water. The condensed fluid finally returns to the evaporator by the capillary action of the wick structure and also by gravity [70]. Through this heat circulation, the hotspots inside each hardware component can be dispersed




Figure 7: (a) The die shot of Intel Xeon E5 CPU [103], and thermogram of the die cooled by the (b) commonly used cold plate and (c) vapor chamber-based cold plate.

automatically and a relatively uniform temperature distribution can be realized.

Attractive characteristics: Compared with the conventional single-phase cooling system, there are several benefits from the hotspot-relievable two-phase vapor chambers. First, because of the huge latent heat of vaporization, the vapor chamber has higher thermal conductivity than the commonly used cold plate and thus reduces the hardware temperature to a lower level. To this extent, the inlet warm water temperature can be raised further to save more cooling energy. Second, as the thermal conductivity of vapor chambers grows with the power density of the heat source, more heat can be absorbed from hotspot areas inside the hardware component [15, 102]. In all, the vapor chamber not only increases the general thermal conductivity, but also smooths the temperature distribution of the heat source in an automatic manner. To verify the benefits, we use the 3D-ICE thermal simulator [84] to obtain the die³ temperature distribution under the commonly used cold plate and vapor chamber-based cold plate. As plotted in Fig. 7, by integrating the vapor chamber into the cold plate, we can realize a much lower and more uniform temperature distribution on the die.

In addition, the two-phase vapor chamber also has many unique properties desirable to edge datacenters:

- **Space-saving:** The vapor chamber is usually as thin as a copper baseplate of the commonly used cold plate (e.g., 2~3 mm) but much lighter.
- **Cheap:** The vapor chamber can be bought for about \$5 from *Alibaba.com* [7]. Besides, about \$1 for the copper baseplate can be saved.
- **Reliable:** The Mean Time Before Failure (MTBF) of the vapor chamber is 80,000 hours (i.e., about 10 years) [70], larger than the hardware lifespan.
- **Environment-friendly:** The vapor chamber consumes no power, and all the raw materials (e.g., purified water as the working fluid) are environment-friendly, significant for today's green datacenter infrastructure.

3.4 Controllers

In CoolEdge, there is a Centralized Controller and each server is directly controlled by an independent subController.

Centralized Controller: The Centralized Controller is composed of two parts, i.e., a monitor and a scheduler. In each adaptation

²The details of the CPU are presented in Section 5.

³A die is a small block of semiconducting material on which a given functional circuit is fabricated [105].

period, the monitor first collects each hardware component's temperature and utilization/power information from subControllers. Then, the scheduler decides the best inlet water temperature and flow rate for each hardware component based on the collected information and the cooling strategy (introduced in Sec. 4). Finally, the scheduler sends the cooling control commands back to each subController.

subController: Each subController periodically collects the temperatures of CPUs and DRAMs with the `lm_sensors` tool, utilization of CPUs and DRAMs by the `/proc` filesystem, and temperatures and powers of GPUs with the `nvidia-smi` tool, and then sends them to the Centralized Controller. Once receiving the control command from the Centralized Controller, each subController sends the control signal to the valves installed on the server. By connecting the valves to a 4-pin power connector on the motherboard, it is convenient to tune each valve by the PWM signal [86], and the relationship between the water flow rate and duty cycle is presented in [86, 91]. As the response time of such a valve is less than 1 s [87], real-time control can be achieved. Besides, such a valve's power consumption is only several Watts [67], which is negligible compared with the IT and cooling equipment.

4 FINE-GRAINED COOLING SOLUTION

In this section, we first theoretically quantify the power saving achieved by the warm water cooling. According to the quantification, we propose two warm water cooling strategies to decide on the cooling setting. At last, we present a custom-designed cooling solution for heterogeneous hardware.

4.1 Key Proposition of Warm Water Cooling

The state-of-the-art literature on warm water cooling states that increasing the inlet water temperature has great potential for saving cooling energy [29, 41, 118]. As the warm water temperature is higher than the ambient temperature, there exist significant natural heat dissipation phenomena in pipes and tanks, lowering the cooling energy consumption of the chiller compared with the cold water cooling. In this section, we provide the first theoretical analysis on the efficiency of warm water cooling from the aspect of natural heat dissipation. In the following part, we take the heat dissipation in pipes as the representative, since the tank can be viewed as a wider pipe and analyzed in a similar way. All the theoretical derivations and detailed discussions are provided in Appendix A.

PROPOSITION 4.1. *The natural heat dissipation efficiency depends on (1) ΔT : the temperature difference between the cooling water in the pipe and the outer air, (2) v : the water flow rate, (3) h : the convective heat transfer coefficient of the air, and (4) ξ and μ : parameters related to physical characteristics of the pipe and water (e.g., the pipe's radius and the density of water), respectively. Based on Fourier's law of heat conduction [18] and Newton's law of cooling [109], the dissipated heat P (in Watts) through the pipe can be calculated by:*

$$P = \xi v \Delta T (1 - \exp(-\mu h / v)), \quad (1)$$

Eq. (1) indicates that compared with the water flow rate, increasing the water temperature contributes significantly to the heat dissipation under the same ambient condition. The convective heat transfer coefficient of the air h represents the thermal resistance




Figure 8: The effect of the inlet water temperature, v , and h on P based on the simulation and estimation.

of a relatively stagnant layer of air between a pipe surface and the air medium [26]. Higher air velocity increases h and thus the dissipated heat. The convective heat transfer coefficient of the free air is usually 2.5~25 W/m²°C, while in the forced convection environment, e.g., with fans blowing around, the value can be up to 10~500 W/m²°C [47].

Using Fluent software [8], we simulate the process that the water flows through a 1-meter copper pipe when the ambient air temperature is stable at 20°C. Fig. 8 shows the effect of the inlet water temperature on the heat dissipation under varying v and h from both the simulation results and estimation results based on Eq. (1). Note that when the inlet water temperature and h are high enough, the amount of dissipated heat can be equal to a CPU's TDP (as listed in Table 1), showing the great potential of warm water cooling to save cooling energy. Although we find that the estimation of heat dissipation is generally accurate, there is still an estimation error of at most 8.4% when the inlet water temperature and h are high. The main reason is that Eq. (1) cannot accurately describe the temperature of the air and the water near the pipe wall [61]. Here we use an attenuation factor β to make up the estimation error, whose detailed expression is introduced in Appendix A.

Proposition 4.1 analyzes the key factors affecting the water cooling efficiency. We find that from the aspect of natural heat dissipation, a higher temperature of the cooling water shows a remarkable effect on increasing cooling efficiency. To sum up, the warm water cooling saves more cooling energy as more heat can be dissipated in a natural way, and the amount of dissipated heat is affected by several factors.

4.2 Heat Dissipation Oriented Strategy

Based on Proposition 4.1, we find that the inlet water temperature and its flow rate are two tunable factors that determine the natural dissipated heat P presented in Eq. (1) and thus the cooling efficiency. To make full use of natural heat dissipation, the cooling system should cool each hardware component with the best cooling setting of the inlet water temperature and flow rate. Let T_{hot} and T_{cold} represent the temperatures of inlet hot water and inlet cold water, respectively. For the hardware component i whose power is P_i , the mixed inlet warm water temperature is denoted by $T_{warm,i}$. As discussed in Sec. 3, the temperature and flow rate of both inlet hot water from Inner Loop and cold water from Outer Loop together determine the temperature of the inlet warm water. We use $v_{hot,i}$ and $v_{cold,i}$ to represent the flow rate of the inlet hot water and

inlet cold water, respectively. The temperature and flow rate of the inlet warm water for the hardware component i is given by $T_{warm,i} = (v_{hot,i}T_{hot} + v_{cold,i}T_{cold}) / (v_{hot,i} + v_{cold,i})$ and $v_{warm,i} = v_{hot,i} + v_{cold,i}$, respectively. Based on the law of conservation of energy [25], the temperature of the hot water from the i -th outlet $T_{out,i}$ can be calculated by:

$$T_{out,i} = T_{warm,i} + \frac{P_i}{\rho v_{warm,i}}, \quad (2)$$

where ρ and c represent the density and specific heat capacity of the inlet water, respectively.

Based on Eq. (1), and $T_{out,i}$ and $v_{warm,i}$ for each hardware component, the dissipated heat (in Joule) of the pipe $E_{pipe} = P_{pipe}t = \sum_i P_{pipe,i}t$ and the tank $E_{tank} = P_{tank}t$ can be calculated accordingly, where t is the time slot. Then, we can obtain the best cooling setting of the water temperature and flow rate for each hardware component to maximize the amount of dissipated heat R_{HDO} :

$$R_{HDO} = \frac{E_{pipe} + E_{tank}}{COP_c} - E_{pump}, \quad (3)$$

where COP_c is the coefficient of performance (COP) [104] of the chiller, and E_{pump} is the power consumption of the pump, whose expression is presented in Sec. 5.2. Since we maximize natural heat dissipation here, we call this warm water cooling strategy as **Heat Dissipation Oriented (HDO)** strategy.

