

REVIEW

Open Access



Enough hot air: the role of immersion cooling

Kawsar Haghshenas^{1*}, Brian Setz¹, Yannis Blosch² and Marco Aiello²

*Correspondence:
k.haghshenas@rug.nl

¹ Bernoulli Institute, University of Groningen, Groningen, The Netherlands

² Service Computing Department, IAAS, University of Stuttgart, Stuttgart, Germany

Abstract

Air cooling is the traditional solution to chill servers in data centers. However, the continuous increase in global data center energy consumption combined with the increase of the racks' power dissipation calls for the use of more efficient alternatives. Immersion cooling is one such alternative. In this paper, we quantitatively examine and compare air cooling and immersion cooling solutions. The examined characteristics include power usage efficiency (PUE), computing and power density, cost, and maintenance overheads. A direct comparison shows a reduction of about 50% in energy consumption and a reduction of about two-thirds of the occupied space, by using immersion cooling. In addition, the higher heat capacity of used liquids in immersion cooling compared to air allows for much higher rack power densities. Moreover, immersion cooling requires less capital and operational expenditures. However, challenging maintenance procedures together with the increased number of IT failures are the main downsides. By selecting immersion cooling, cloud providers must trade-off the decrease in energy and cost and the increase in power density with its higher maintenance and reliability concerns. Finally, we argue that retrofitting an air-cooled data center with immersion cooling will result in high costs and is generally not recommended.

Keywords: Air cooling, Data center, Immersion cooling, Power density, Maintenance, Power usage efficiency

Introduction

The scale and the number of data centers are increasing worldwide to meet the rising demand for IT applications and cloud services. High energy consumption (operational cost) and environmental effects are among the main concerns, which both require a great deal of attention. Due to the fact that IT equipment produces heat under load, they have to be artificially cooled to ensure availability and reliability. Therefore, the cooling system is one of the main energy consumers of a data center. Presently, air-cooling is the most prominent technique used in data centers. Using this technique entails that about 40% of the total energy consumption is dedicated to cooling (Ni and Bai 2017).

In an air-cooled data center, cold air is circulating through the perforated tiles up and into the front of the servers and hot air is pushed out by new cold air coming into the servers. By this method, it is possible to cool server racks with at most 50kW power

density (Kheirabadi and Groulx 2016). In 2018, only 10% of respondents to a survey reported that the power density of some of their racks was above 40 kW (Smolaks 2019). However, growing power-hungry cluster workloads, such as machine learning ones, produce higher heat densities. These workloads are run on dozens of GPUs, each having thousands of cores that need to be supplied by power. The GPUs increase the power density of the rack, requesting more cooling capacity. Moreover, recent studies show that as the end of Dennard scaling is reached (Dennard et al. 1974), and transistor sizes approach their practical limitations, new cooling solutions are needed to keep the traditional performance improvement trend. Therefore, CPUs and GPUs with higher power consumption are expected to be manufactured in near future (Sun et al. 2019; Fan et al. 2018; Intel 2017), and consequently, the power density of racks will increase as well.

To address both high energy consumption and low power density problems associated with air cooling, researchers and cloud providers have started to explore alternative solutions. Liquid immersion cooling is a viable one that has attracted attention in the last decade. In immersion cooling, components are fully immersed into a dielectric fluid that conducts heat and does not conduct electricity, therefore, the heat of all IT components is fully removed by liquid, which reduces the power usage efficiency (PUE) of the data center. The PUE is defined as the ratio between total energy and the energy used for the IT equipment of a data center.

The PUE of immersion-cooled data centers is close to perfection, about 1.02–1.04 (Matsuoka et al. 2017; An et al. 2018; Eiland et al. 2014; Chandrasekaran et al. 2017; Shah 2018) which shows that these centers consume 10–50% less energy compared to their air-cooled counterparts, for the same amount of computational load. In addition, compared to the air, common dielectric fluids have a much higher heat capacity. Therefore, immersion cooling allows for more computing power in less space. While the maximum power density per rack for an air-cooled data center is around 50 kW, immersion-cooled counterpart allows for up to 250 kW per rack (Kheirabadi and Groulx 2016; Two-phase immersion cooling a revolution 2015).