4.3 Chiller Power Oriented Strategy

Although the HDO strategy provides the precise warm water cooling setting, it could be impractical to achieve in reality. For instance, some parameters in Eq. (1), such as the size of pipes and h , are inaccessible or varying dynamically. Moreover, the heavy computational overhead would also affect the edge's performance. To ensure the practicability of the cooling management, we present a simple and more general strategy.

It is obvious that the lower warm water temperature and higher flow rate we set, the larger volume of chilled water will be used to cool the hardware component within a given time slot, and hence the temperature of outlet hot water is lower. In this case, according to Proposition 4.1, the amount of dissipated heat would fall, reducing the overall cooling efficiency. As a result, we regard the cooling setting with the least chilled water provision, i.e., the least energy consumption of the chiller as the best choice and call this **Chiller Power Oriented (CPO)** strategy. Here, we can obtain the best cooling setting of water temperature and flow rate for each hardware component to minimize the chilled water provision R_{CPO} :

$$R_{CPO} = \sum_i v_{cold,i} = \sum_i v_{warm,i} \cdot \frac{T_{hot} - T_{warm,i}}{T_{hot} - T_{cold}}. \quad (4)$$

4.4 The Cooling Control Solution

After introducing the two strategies for improving the efficiency of warm water cooling, we present a custom-designed cooling control solution for heterogeneous hardware here. Since the computing hardware (e.g., CPU and GPU) shows different thermal specifications and dynamic characteristics from the memory hardware (e.g.,




Figure 9: Measurement results of CPU temperature with different inlet water temperature, water flow rate, and CPU utilization.




Figure 10: Measurement results of GPU temperature with different inlet water temperature, water flow rate, and GPU power.

DRAM) as discussed in Sec. 2.2, we quantify their thermal profiles at first.

According to our findings in Sec. 2.2, the inlet water temperature and flow rate directly impact the instantaneous temperature of computing hardware, so we only consider their real-time utilization/power to make cooling decisions. Measurement results of the CPU and GPU temperature with different inlet water temperature, water flow rate, and hardware utilization/power are plotted in Fig. 9 and Fig. 10, respectively.

For DRAMs, since their power is excessively low as shown in Table 1, the generated heat can be timely taken away whatever the inlet water flow rate is, and the inlet water temperature directly influences the DRAM temperature variation rate instead of the instantaneous DRAM temperature. Here, we define the temperature variation rate as the reciprocal of the time spent to raise or reduce the hardware temperature by 1°C under specific inlet water temperature, hardware temperature, and hardware utilization. Next, we measure the relationships between the DRAM temperature variation rate and these three variables as shown in Fig. 11. The positive and negative values of the temperature variation rate mean an increase and decrease in the hardware temperature, respectively. The absolute value of the temperature variation rate grows exponentially with the temperature deviation from the stable temperature, making the temperature quickly reach near the stable temperature.

Based on the above hardware profiles and the received information from subControllers, the scheduler decides the best cooling




Figure 11: DRAM temperature variation rate (the solid points represent the stable temperatures).

setting for each heterogeneous hardware component periodically with a custom-designed solution as follows:

- **Computing hardware:** For each hardware component, the scheduler first selects possible choices of the inlet water temperature and flow rate from the profiles (e.g., the measurement results in Fig. 9 and Fig. 10) which guarantee the hardware temperature lower than the safe operating temperature. Leveraging the HDO or CPO strategy, the scheduler can obtain the best warm water cooling setting by maximizing Eq. (3) or minimizing Eq. (4).
- **Memory hardware:** Since only the inlet water temperature is tunable for DRAMs, both HDO and CPO strategies always choose the same cooling setting—maximizing the inlet water temperature. Thus, for each inlet water temperature, the scheduler calculates the DRAM temperature at the next adaptation period based on the measurement results in Fig. 11 and chooses the best cooling setting of the water temperature with the safe operating temperature guarantee.

According to our single-threaded testing on a server with an Intel Xeon E5-2697 v4 CPU, it takes about 20 ms for the HDO or CPO strategy to obtain the best cooling setting for each hardware component. Leveraging parallel computing, we can easily scale up the computing speed when the number of hardware components increases greatly. After determining all the cooling settings, the scheduler will send cooling control commands to each subController. Then the subController will send the control signal to each valve at once to adjust the amounts of hot water and cold water sent to each hardware component accordingly.

5 EVALUATION

In this section, we first introduce our well-established hardware prototype. Based on the collected hardware profiles, we conduct extensive simulations with real-world traces to evaluate CoolEdge. In the end, we present a cost saving analysis and briefly discuss the system reliability.

5.1 Hardware Prototype

To verify the practicability of CoolEdge and collecting thermal profiles that are presented in Fig. 9, Fig. 10, and Fig. 11, we build up a hotspot-relievable warm water cooling prototype in a Dell Precision Tower 7910 Workstation [22], as illustrated in Fig. 12a. The cooling




Figure 12: Hardware prototype of CoolEdge.

part contains Inner Loop and Outer Loop. Inner Loop is composed of a water tank, a pump, a flowmeter for monitoring the water flow rate, and a thermosensor for monitoring the temperature of inlet warm water from the water tank. Outer Loop consists of a pump, a flowmeter, and a chiller for providing chilled cold water. After mixing the hot water and cold water via the valves, customized warm water is sent to each hardware component. Fig. 12b shows the items in the server, including an Intel Xeon E5-2680 v4 CPU, an Nvidia GeForce RTX 2080 Ti GPU, and four G.Skill DDR4 DRAMs. Note that to make the illustration clear, we do not connect water pipes to all the hardware components, and take the CPU as a representative.

5.2 Evaluation Setup

As real-world edge datacenters and edge traces are not accessible at the moment, to evaluate the performance of CoolEdge from the industrial perspective, we simulate the energy usage using cluster traces from SURFsara [48]. The traces include the utilization or power information of 341 CPUs, 341 DRAMs, and 57 GPUs for about 3 months. To make the trace conform to the high-density and high-utilization characteristics of edge servers, we regard each server as a 2-way server and select the top-20 traces of CPUs, GPUs, and DRAMs separately based on their highest utilization or power within the selected 10 hours. That is, there are 10 servers considered in the evaluation, each equipped with 2 CPUs, 2 GPUs, and 2 DRAMs. Later Fig. 16 shows their utilization during the 10 hours. In the datacenter-level simulation, all the thermal profiles of CPU, GPU, and DRAM required by the Centralized Controller are collected from the hardware prototype. This would not influence the feasibility and availability of the profiles considering the differences between servers and workstations are irrelevant to water cooling efficiency [63]. We also take into account necessary physical infrastructure of edge datacenters in the simulation, including the length of pipes, and the sharing of one centralized chiller and two pumps in Inner Loop and Outer Loop, respectively. To validate the effectiveness of the fine-grained cooling solution, we integrate vapor chambers for all the baselines in the simulation. Then, we present a brief analysis of energy savings and cost savings from vapor chambers in Sec. 5.4. Since air cooling cannot meet the cooling demands of edge datacenters, we only consider two water cooling baseline strategies as follows:

Table 2: Safe operating temperature of hardware

Hardware type	CPU	GPU	DRAM
Safe operating temperature (°C)	77	76	43

Table 3: Other parameters

Parameter	h	COP_c	Ambient temperature
Value	10 W/m ² °C [45]	3.6 [60]	20°C

- Conventional coarse-grained water cooling baseline (*Coarse-grained*): For this baseline, we set the global water temperature and flow rate according to the highest cooling demand of all the hardware components.
- State-of-the-art TEC baseline (*TEC*): Jiang et al. [41] equip each CPU with a TEC in a datacenter. Since this solution is infeasible in a heterogeneous edge datacenter as discussed in Sec. 2.3, we only use the CPU trace in the comparison, and suppose there are two CPUs in each server. We use E_{TEC} to indicate the energy consumption of TECs.



We also use fans to maintain the ambient temperature and improve the natural heat dissipation by increasing h . Considering that their energy consumption is much lower than the water cooling equipment, during the simulation, the total energy consumption E_{total} is only the summation of the energy consumption of the centralized chiller $E_{chiller}$ and two pumps E_{pump} . The former can be calculated by $E_{chiller} = \frac{\rho c t (T_{hot} - T_{cold}) \sum_i v_{cold,i}}{COP_c}$, and the latter can be approximately calculated by $E_{pump} = 1.3674vt$ according to our experimental results, where $v = \sum_i v_{warm,i}$ is the total flow rate of every water branch, and t is the time slot. As prolonged operation at near MOT may degrade performance and shorten hardware lifespan [11, 20], for the CPU, we set the safe operating temperature as 90% of its MOT, while for the GPU, the value is set as 85% of its MOT because of its higher temperature variation rate as shown in Fig. 3 and thus higher probability of overheating. As for the DRAM, the value is set at a lower level to maintain high reliability [55]. All the safe operating temperatures and other parameters are listed in Table 2 and Table 3, respectively.

5.3 Evaluation Results

We analyze the simulation results of CoolEdge (applying the HDO or CPO strategy), and two baseline strategies of *Coarse-grained* and *TEC* from several aspects as follows.

Total cooling energy consumption: As shown in Fig. 13a, the HDO and CPO strategies reduce the cooling energy by 49.57% and 60.08%, respectively, as compared to the *Coarse-grained* strategy. HDO and CPO strategies also reduce the partial power usage effectiveness (pPUE)⁴ from 1.17 achieved by the *Coarse-grained* strategy to 1.09 and 1.07, respectively. It should be noted that although some recent cloud datacenters have achieved a low PUE thanks to the ideal climate and/or well-designed cooling techniques, the PUE of

⁴The pPUE is defined as (IT hardware energy + cooling energy) / IT hardware energy [41].

**Figure 13: Total energy consumption.****Figure 14: Cooling energy consumption patterns.**

edge datacenters is typically close to 2 [98] owing to the critical challenges introduced in Sec. 2. For the results under the CPU trace, as plotted in Fig. 13b, the proposed strategies reduce up to 81.81% and 71.92% of the cooling energy, respectively, as compared with the *Coarse-grained* strategy and *TEC* strategy. The pPUE of the HDO, CPO, *TEC*, and *Coarse-grained* strategies are 1.05, 1.02, 1.10, and 1.06, respectively. In both figures, the CPO strategy consumes less cooling energy than the HDO strategy. On the one hand, the CPO strategy has a slightly higher E_{pump} since it does not optimize E_{pump} . On the other hand, the CPO strategy reduces $E_{chiller}$ considerably as it aims at minimizing the amount of required chilled water. Note that in the following analysis, we consider the heterogeneous system under all the traces and compare the proposed strategies with only the *Coarse-grained* strategy.

Cooling energy consumption patterns: Fig. 14 depicts the total IT hardware power and the cooling energy consumption patterns of HDO, CPO, and *Coarse-grained* strategies. As we can see,




Figure 15: CDFs of the maximum hardware temperature (the gray vertical lines represent safe operating temperatures).




Figure 16: Utilization patterns of CPUs, GPUs, and DRAMs.

Table 4: Cost saving calculation (unit: \$/(server×year))

Description	TEC		CoolEdge	
ExCapEx	TEC	0.4	Valve	18.0
	Copper plate	0.2	Vapor chamber	1.0
	Additional cold plate	2.1	Copper baseplate	-0.2
ChiSav	1.25		5.32	
EnerSav	2.68		35.44	
CoSav	1.23		21.96	

both $E_{chiller}$ and E_{pump} increase synchronously when the IT hardware power boosts, such as time = 1,800 and time = 2,200, because more chilled water has to be pumped to cool hardware components with higher power consumption and thus higher temperature. Compared with the HDC and CPO strategies, the *Coarse-grained* strategy incurs much more cooling energy fluctuation since it needs to consume substantial extra energy to eliminate even one hotspot in each period. The higher peak cooling demand usually means a higher capital expenditure of the chiller, which will be discussed in later Sec. 5.4.

Hotspot elimination for heterogeneous hardware: The cumulative distribution functions (CDFs) of the maximum temperatures in each period of CPUs, GPUs, and DRAMs are plotted in Fig. 15 separately, and Fig. 16 plots their utilization patterns during the 10 hours. As we can see, the utilization of CPUs remains low almost all the time except for the No. 17 CPU, whose utilization exceeds 80% about half the time. By contrast, a large proportion

of GPUs and DRAMs run close at their maximum utilization frequently and thus there are several hotspots in each period. That is why the maximum temperatures of GPUs and DRAMs remain high almost all the time. In conclusion, as plotted in Fig. 15, although the maximum temperatures get higher by applying the proposed strategies as compared with the *Coarse-grained* strategy, no hardware component is overheated, which demonstrates the effectiveness of CoolEdge from the aspect of hotspot elimination.

5.4 Cost Saving Analysis

Here, we estimate the cost savings from CoolEdge and the *TEC* strategy as compared with the *Coarse-grained* strategy. We consider extra capital expenditures (ExCapEx), capital expenditure savings of the chiller (ChiSav), and cooling energy savings (EnerSav) in the analysis. ExCapEx mainly depends on the additions and can be calculated according to their purchase prices and lifespans [5–7, 41, 70, 90]. Note that the price of valves is irrelevant to hardware components, and the price of vapor chambers is mainly determined by the component size. Hence, these hardware prototype-based calculation results can reflect the actual ExCapEx for servers. Since the price of the chiller is mainly determined by its cooling capacity [28], ChiSav can be calculated by the peak cooling energy consumption as depicted in Fig. 14. As for EnerSav, it can be calculated from Fig. 13, where the electricity price for industrial consumers is about 15 cents/kWh [85]. Ultimately, Cost Savings (CoSav) can be calculated by ChiSav + EnerSav – ExCapEx. Based on the experimental results shown in Fig. 5, as the CPU temperature is reduced by 4.36°C on average through integrating the vapor chambers, EnerSav and CoSav can be further improved by 4.5% and 3.5%, respectively [59]. All the calculation results are listed in Table 4. According to the estimation, for 2,000 small-scale edge datacenters (each equipped with 80 servers) in a city [16, 96, 99], the cost savings of the *TEC* strategy is about \$196,800/year, while CoolEdge can save up to \$3,513,600/year, 17.85× as much as the *TEC* strategy, showing the great potential to widely deploy CoolEdge. Besides, the *TEC* strategy is designed for homogeneous datacenters, while CoolEdge is tailored to heterogeneous edge datacenters, with good generalizability in supporting heterogeneity.

5.5 System Reliability

Since the control system plays a key role in ensuring the safety of the whole datacenter infrastructure, CoolEdge also provides several measures to guarantee the system reliability in case of a control system failure. For the hardware-level failure, i.e., a valve

breaking down, the Centralized Controller will immediately migrate the workload running on this broken-valve-related hardware component to another idle hardware component and send alarms to datacenter operators for a repair or replacement. For achieving higher reliability of valves, there could also be a bypass water branch with a simple on-off valve at an additional capital expenditure. The subController will control this valve to directly send cold water in case of a proportional valve failure. At the software level, there are three kinds of possible failure, i.e., the unresponsive Centralized Controller, unresponsive subController, and instant hardware temperature rise that may lead to a hardware failure. First of all, when the Centralized Controller cannot respond timely (e.g., network failure), the subController will set valves to directly provide cold water. Meanwhile, an idle or low-utilization subController will perform as a substitute until the Centralized Controller recovers. Secondly, when a local subController loses connections (e.g., OS malfunctions), the valves will be fully open by default and the Centralized Controller will notify the operators to repair. Thirdly, although the Centralized Controller sends cooling control command to each subController periodically, when the hardware temperature reaches the alarming safe operating temperature abruptly, the subController will ask for an instant decision from the Centralized Controller rather than waiting for the next adaptation period.

6 RELATED WORK

Warm water cooling: Warm water cooling has been utilized by many works to reduce cooling costs [3, 21, 39, 41, 46, 77, 118]. Jiang et al. [41] propose a fine-grained warm water cooling solution to improve cooling efficiency by eliminating hotspots with TECs. However, as discussed in Sec. 2.3, there exist several limitations when dealing with high density and heterogeneity of edge datacenters. By contrast, CoolEdge provides general cooling supports for heterogeneous hardware components. Considering the limited space inside high-density servers, our newly-designed cold plates can directly replace original ones and all valves can be easily installed outside servers. Zhu et al. [118] illustrate another benefit of warm water cooling that the waste heat can be not only used for district heating, but also turned into electricity with thermoelectric generators. To the best of our knowledge, we are the first to theoretically analyze and formulate the profit of warm water cooling, and adopt it into edge datacenters with strict requirements and critical challenges.