Immersion cooling is certainly an efficient solution in terms of computing efficiency and power density. In addition, the capital expenditure for constructing a data center, keeping constant the power density, is lower with immersion cooling compared to air cooling (Bunger et al. 2020). But why it is not deployed in today's typical data centers? Maintenance and reliability, and specifically the lack of practical information on these aspects, are the primary concerns for adapting immersion cooling technologies (Coles and Herrlin 2016; Jalili et al. 2021; Villa 2020; Alibaba 2018; Varma 2019; Ramdas et al. 2019). Increased number of IT equipment failures, liquid leakage, and liquid evaporation are the main challenges associated with its maintenance procedure, which also impose additional operational costs. However, there are several successful implementations of immersion cooling in data centers. For example, one of the biggest players in cloud computing, Microsoft, has already constructed its first immersion-cooled data center (Weston 2021).

In this paper, we explore quantitatively the trade-offs between air and immersion cooling technologies and evaluate several aspects of both methods. For this evaluation, we refer to various references, from research studies to practical implementations. Additionally, quantitative analyses on data center efficiency and computing power are

presented throughout the paper. This paper will help data center operators to have a better perspective while selecting the best solution for their specific application. Our investigations show that immersion cooling has significant advantages specifically for high power applications and we expect its adoption to grow. Improving the efficiency of a big data center by a small percentage would considerably impact total energy consumption. In addition, the migration from air to liquid cooling could influence operators of smaller data centers in trusting immersion cooling. However, we also argue that retrofitting an air-cooled data center with liquid cooling will result in high costs and is generally not recommended.

The remainder of this paper is organized as follows. The basic concepts of air cooling and immersion cooling, together with a literature review on improving data center efficiency, are presented in Sect. "Background". Then, Sect. "Immersion cooling in practice" contains an overview of several practical implementations of immersion-cooled data centers and several companies that offer immersion cooling server systems. Section "Efficiency, Density, Cost, and Maintenance" presents an analysis and comparison of the two cooling solutions including in terms of efficiency, density, cost, and maintenance. Finally, concluding remarks are provided in Sect. "Conclusions".

Background

Computing hardware consumes electricity, which is dissipated as heat due to resistances in the electrical circuits. In turn, this heat needs to be dispersed to ensure the correct functioning of the equipment. In this section, we provide some background on air cooling and immersion cooling solutions for data centers. In addition, we present a literature review (of reviews) on data centers' cooling solutions.

Air cooling

Data center cooling was not an issue when server rooms were still small and computing hardware did not exceed a rack power density of 5 kW (Voices of the Industry 2020). These relatively low quantities of heat were removed by basic air conditioning and server fans, which was the easiest solution at the time. Being cheap and easy to implement, air cooling quickly became the standard solution for server cooling and was further developed.

An air cooling system is fundamentally described by the heat removal method, air distribution type, and the location of the cooling unit which directly supplies cool air to the IT equipment. Heat removal is the process of moving heat energy from the IT space to the outdoors. Moving heat is accomplished by using a heat exchanger to transfer heat from one fluid to another (e.g. from air to water) (Evans 2012). The indoor heat exchanger can be a computer room air conditioning (CRAC), a computer room air handling (CRAH), or a self-contained system and the outdoor heat exchange can be a chiller, a cooling tower, a dry cooler, etc. The heat exchanger is responsible for cooling the fluid to a desirable temperature (Geng 2014). In favourable outdoors conditions, it is possible to bypass indoor/outdoor heat exchangers and directly use natural air to cool the transport fluid. This is referred to as free cooling.

The principal aim of any type of air distribution is to regulate airflow. Air distribution types such as flooded, targeted, and contained differentiate various cooling systems

in data centers. The flooded type allows the airflow to be primarily constrained by the physical boundaries of the room, including the walls, ceiling, and floor. This results in substantial mixing of hot and cold air streams. The targeted type uses a specific mechanism to direct both supply and return airflow within three meters of the IT equipment's intake and exhaust. Lastly, the contained type fully encloses the supply and return airflow of the IT equipment, thereby avoiding any mixing between the two air streams. All these methods can be utilized in either the supply or return path. The type of air distribution and its management has a significant impact on the maximum potential power density of clusters (Rasmussen 2017).

The cooling unit, which supplies chilled air to the IT equipment, comes in three variants: row-oriented, rack-oriented, and room-oriented. The placement of the cooling unit is crucial in aspects of data center design, cooling efficacy, power density, and space usage (Dunlap and Rasmussen 2006). An efficient system aims at designing an air distribution system that can handle average power density while being capable of managing peak power density. The challenge of mitigating potential "hot spots" using traditional designs often leads to the cooling plant and air distribution system being oversized, thereby incurring additional capital and operational expenses. A promising solution lies in containing the return air to manage high-density areas and to prevent extra costs.