Immersion cooling: Nowadays, immersion cooling has emerged as an efficient water cooling technique to hold a higher power density [2, 38]. Sugon has exhibited its newly designed two-phase immersion cooling server at Supercomputing 2018 [43]. As compared with direct-to-chip cooling, immersion cooling can reduce PUE further, and eliminate the use of cold plates and the corresponding engineering costs to design a new cold plate when the physical structure of new hardware components changes. However, there are still some technical problems of immersion cooling that may prevent its adoption in edge datacenters. Firstly, due to the strict requirements of fluid, such as high thermal conductivity, low electrical conductivity, fixed boiling point, etc., the fluid cost is much higher than direct-to-chip cooling [68]. Secondly, highly sealed sleds and stable gas pressure are necessary to ensure safety, which increases the tank cost. Thirdly, immersion cooling usually

occupies more space since racks are usually placed horizontally rather than vertically [97], worsening the problem of land scarcity at the edge. Last but not least, immersion cooling may be over qualified in most cases at present [97], where the cheaper direct-to-chip cooling probably fits better.

Phase change solutions: There are several works [36, 80, 82, 83, 112] utilizing phase change materials (PCMs) to handle cooling mismatch and/or achieve higher thermal conductivity. Skach et al. [82] shave the peak cooling load in an air-cooled datacenter with commercial paraffin that absorbs heat by melting at daytime and releases heat by solidification at nighttime. They further advance a job placement strategy for active control of PCMs' melting temperature [83]. Seuret et al. [80] are the first to present a two-phase thermosyphon entity for high heat flux hardware. However, neither the paraffin nor the thermosyphon can be easily integrated into high-density edge datacenters because of space limitation.

Vapor chamber solutions: Several studies propose to use vapor chambers to cool heat sources like processors. Tsai et al. [92] and Liu et al. [56] focus on the structures (e.g., flow mechanism) and properties (e.g., thermal resistance) of vapor chambers, and the influencing factors in heat conduction like installation orientations. Parhizi et al. [66] and Yuan et al. [110] evaluate the performance of vapor chambers by developing simulation models. Different from these studies, our work digs into thermal specifications of real hardware components used in datacenters. To our best knowledge, we are the first to integrate the vapor chamber with the commonly used cold plate to dissipate heat from server hardware components in a practical and cost-effective manner.

Power and thermal management in datacenters: Due to the rapid growth of cloud and edge computing, Internet of things, deep learning, etc., the power consumption of IT and cooling equipment in datacenters increases dramatically in recent years [23, 37]. To improve the overall efficiency, many works pay close attention to power and thermal management in datacenters, such as workload management [34, 40, 49, 53, 57, 72, 113, 115, 116], hotspot elimination [41, 55], underprovisioning [58, 111], heat harvesting [14, 54, 118], and demand response [13, 114]. In this section, we summarize some closely related works on the hotspot issue and/or cooling efficiency. Liu et al. [55] notice that the DRAM temperature can rise to alarming 95°C in a high-performance computing IT system. To avoid throttling and maintain high performance, they propose three schemes to reduce peak temperature and temperature variation in an air-cooled datacenter. Zhou et al. [113] develop a power management framework to save CPU power with the dynamic voltage and frequency scaling technique. Instead of a single hardware type, other works focus on the whole datacenter infrastructure including the cooling system. Intel [34] devises a scheme of improving the ambient temperature in a low-power-density air-cooled datacenter to increase cooling efficiency. Ran et al. [72] propose a deep reinforcement learning based framework to schedule CPU jobs and adjust airflow for saving cooling energy while reducing hotspots. However, these works mainly focus on either power and thermal management of IT hardware components, or system-level cooling control in a homogeneous cloud datacenter, while our work enables component-level cooling control tailored to high-density and heterogeneous edge datacenters.

7 CONCLUSION AND FUTURE WORK

In this paper, we propose CoolEdge, a hotspot-relievable warm water cooling system for owner-operated edge datacenters. CoolEdge integrates the vapor chamber into the commonly used cold plate to disperse chip-level hotspots without manual intervention. Through the elaborate design of the water circulation system with a fine-grained cooling solution, CoolEdge can eliminate hardware-level hotspots efficiently. The evaluation results indicate that CoolEdge achieves the best cooling efficiency without the risk of overheating any hardware components. Compared with the conventional and state-of-the-art water cooling systems, CoolEdge saves the cooling energy by 81.81% and 71.92%, respectively. For a city with 2,000 edge datacenters, CoolEdge can save up to \$3,513,600 yearly based on the estimation. It is worth noting that CoolEdge is also applicable to cloud datacenters although their requirements are less stringent.

Our future work includes custom-designed vapor chambers, and self-adaptive and power-free “valves” to further increase the flexibility of CoolEdge. In particular, through comprehensive consideration of the size, shape, filling ratio, etc., the vapor chambers can cater to more hardware types. By coupling two micro pipes of inlet cold water and outlet hot water with negative thermal expansion materials [101], the “valves” can adapt to varied cooling demands automatically without any power supply. To further demonstrate the advantages of CoolEdge, we will also keep track of and leverage latest released traces from public edge datacenters. Besides datacenter cooling, we believe our study on the hotspot issue discussed in Sec. 2.2 provides insights into power and thermal management at multiple levels. These insights may further influence the integrated circuit design and datacenter infrastructure design, including both IT and non-IT systems. We argue that it is no longer efficient to neglect the cooling system’s impacts on the overall efficiency of datacenters. As carbon neutrality of datacenters has become a critical issue recently, our work also intends to motivate more green computing research on datacenters, especially the emerging edge datacenters.

ACKNOWLEDGMENTS

We thanks the anonymous reviewers and our shepherd Ricardo Bianchini for their valuable feedback. This work was funded in part by National Key Research & Development (R&D) Plan under grant 2017YFB1001703, by NSFC under grant 61722206 and 61761136014, and by National Program for Support of Top-notch Young Professionals in National Program for Special Support of Eminent Professionals.

A DETAILED DISCUSSION ON NATURAL HEAT DISSIPATION

To derive Proposition 4.1, we consider a hollow-cylinder pipe of length L and thermal conductivity λ . r_i , r_o , T_i , and T_o denote the inner radius, outer radius, inner surface temperature, and outer surface temperature, respectively. For steady-state heat conduction, the problem can be formulated as $\frac{L}{r} \frac{d}{dr} (r \frac{dT}{dr}) = 0$ [24]. By applying the boundary conditions $T|_{r=r_i} = T_i$ and $T|_{r=r_o} = T_o$, temperature distribution along the radial direction is given by

$T_r = T_i - \frac{T_i - T_o}{\ln(\frac{r_o}{r_i})} \ln(\frac{r}{r_i})$. Based on Fourier’s law of heat conduction [18] and Newton’s law of cooling [109], the dissipated heat P (in Watts) through the pipe is expressed by:

$$P = (T_i - T_o) \left(\frac{\ln(\frac{r_o}{r_i})}{2\pi\lambda L} + \frac{1}{2\pi h r_o L} \right). \quad (5)$$

Eq. (5) assumes that the water temperature along the axis remains unchanged. To obtain a more accurate solution, we further adopt the infinitesimal calculus approach. Specifically, one water element with length dl is considered, whose initial temperature is T_i (i.e., the inlet water temperature of the pipe). Based on Eq. (5), the dissipated heat through the pipe in a time slot dt can be obtained by

$Q_{diss} = Pdt = (T_i - T_o) / (\frac{\ln(\frac{r_o}{r_i})}{2\pi\lambda dl} + \frac{1}{2\pi h r_o dl}) dt$. According to the law of conservation of energy [25], Q_{diss} equals to the heat loss of the water $Q_{loss} = -c\pi r_i^2 \rho dl \cdot dT$, where ρ and c are the density and specific heat capacity of the water, respectively, and dT denotes the temperature reduction within dt . Hence, we have:

$$(T - T_o) \left(\frac{\ln(\frac{r_o}{r_i})}{2\pi\lambda dl} + \frac{1}{2\pi h r_o dl} \right) dt = -c\pi r_i^2 \rho dl \cdot dT \\ \Leftrightarrow \frac{dT}{dt} + \frac{1}{c r_i^2 \rho (\frac{\ln(\frac{r_o}{r_i})}{2\lambda} + \frac{1}{2hr_o})} T = \frac{1}{c r_i^2 \rho (\frac{\ln(\frac{r_o}{r_i})}{2\lambda} + \frac{1}{2hr_o})} T_o.$$

By defining $\alpha = \pi / (\frac{\ln(\frac{r_o}{r_i})}{2\lambda} + \frac{1}{2hr_o})$, we have $T_{v,l} = (T_i - T_o)e^{-\frac{\alpha}{c\rho v}l} + T_o$, where $T_{v,l}$ refers to the water temperature at a distance of l from the inlet when the flow rate is v . Hence, we can obtain the dissipated heat through the pipe with Eq. (5):

$$P = \int_0^L dP = c\rho v(T_i - T_o)(1 - \exp(-\frac{\alpha}{c\rho v}L)). \quad (6)$$

To verify the accuracy of Eq. (6), we use Fluent software [8] to simulate the process of heat dissipation through pipes. We choose a 1-meter copper pipe whose inner radius, outer radius, and thermal conductivity are 4 mm, 5 mm, and 401 W/m°C, respectively. The ambient temperature is assumed to be stable at 20°C. From the results in Fig. 8, we find that the estimation of the dissipated heat is generally accurate, while there is still an estimation error of at most 8.4% when T_i and h are high. The main reason is that we do not consider the water temperature distribution along the radius direction though it is generally uniform. Here, we use an attenuation factor $\beta = (1 - 0.0008h) \exp(-\frac{1}{26.5\sigma + 5.2})$ to represent this effect, where $h \leq 100$ W/m²°C and the water velocity $\sigma \leq 1$ m/s.

REFERENCES

- [1] Catalogue of our peltier modules. Retrieved December 30, 2021 from <https://peltiermodules.com/?p=product>
- [2] 3M. 2021. Immersion cooling for data centers. Retrieved December 30, 2021 from https://www.3m.com/3M/en_US/data-center-us/applications/immersion-cooling/
- [3] Alekhya Addagatla, John Fernandes, Divya Mani, Dereje Agonafer, and Veerendra Mulay. 2015. Effect of warm water cooling for an isolated hybrid liquid cooled server. In *Proceedings of the 31st Thermal Measurement, Modeling & Management Symposium*. <https://doi.org/10.1109/SEMI-THERM.2015.7100161>.
- [4] Muhammad Abdullah Adnan, Ryo Sugihara, and Rajesh K Gupta. 2012. Energy efficient geographical load balancing via dynamic deferral of workload. In *Proceedings of the 5th IEEE International Conference on Cloud Computing*. <https://doi.org/10.1109/cloud49709.2020.00079>.
- [5] Alibaba. 2021. Price of the cold plate. Retrieved December 30, 2021 from https://www.alibaba.com/product-detail/Standard-40-40-mm-copper-liquid_62462733833.html

- [6] Alibaba. 2021. Price of the proportional solenoid valve. Retrieved December 30, 2021 from https://www.alibaba.com/product-detail/proportional-solenoid-valve_60730339977.html
- [7] Alibaba. 2021. Price of the vapor chamber. Retrieved December 30, 2021 from https://www.alibaba.com/product-detail/China-custom-copper-cu1100-VC-vapor_62410320227.html
- [8] Ansys. 2021. Ansys Fluent Fluid Simulation Software. Retrieved December 30, 2021 from <https://www.ansys.com/products/fluids/ansys-fluent>
- [9] Asetek. 2021. Liquid Cooling for Data Centers. Retrieved December 30, 2021 from <https://www.asetek.com/data-center/technology-for-data-centers>
- [10] Raid Ayoub, Rajib Nath, and Tajana Rosing. 2012. JETC: Joint energy thermal and cooling management for memory and CPU subsystems in servers. In *Proceedings of the 18th International Symposium on High Performance Computer Architecture*. <https://doi.org/10.1109/hpca.2012.6169035>.
- [11] Brian Stone. Why You Should Prolong GPU Lifespan (And Ways to Do it Right). Retrieved December 30, 2021 from <https://www.computer.org/publications/tech-news/trends/why-you-should-prolong-gpu-lifespan>
- [12] Charles Byers. 2021. Heterogeneous Computing in the Edge. Retrieved December 30, 2021 from https://www.iiconsortium.org/news/joi-articles/2021_June_Jol_Edge_Hetero_Compute_Byers_Final.pdf
- [13] Shutong Chen, Lei Jiao, Fangming Liu, and Lin Wang. 2022. EdgeDR: An Online Mechanism Design for Demand Response in Edge Clouds. *IEEE Transactions on Parallel and Distributed Systems* 33, 2 (2022), 343–358. <https://doi.org/10.1109/TPDS.2021.3087360>.
- [14] Shutong Chen, Zhi Zhou, Fangming Liu, Zongpeng Li, and Shaolei Ren. 2018. CloudHeat: An Efficient Online Market Mechanism for Datacenter Heat Harvesting. *ACM Transactions on Modeling and Performance Evaluation of Computing Systems* 3, 3 (2018), 31. <https://doi.org/10.1145/3199675>.
- [15] Yen-Shu Chen, Kuo-Hsiang Chien, Tzu-Chen Hung, Chi-Chuan Wang, Yuh-Ming Ferng, and Bau-Shei Pei. 2009. Numerical simulation of a heat sink embedded with a vapor chamber and calculation of effective thermal conductivity of a vapor chamber. *Applied Thermal Engineering* 29, 13 (2009), 2655–2664. <https://doi.org/10.1016/j.applthermaleng.2008.12.009>.
- [16] City Mayors. 2007. The largest cities in the world by land area, population and density. Retrieved December 30, 2021 from <http://www.citymayors.com/statistics/largest-cities-area-125.html>
- [17] cnTechPost. 2020. Tencent Cloud enables first 5G edge computing center. Retrieved December 30, 2021 from <https://cnetechpost.com/2020/10/15/tencent-cloud-enables-first-5g-edge-computing-center/>
- [18] Nick Connor. What is Fourier's Law of Thermal Conduction — Definition. Retrieved December 30, 2021 from <https://www.thermal-engineering.org/what-is-fouriers-law-of-thermal-conduction-definition/>
- [19] CoolIT Systems. 2015. Why Centralized Pumping Outperforms Distributed Systems for Rack-based Direct Liquid Cooling. Retrieved December 30, 2021 from <https://www.hpcwire.com/2015/10/26/why-centralized-pumping-outperforms-distributed-systems-for-rack-based-direct-liquid-cooling/>
- [20] CPU World. 2018. Minimum/Maximum operating temperatures. Retrieved December 30, 2021 from https://www.cpu-world.com/Glossary/M/Minimum_Maximum_operating_temperatures.html
- [21] Milnes P David, Madhusudan Iyengar, Prithvi Parida, Robert Simons, Mark Schultz, Michael Gaynes, Roger Schmidt, and Timothy Chainer. 2012. Experimental characterization of an energy efficient chiller-less data center test facility with warm water cooled servers. In *Proceedings of the 28th Annual IEEE Semiconductor Thermal Measurement and Management Symposium*. <https://doi.org/10.1109/stherm.2012.6188853>.
- [22] Dell. Dell Precision Tower 7000 Series (7910). Retrieved December 30, 2021 from <https://i.dell.com/sites/doccontent/shared-content/data-sheets/en/Documents/Dell-Precision-Tower-7000-Series-7910-Spec-Sheet.pdf>
- [23] Wei Deng, Fangming Liu, Hai Jin, Bo Li, and Dan Li. 2014. Harnessing renewable energy in cloud datacenters: opportunities and challenges. *IEEE Network* 28, 1 (2014), 48–55. <https://doi.org/10.1109/MNET.2014.6724106>.
- [24] Dr. Scott K. Thomas. Heat Conduction Equation. Retrieved December 30, 2021 from <http://cecs.wright.edu/~stthomas/hchapter02.pdf>
- [25] Energy Education. Law of Conservation of Energy. Retrieved December 30, 2021 from https://energyeducation.ca/encyclopedia/Law_of_conservation_of_energy
- [26] Engineers Edge. Convective Heat Transfer Convection Equation and Calculator. Retrieved December 30, 2021 from https://www.engineersedge.com/heat_transfer/convection.htm
- [27] Courtney Fingar. Playing it cool pays off for Iceland. Retrieved December 30, 2021 from <https://www.fdiintelligence.com/article/57142>
- [28] FPL. Air-Cooled Chillers. Retrieved December 30, 2021 from <https://infpl.fpl.com/business/pdf/air-cooled-chillers-primer.pdf>