In fact, the prevention of hot and cold air mixing is crucial for efficient air cooling in data centers (Rasmussen 2017). This can be achieved by strategies such as the hot aisle containment system (HACS) and the cold aisle containment system (CACS). These strategies enhance power density and efficiency in data centers when compared to conventional cooling designs. The CACS system encloses the cold aisle, allowing the rest of the data center to function as a large hot-air return plenum. Conversely, the HACS system encloses the hot aisle, capturing hot air and enabling the rest of the room to serve as a large cold-air supply plenum. Generally, HACS is more efficient than CACS as it enables higher temperatures and increased chilled water temperatures, which result in more economizer hours and savings on electricity. However, it is easier and less expensive to retrofit a CACS in an existing raised floor and room-oriented cooling system (Niemann et al. 2011).

Fig. 1 offers a basic illustration of the air cooling methodology used for racks in data centers. It depicts the most elementary forms featuring straightforward air circulation on the left and more advanced systems, like CRAC cooling, on the right.

Immersion cooling

Immersion cooling is an approach that uses liquid instead of air to remove heat from computing hardware. In this method, components are fully immersed into a dielectric fluid. A dielectric fluid conducts heat and does not conduct electricity at all; instead, it acts as an insulator. Therefore, the heat of all the components is fully removed by the liquid, eliminating the need for air cooling and moving parts (i.e. server fans). The commonly used dielectric fluids have a much higher heat capacity than air (Shinde 2019; Eiland et al. 2014). Most liquids used in immersion cooling are white mineral oil, engineered electric fluids, and other oils (Shinde 2019; Jalili et al. 2021). There are two types of liquid immersion cooling: single-phase and two-phase, Fig. 2.

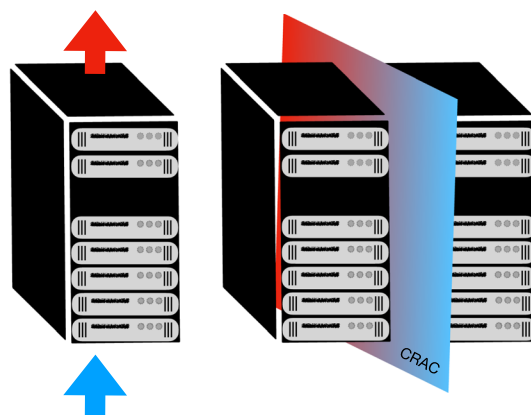


Fig. 1 Air-based cooling techniques for data centers

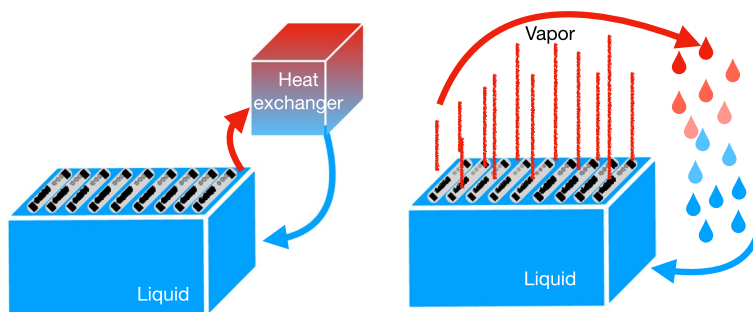


Fig. 2 Liquid-based cooling techniques for data centers

In single-phase immersion cooling, the fluid stays in liquid form and there is no phase change. Heat emitting components are cooled by the fluid flowing over them, and the heated fluid is transported away. The circulation of the fluid is driven by a pump or by natural convection. In natural convection-driven systems, the heated fluid floats to the top of the tank because it has a higher volume than colder fluid. It then flows, due to more fluid rising to the top, to the side of the tank where it is cooled by a heat exchanger connected to an external loop. The density of the cooled liquid increases again and it sinks back to the bottom of the tank by gravity. In a pump-driven system, the convection is driven by a pump. The pump forces the liquid through an inlet inside the tank and out through an outlet on the opposite side. The liquid is then cooled by flowing through a coolant-to-water heat exchanger.

In two-phase immersion cooling, the coolant changes phase whenever it gets in contact with a heat-producing component. In order to avoid damaging the components, the boiling point of the coolant has to be lower than the critical temperature of the components. While evaporating on a hot component, the gas will float to the top of the tank and will make room for the new colder liquid coolant. A condenser is located inside the tank above the liquid. Cooling water is flowing through the condenser to transport the heat away. The coolant condenses there and will fall back into the tank where it can absorb heat again (Kanbur et al. 2020).