- [29] Maurizio Frizziero. 2016. Rethinking Chilled Water Temperatures Can Bring Big Savings in Data Center Cooling. Retrieved December 30, 2021 from <https://blog.se.com/datacenter/2016/08/17/water-temperatures-data-center-cooling/>
- [30] Gartner. 2018. What Edge Computing Means for Infrastructure and Operations Leaders. Retrieved December 30, 2021 from <https://www.gartner.com/smarterwithgartner/what-edge-computing-means-for-infrastructure-and-operations-leaders/>
- [31] Íñigo Goiri, Thu D. Nguyen, and Ricardo Bianchini. 2015. CoolAir: Temperature- and Variation-Aware Management for Free-Cooled Datacenters. In *Proceedings of the 20th International Conference on Architectural Support for Programming Languages and Operating Systems*. <https://doi.org/10.1145/2775054.2694378>.
- [32] Taliver Heath, Ana Paula Centeno, Pradeep George, Luiz Ramos, Yogesh Jaluria, and Ricardo Bianchini. 2006. Mercury and Freon: Temperature emulation and management for server systems. In *Proceedings of the 12th International Conference on Architectural Support for Programming Languages and Operating Systems*. <https://doi.org/10.1145/1168918.1168872>.
- [33] Wei Huang, Shougata Ghosh, Sivakumar Velusamy, Karthik Sankaranarayanan, Kevin Skadron, and Mircea R Stan. 2006. HotSpot: A compact thermal modeling methodology for early-stage VLSI design. *IEEE Transactions on very large scale integration systems* 14, 5 (2006), 501–513. <https://doi.org/10.1109/tvlsi.2006.876103>.
- [34] Intel. The Efficient Datacenter. Retrieved December 30, 2021 from <https://www.intel.co.jp/content/dam/doc/technology-brief/efficient-datacenter-high-ambient-temperature-operation-brief.pdf>
- [35] Intel. Intel Xeon Processor E5-2680 v4. Retrieved December 30, 2021 from <https://ark.intel.com/content/www/us/en/ark/products/91754/intel-xeon-processor-e5-2680-v4-35m-cache-2-40-ghz.html>
- [36] Arman Iranfar, Ali Pahlevan, Marina Zapater, and David Atienza. 2019. Enhancing Two-Phase Cooling Efficiency through Thermal-Aware Workload Mapping for Power-Hungry Servers. In *Proceedings of the Design, Automation & Test in Europe Conference & Exhibition*. <https://doi.org/10.23919/DATE.2019.8715033>.
- [37] Leila Ismail and Huneed Materwala. 2020. Computing Server Power Modeling in a Data Center: Survey, Taxonomy, and Performance Evaluation. *Comput. Surveys* 53, 3 (2020), 34 pages. <https://doi.org/10.1145/3390605>.
- [38] Majid Jalili, Ioannis Manousakis, Íñigo Goiri, Pulkit A Misra, Ashish Raniwala, Hosam Alissa, Bharath Ramakrishnan, Phillip Tuma, Christian Belady, Marcus Foutoura, and Ricardo Bianchini. 2021. Cost-efficient overcooling in immersion-cooled datacenters. In *Proceedings of the 48th Annual International Symposium on Computer Architecture*. <https://doi.org/10.1109/ISCA52012.2021.00055>.
- [39] Radosław Januszewski, Norbert Meyer, and Joanna Nowicka. 2014. Evaluation of the impact of direct warm-water cooling of the HPC servers on the data center ecosystem. In *Proceedings of the International Supercomputing Conference*. https://doi.org/10.1007/978-3-319-07518-1_24.
- [40] Kaixuan Ji, Ce Chi, Avinab Marahatta, Fa Zhang, and Zhiyong Liu. 2020. Energy Efficient Scheduling Based on Marginal Cost and Task Grouping in Data Centers. In *Proceedings of the 11th ACM International Conference on Future Energy Systems*. <https://doi.org/10.1145/3396851.3402657>.
- [41] Weixiang Jiang, Ziyang Jia, Sirui Feng, Fangming Liu, and Hai Jin. 2019. Fine-grained warm water cooling for improving datacenter economy. In *Proceedings of the 46th Annual International Symposium on Computer Architecture*. <https://doi.org/10.1145/3307650.3322236>.
- [42] George Kamiya. 2021. Data Centres and Data Transmission Networks. Retrieved December 30, 2021 from <https://www.iea.org/reports/data-centres-and-data-transmission-networks>
- [43] Patrick Kennedy. 2018. Sugon Nebula Phase Change Immersion Cooling a Unique Platform. Retrieved December 30, 2021 from <https://www.servethehome.com/sugon-nebula-phase-change-immersion-cooling-a-unique-platform/>
- [44] S Karen Khatamifard, Longfei Wang, Weize Yu, Selçuk Köse, and Ulya R Karpuzcu. 2017. ThermoGater: Thermally-aware on-chip voltage regulation. In *Proceedings of the 44th Annual International Symposium on Computer Architecture*. <https://doi.org/10.1145/3079856.3080250>.
- [45] Hassan Khurshid, Karthik Silaipillayarputhur, and Tawfiq Al Mughnam. 2018. Design of a Heat Sink for an misc Component in ABB Drive using Different Types of Fins. In *MATEC Web of Conferences*, Vol. 249. 03009. <https://doi.org/10.1051/mateconf/201824903009>.
- [46] Min-Hwi Kim, Sang-Woo Ham, Jun-Seok Park, and Jae-Weon Jeong. 2014. Impact of integrated hot water cooling and desiccant-assisted evaporative cooling systems on energy savings in a data center. *Energy* 78 (2014), 384–396. <https://doi.org/10.1016/j.energy.2014.10.023>.
- [47] Philip Kosky, Robert Balmer, William Keat, and George Wise. 2021. Chapter 14 - Mechanical Engineering. In *Exploring Engineering (Fifth Edition)* (fifth edition ed.), Philip Kosky, Robert Balmer, William Keat, and George Wise (Eds.). Academic Press, 317 – 340. <https://doi.org/10.1016/B978-0-12-815073-3.00014-4>.
- [48] Kristian Valur Laursen Olason, Alexandru Uta, Alexandru Iosup, Paul Melis, Damian Podareanu, and Valeriu Codreanu. 2020. Beneath the SURFace: An MRI-like View into the Life of a 21st Century Datacenter. <https://doi.org/10.5281/zenodo.3878143>
- [49] Seunghak Lee, Ki-Dong Kang, Hwanjun Lee, Hyungwon Park, Younghoon Son, Nam Sung Kim, and Daehoon Kim. 2021. GreenDIMM: OS-assisted DRAM Power Management for DRAM with a Sub-array Granularity Power-Down State. In *Proceedings of the 54th Annual International Symposium on Microarchitecture*. <https://doi.org/10.1145/3466752.3480089>.
- [50] LF Edge. State of the Edge 2020. Retrieved December 30, 2021 from <https://www.lfedge.org/wp-content/uploads/2020/04/SOTE2020.pdf>

- [51] LF Edge. State of the Edge 2021. Retrieved December 30, 2021 from <https://www.lfedge.org/2021/03/12/state-of-the-edge-2021-report/>
- [52] Jiang Lin, Hongzhong Zheng, Zhichun Zhu, Howard David, and Zhao Zhang. 2007. Thermal modeling and management of DRAM memory systems. In *Proceedings of the 34th Annual International Symposium on Computer Architecture*. <https://doi.org/10.31274/etd-180810-2677>.
- [53] Fangming Liu, Zhi Zhou, Hai Jin, Bo Li, Baochun Li, and Hongbo Jiang. 2014. On Arbitrating the Power-Performance Tradeoff in SaaS Clouds. *IEEE Transactions on Parallel and Distributed Systems* 25, 10 (2014), 2648–2658. <https://doi.org/10.1109/TPDS.2013.208>.
- [54] Jie Liu, Michel Goraczko, Sean James, Christian Belady, Jiakang Lu, and Kamin Whitehouse. 2011. The Data Furnace: Heating Up with Cloud Computing. In *Proceedings of the 3rd USENIX Workshop on Hot Topics in Cloud Computing*. <https://doi.org/10.1.1.233.1649>.
- [55] Song Liu, Brian Leung, Alexander Neckar, Seda Oğrenci Memik, Gokhan Memik, and Nikos Hardavellas. 2011. Hardware/software techniques for DRAM thermal management. In *Proceedings of the 17th International Symposium on High Performance Computer Architecture*. <https://doi.org/10.1109/hpca.2011.5749756>.
- [56] Wangyu Liu, Jingren Gou, Yuanqiang Luo, and Min Zhang. 2018. The experimental investigation of a vapor chamber with compound columns under the influence of gravity. *Applied Thermal Engineering* 140 (2018), 131–138. <https://doi.org/10.1016/j.applthermaleng.2018.05.010>.
- [57] Zhenhua Liu, Yuan Chen, Cullen Bash, Adam Wierman, Daniel Gmach, Zhikui Wang, Manish Marwah, and Chris Hyser. 2012. Renewable and cooling aware workload management for sustainable data centers. In *Proceedings of the 12th ACM SIGMETRICS/PERFORMANCE joint international conference on Measurement and Modeling of Computer Systems*. <https://doi.org/10.1145/2318857.2254779>.
- [58] Ioannis Manousakis, Inigo Goiri, Sriram Sankar, Thu D. Nguyen, and Ricardo Bianchini. 2015. CoolProvision: Underprovisioning datacenter cooling. In *Proceedings of the 6th ACM Symposium on Cloud Computing*. <https://doi.org/10.1145/2806777.2806938>.
- [59] Energy Efficiency Manual. 2015. Keep the chilled water supply temperature as high as possible. Retrieved December 30, 2021 from <http://energybooks.com/wp-content/uploads/2015/07/264266.pdf>
- [60] McQuay. Air-cooled Screw Chiller. Retrieved December 30, 2021 from <http://www.mcquay.com.hk/my-product/air-cooled-screw-chiller-awsmcs/>
- [61] SK Mishra, H Chandra, and A Arora. 2019. Effects on heat transfer and radial temperature profile of non-isoviscous vibrational flow with varying Reynolds number. *Journal of Applied Fluid Mechanics* 12, 1 (2019), 135–144. <https://doi.org/10.29252/jafm.75.253.28952>.
- [62] John Musilli and Paul Vaccaro. 2014. Intel IT Redefines the High-Density Data Center: 1,100 Watts/Sq Ft. Retrieved December 30, 2021 from <https://www.intel.com/content/dam/www/public/us/en/documents/white-papers/intel-it-redefines-high-density-data-center-1100-watts-per-sqft-paper.pdf>
- [63] N-able. 2019. Servers vs. Workstations Overview. Retrieved December 30, 2021 from <https://www.n-able.com/blog/servers-vs-workstations>
- [64] NVIDIA. NVIDIA A100 Tensor Core GPU. Retrieved December 30, 2021 from <https://www.nvidia.com/en-us/data-center/a100/>
- [65] Open Compute Project. Open Rack/SpecsAndDesigns. Retrieved December 30, 2021 from https://www.opencompute.org/wiki/Open_Rack/SpecsAndDesigns
- [66] Mohammad Parhizi, Ali Akbar Merrikh, and Ankur Jain. 2014. Investigation of Two-Phase, Vapor Chamber Based Thermal Management of Multiple Microserver Chips. In *Proceedings of the International Mechanical Engineering Congress and Exposition*. <https://doi.org/10.1115/imece2014-39928>.
- [67] Parker. Miniature Proportional Valve. Retrieved December 30, 2021 from <https://ph.parker.com/us/12051/en/hf-pro-miniature-proportional-valve/920-000047-001>
- [68] Pol. What is Immersion Cooling. Retrieved December 30, 2021 from <https://submer.com/blog/what-is-immersion-cooling/>
- [69] Rounak Prashnani. 2018. Task Migration Algorithm to Reduce Temperature Imbalance Amongst Cores in Linux Based Multi-Core Processor Systems. In *Proceedings of the 3rd International Conference for Convergence in Technology*. <https://doi.org/10.1109/i2ct.2018.8529564>.
- [70] Radian. What is a Vapor Chamber Heatsink? Retrieved December 30, 2021 from <https://www.radianheatsinks.com/vapor-chamber-heatsink/>
- [71] Luiz Ramos and Ricardo Bianchini. 2008. C-Oracle: Predictive thermal management for data centers. In *Proceedings of the 14th International Symposium on High Performance Computer Architecture*. <https://doi.org/10.1109/hpca.2008.4658632>.
- [72] Yongyi Ran, Han Hu, Xin Zhou, and Yonggang Wen. 2019. DeepEE: Joint optimization of job scheduling and cooling control for data center energy efficiency using deep reinforcement learning. In *Proceedings of the 39th International Conference on Distributed Computing Systems*. <https://doi.org/10.1109/icdc.2019.00070>.
- [73] Rittal. 2020. Data Center Cooling: 4 Effective Types of Liquid Cooling. Retrieved December 30, 2021 from https://zutacore.com/wp-content/uploads/2020/06/Data_Center_Cooling_Liquid_Cooling_US528.pdf
- [74] John Roach. 2020. Microsoft finds underwater datacenters are reliable, practical and use energy sustainably. Retrieved December 30, 2021 from <https://news.microsoft.com/innovation-stories/project-natick-underwater-datacenter/>
- [75] Safe Temp. 2019. Nvidia GeForce RTX 2080 Ti Max Temp. Retrieved December 30, 2021 from <https://safetemp.blogspot.com/2019/01/geforce-rtx-2080-ti-max-temp.html>
- [76] Safe Temp. 2021. Nvidia A100 PCIe 80GB Max Temp. Retrieved December 30, 2021 from <https://safetemp.blogspot.com/2021/10/a100-pcie-80-gb-max-temp.html>
- [77] Manasa Sahini, Chinmay Kshirsagar, Mathan Kumar, Dereje Agonafer, John Fernandes, Jacob Na, Veerendra Mulay, Patrick McGinn, and Michael Soares. 2017. Rack-level study of hybrid cooled servers using warm water cooling for distributed vs. centralized pumping systems. In *Proceedings of the 33rd Thermal Measurement, Modeling & Management Symposium*. <https://doi.org/10.1109/semi-therm.2017.7896924>.
- [78] Samsung. 983 DCT. Retrieved December 30, 2021 from <https://www.samsung.com/semiconductor/minisite/ssd/product/data-center/983dct/>
- [79] Schneider Electric. Liquid Cooling Technologies for Data Centers and Edge Applications. Retrieved December 30, 2021 from https://download.schneider-electric.com/files?p_Doc_Ref=SPD_VAVR-AQKM3N_EN
- [80] André Seuret, Arman Iranfar, Marina Zapater, John Thome, and David Atienza. 2018. Design of a two-phase gravity-driven micro-scale thermosyphon cooling system for high-performance computing data centers. In *Proceedings of the 17th IEEE Intersociety Conference on Thermal and Thermomechanical Phenomena in mixed Systems*. <https://doi.org/10.1109/ITHERM.2018.8419531>.
- [81] Shriganesh Shurpali. Role of Edge Computing in Connected and Autonomous Vehicles. Retrieved December 30, 2021 from <https://www.einfochips.com/blog/role-of-edge-computing-in-connected-and-autonomous-vehicles/>
- [82] Matt Skach, Manish Arora, Chang-Hong Hsu, Qi Li, Dean Tullsen, Lingjia Tang, and Jason Mars. 2015. Thermal time shifting: Leveraging phase change materials to reduce cooling costs in warehouse-scale computers. In *Proceedings of the 42nd Annual International Symposium on Computer Architecture*. <https://doi.org/10.1145/2749469.2749474>.
- [83] Matt Skach, Manish Arora, Dean Tullsen, Lingjia Tang, and Jason Mars. 2018. Virtual melting temperature: Managing server load to minimize cooling overhead with phase change materials. In *Proceedings of the 45th Annual International Symposium on Computer Architecture*. <https://doi.org/10.1109/isca.2018.00013>.
- [84] Arvind Sridhar, Alessandro Vincenzi, Martino Ruggiero, Thomas Brunschweiler, and David Atienza. 2010. 3D-ICE: Fast compact transient thermal modeling for 3D ICs with inter-tier liquid cooling. In *Proceedings of the 28th International Conference on Computer-Aided Design*. <https://doi.org/10.1109/ICCAD.2010.5653749>.
- [85] Statista. 2020. Electricity prices for industries in the European Union in 2019, by country. Retrieved December 30, 2021 from <https://www.statista.com/statistics/1046605/industry-electricity-prices-european-union-country/>
- [86] Tameson. Proportional Solenoid Valve – How They Work. Retrieved December 30, 2021 from <https://tameson.com/proportional-solenoid-control-valve.html>
- [87] Tameson. Solenoid valve response time. Retrieved December 30, 2021 from <https://tameson.com/solenoid-valve-response-time.html>
- [88] Tech Power Up. Nvidia GeForce RTX 2080 Ti. Retrieved December 30, 2021 from <https://www.techpowerup.com/gpu-specs/geforce-rtx-2080-ti-3x305>
- [89] Trijo Tharayil, Lazarus Godson Asirvatham, Vysakh Ravindran, and Somchai Wongwises. 2016. Effect of filling ratio on the performance of a novel miniature loop heat pipe having different diameter transport lines. *Applied Thermal Engineering* 106 (2016), 588–600. <https://doi.org/10.1016/j.applthermaleng.2016.05.125>.
- [90] Thorne & Derrick. ASCO 220 Solenoid Valves. Retrieved December 30, 2021 from <https://www.heatingandprocess.com/product/product-category/asco-220-solenoid-valves-steam/>
- [91] TLX Technologies. PWM Solenoid Theory. Retrieved December 30, 2021 from <https://www.tlxtech.com/understanding-solenoids/theory-operation/pwm-solenoid-theory>
- [92] Meng-Chang Tsai, Shung-Wen Kang, and Kleber Vieira de Paiva. 2013. Experimental studies of thermal resistance in a vapor chamber heat spreader. *Applied Thermal Engineering* 56, 1–2 (2013), 38–44. <https://doi.org/10.1016/j.applthermaleng.2013.02.034>.
- [93] Vapor IO. The Vapor Chamber. Retrieved December 30, 2021 from <https://www.vapor.io/chamber/>
- [94] Vapor IO. Vapor Chamber Technical Specifications. Retrieved December 30, 2021 from <https://www.vapor.io/wp-content/uploads/2018/05/Vapor-Chamber-Tech-Specs-V3.pdf>
- [95] Vertiv. Liquid Cooling Options for Data Centers. Retrieved December 30, 2021 from <https://www.vertiv.com/en-us/solutions/learn-about/liquid-cooling-options-for-data-centers/>
- [96] VIAVI Solutions. What is 5G Technology? Retrieved December 30, 2021 from <https://www.viavisolutions.com/en-us/5g-technology>
- [97] Herb Villa. 2020. Liquid Cooling vs. Immersion Cooling Deployment. Retrieved December 30, 2021 from <https://blog.rittal.us/liquid-cooling-vs-immersion-cooling-deployment>