Literature

There is a vast body of literature about increasing the efficiency of data centers. These works mostly focus on the scheduling of the workload and its optimization, e.g., Son et al. (2017); Haghshenas et al. (2020). The specific aspect of cooling has also attracted some attention and several reviews on cooling solutions for data centers have appeared. One of the first reviews considers the thermal aspects for air cooling in data centers (Patankar 2010). Even though some authors were already advocating a role for immersion cooling (Tuma 2010), the review in Patankar (2010) only considers air-based solutions. In Ebrahimi et al. (2014), a review of data center cooling technologies is offered, and particular attention is devoted to the opportunity of recovering the heat. For example, district heating can be fed from recovered hot air from a data center.

Li et al. offer a detailed thermal analysis of cooling solutions, including several strategies based on cold plates, waste heat recovery, and heat pipes (Li and Kandlikar 2015). Immersion cooling is not considered as one of the possible solutions. Kheirabadi and Groulx also focus on the thermal aspects of server cooling (Kheirabadi and Groulx 2016). They classify the solutions as air-based (CRAC, CRAH, RDHx, and SCHx) and liquid-based (indirect: single-phase, two-phase, heat pipe; and direct: pool boiling, spray cooling, and jet impingement). The results of the review and technology comparison show a higher efficiency for liquid cooling-based solutions while maintaining that air-based solutions are a valid option especially if high efficiency is not a top requirement. The chemical and physical effects of immersion cooling on IT equipment are investigated in Shah et al. (2016). The results are particularly relevant to understand the maintenance needs and life expectancy of the equipment. In addition, a number of advantages of immersion cooling are identified, such as a decrease in overheating and temperature swings, elimination of fan failures, elimination of dust and moisture-related failures, and reduced corrosion.

A review of the literature and product descriptions is presented in Watson and Venkiteswaran (2017). The authors select two prototypical systems, one based on air cooling and one on immersion cooling, and provide a comparison of the two based on simulations especially focusing on scalability aspects. The results indicate a preference for immersion cooling. A review of thermal management in data centers is offered in Nadjahi et al. (2018). The authors recognize that open aisle air-based cooling is the dominant one in practice, though focus on novel techniques and technologies that have good promise in terms of energy efficiency. Kuncoro et al. review exclusively immersion cooling solutions and compare the performance of different types of cooling liquids (Kuncoro 2019).

Building on the current state of the art, our aim with the present work is to provide a quantitative comparison of cooling solutions from the computational point of view and consider not only issues of energy efficiency, but also computational density, investment, and maintenance costs. To the best of our knowledge, our paper is the first work that compares air and liquid cooling for data centers quantitatively and comprehensively.

Immersion cooling in practice

Although air cooling is the dominant solution in data centers, several large companies are adopting the immersion cooling technology and a number of startups have appeared offering innovative systems. At the time of writing, *Microsoft* has announced its first liquid immersion-cooled data center in Washington, USA (Weston 2021). This data center

is used for cloud-based communication platforms such as Microsoft Teams. *Alibaba* also uses single-phase immersion cooling (1PIC) tanks in its data centers. They have shown that using immersion cooling reduces the total power consumption by 36% and helps to achieve a PUE of 1.07 (Zhong 2019). Another example comes from the *BitFury* group that built a 40+ MW data center that comprises 160 tanks running with a PUE of 1.02 using two-phase immersion cooling (2PIC) (two-phase immersion cooling a revolution 2015).

Furthermore, some companies are offering immersion-cooled server systems. *Asperitas*, for example, is a Dutch company located in Amsterdam that offers complete liquid immersion cooling solutions to its customers. Their server enclosures operate with single-phase immersion cooling and natural convection, thus avoiding the use of any mechanical parts. The immersion-cooled servers of *Asperitas* are insulated; this is done to capture all heat produced by the servers in the fluid and to allow for maximum waste heat reutilization. Each of their AIC24 enclosures can contain up to 48 servers or 288 GPUs with a footprint of only 60 cm x 120 cm (Asperitas 2021).

An alternative interesting approach was that proposed by the Dutch company *Nerdalize*. The idea was to offer a distributed data center by displacing servers in the residential buildings (Ngoko et al. 2018). The immersion cooled servers would exchange heat with water that was then used for indoor heating and hot tap water. When the energy savings of (water) heating are taken into account, the PUE of such a system would be less than 1.0. The company deployed several servers before bankruptcy in 2018. Interestingly, the company which restarted *Nerdalize*, *LeafCloud*, decided to opt for air cooling, mostly due to its lower maintenance costs and IT failure rates.