- [98] Victor Avelar Vincent Petit, Steven Carlini. Digital economy and climate impact. Retrieved December 30, 2021 from https://download.schneider-electric.com/files?p_File_Name=998-21202519.pdf
- [99] Béla Waldhauser. 2019. The Edge & Edge Data Centers: Gaining Clarity. Retrieved December 30, 2021 from <https://www.dotmagazine.online/issues/on-the-edge-building-the-foundations-for-the-future/edge-data-centers>
- [100] Jianxiong Wan, Xiang Gui, Shoji Kasahara, Yuanyu Zhang, and Ran Zhang. 2018. Air flow measurement and management for improving cooling and energy efficiency in raised-floor data centers: A survey. *IEEE Access* 6 (2018), 48867–48901. <https://doi.org/10.1109/ACCESS.2018.2866840>.
- [101] Qiming Wang, Julie A Jackson, Qi Ge, Jonathan B Hopkins, Christopher M Spadaccini, and Nicholas X Fang. 2016. Lightweight mechanical metamaterials with tunable negative thermal expansion. *Physical review letters* 117, 17 (2016), 175901. <https://doi.org/10.1103/physrevlett.117.175901>.
- [102] Rong-Tsu Wang, Jung-Chang Wang, and Tien-Li Chang. 2011. Experimental analysis for thermal performance of a vapor chamber applied to high-performance servers. *Journal of Marine Science and Technology* 19, 4 (2011), 353–360. <https://doi.org/10.51400/2709-6998.2175>.
- [103] WikiChip. Die floorplan of Intel Xeon E5. Retrieved December 30, 2021 from [https://en.wikichip.org/wiki/intel/microarchitectures/broadwell_\(client\)](https://en.wikichip.org/wiki/intel/microarchitectures/broadwell_(client))
- [104] Wikipedia. Coefficient of performance. Retrieved December 30, 2021 from https://en.wikipedia.org/wiki/Coefficient_of_performance
- [105] Wikipedia. Die (integrated circuit). Retrieved December 30, 2021 from [https://en.wikipedia.org/wiki/Die_\(integrated_circuit\)](https://en.wikipedia.org/wiki/Die_(integrated_circuit))
- [106] Wikipedia. Thermal design power. Retrieved December 30, 2021 from https://en.wikipedia.org/wiki/Thermal_design_power
- [107] Songkran Wiriyaart and Paisarn Naphon. 2018. Fill ratio effects on vapor chamber thermal resistance with different configuration structures. *International Journal of Heat and Mass Transfer* 127 (2018), 164–171. <https://doi.org/10.1016/j.ijheatmasstransfer.2018.07.029>.
- [108] Mengwei Xu, Zhe Fu, Xiao Ma, Li Zhang, Yanan Li, Feng Qian, Shangguang Wang, Ke Li, Jingyu Yang, and Xuanzhe Liu. 2021. From cloud to edge: a first look at public edge platforms. In *Proceedings of the 21st ACM Internet Measurement Conference*. <https://doi.org/10.1145/3487552.3487815>.
- [109] Michael Yoshizawa. Newton's law of cooling or heating. Retrieved December 30, 2021 from <http://web.math.ucsb.edu/~myoshi/cooling.pdf>
- [110] Zihao Yuan, Geoffrey Vaartstra, Prachi Shukla, Sherief Reda, Evelyn Wang, and Ayse K Coskun. 2019. Modeling and optimization of chip cooling with two-phase vapor chambers. In *Proceedings of the IEEE/ACM International Symposium on Low Power Design and Design*. <https://doi.org/10.1109/islped.2019.8824965>.
- [111] Chaojie Zhang, Alok Gautam Kumbhare, Ioannis Manousakis, Deli Zhang, Pulkit A. Misra, Rod Assis, Kyle Woolcock, Nithish Mahalingam, Brijesh Warrier, David Gauthier, Lalu Kunnath, Steve Solomon, Osvaldo Morales, Marcus Fontoura, and Ricardo Bianchini. 2021. Flex: High-Availability Datacenters With Zero Reserved Power. In *Proceedings of the 48th Annual International Symposium on Computer Architecture*. <https://doi.org/10.1109/ISCA52012.2021.00033>.
- [112] Wenli Zheng, Kai Ma, and Xiaorui Wang. 2014. Exploiting thermal energy storage to reduce data center capital and operating expenses. In *Proceedings of the 20th International Symposium on High Performance Computer Architecture*. <https://doi.org/10.1109/hpca.2014.6835924>.
- [113] Liang Zhou, Laxmi N Bhuyan, and KK Ramakrishnan. 2020. Gemini: Learning to Manage CPU Power for Latency-Critical Search Engines. In *Proceedings of the 53rd Annual International Symposium on Microarchitecture*. <https://doi.org/10.1109/micro50266.2020.00059>.
- [114] Zhi Zhou, Fangming Liu, Shutong Chen, and Zongpeng Li. 2020. A Truthful and Efficient Incentive Mechanism for Demand Response in Green Datacenters. *IEEE Transactions on Parallel and Distributed Systems* 31, 1 (2020), 1–15. <https://doi.org/10.1109/TPDS.2018.2882174>.
- [115] Zhi Zhou, Fangming Liu, Bo Li, Baochun Li, Hai Jin, Ruolan Zou, and Zhiyong Liu. 2014. Fuel cell generation in geo-distributed cloud services: A quantitative study. In *Proceedings of the 34th International Conference on Distributed Computing Systems*. <https://doi.org/10.1109/ICDCS.2014.14>.
- [116] Zhi Zhou, Fangming Liu, Ruolan Zou, Jiangchuan Liu, Hong Xu, and Hai Jin. 2016. Carbon-Aware Online Control of Geo-Distributed Cloud Services. *IEEE Transactions on Parallel and Distributed Systems* 27, 9 (2016), 2506–2519. <https://doi.org/10.1109/TPDS.2015.2504978>.
- [117] Qi Zhu, Xiang Li, and Yinan Wu. 2008. Thermal management of high power memory module for server platforms. In *Proceedings of the 11th Intersociety Conference on Thermal and Thermomechanical Phenomena in misc Systems*. <https://doi.org/10.1109/therm.2008.4544319>.
- [118] Xinhui Zhu, Weixiang Jiang, Fangming Liu, Qixia Zhang, Li Pan, Qiong Chen, and Ziyang Jia. 2020. Heat to power: Thermal energy harvesting and recycling for warm water-cooled datacenters. In *Proceedings of the 47th Annual International Symposium on Computer Architecture*. <https://doi.org/10.1109/isca45697.2020.00042>.
- [119] Pengfei Zou, LP Ang, Kevin Barker, and Rong Ge. 2020. Indicator-Directed Dynamic Power Management for Iterative Workloads on GPU-Accelerated Systems. In *Proceedings of the 20th IEEE/ACM International Symposium on Cluster, Cloud and Internet Computing*. <https://doi.org/10.1109/ccgrid49817.2020.00-37>.