Another company offering liquid immersion cooling enclosures in various sizes is *Submer* (Submer 2021). All of *Submer*'s products use single-phase immersion cooling and are ranging from cabinet-sized enclosures called microPod, up to megaPod, which are set up inside shipping containers. The microPods are capable of cooling 5 kW of components even in direct sunlight which makes them suitable for companies who want to cool their in-house equipment efficiently. On the other hand, megaPods are targeted for higher computing powers. They can be put in almost any place since there is only electricity and network connection needed. For example, it would be possible to install a megaPod onto or near a building which then supplies the building with heat. Similar to *Nerdalize* and *LeafCloud*, their products are excellent for waste heat reutilization.

Efficiency, density, cost, and maintenance

Cloud providers are generally profit-oriented. This means that data centers are implemented and optimized for maximum computing output using the least amount of investment and running costs. While some operators make decisions with sustainability in mind, the majority of them choose the most cost-efficient solutions.

In this section, we compare air cooling and immersion cooling solutions on several dimensions including computing efficiency, computing density, power density, cost, and maintenance. From the perspective of a data center operator, there is a clear link between each of these factors and the profit margin. The efficiency of the cooling method affects each of the mentioned factors, in turn reinforcing the importance of choosing the appropriate cooling solution.

Computing efficiency

The most popular metric for measuring the energy efficiency of a data center is the PUE (Reddy et al. 2017). PUE, originally proposed in 2006, and standardized as ISO/IEC 30134-2:2016 in 2016. The definition of PUE is given in Eq. 1.

$$PUE = \frac{\text{total energy}}{\text{IT energy}} \geq 1 \tag{1}$$

Since the IT energy consumption is included in the total energy, the value of PUE will typically be greater or equal to one. Several studies report on the PUE of data centers with specific installations of immersion and air cooling solutions, these are shown in Table 1.

For air-cooled data centers, reported PUE values range from 1.1 to 2.9 (Matsuoka et al. 2017; Shah 2018; McNevin 2013; Miller 2014). Values close to 1.1 can only be achieved by hyper-scale data center facilities, which are especially optimized for efficient cooling. For example, the state-of-the-art air-cooled data centers of Google have the PUE of 1.12 (Miller 2014). This means that in these data centers, 89% of the total energy is consumed by IT equipment. However, high efficient data centers with state-of-the-art design make up for only a small portion of data centers worldwide (Jones 2018).

In addition, average air-cooled data centers have a much higher PUE compared to the most efficient ones. The average PUE has been reported to be between 2.2 and 2.61 for the data centers in Singapore, Japan, Hong Kong, and Australia. This shows a significant difference between state-of-the-art hyper-scale and average data centers (McNevin 2013). Moreover, as Fig. 4 shows, the progress in improving the PUE of data centers worldwide has stalled. The average annual PUE in 2022 was 1.55, which is consistent with the trend observed in recent years and improvement slowed considerably since 2014 (Davis et al. 2022).

Studies regarding immersion cooling data centers have reported consistently better results. For these, the PUE falls in the range 1.02 to 1.04 (Matsuoka et al. 2017; An et al. 2018; Eiland et al. 2014; Chandrasekaran et al. 2017; Shah 2018). The PUE of 1.02 seems to be the sweet spot for immersion cooling. Table 1 shows how the various studies agree on the fact that PUE values around 1.02 are achievable with immersion cooling. With a PUE of 1.02, about 98% of the energy consumed by the data center goes to the IT equipment. This is close to the perfect efficiency. The maximum reported PUE for immersion-cooled data centers is 1.17 (Eiland et al. 2014). This case relates to an experiment aiming

Table 1 PUE values reported in the literature

| Source | Reported PUE Air cooling | Reported PUE Immersion cooling |
|------------------------------|-----------------------------|-----------------------------------|
| (Matsuoka et al. 2017) | 1.1 | < 1.04 |
| (An et al. 2018) | | 1.02 |
| (Eiland et al. 2014) | | 1.03–1.17 |
| (Chandrasekaran et al. 2017) | | 1.02–1.03 |
| (Shah 2018) | 1.7–2.9 | 1.02–1.03 |
| (Miller 2014) | 1.12, 1.18 | |
| (McNevin 2013) | 2.2–2.61 | |

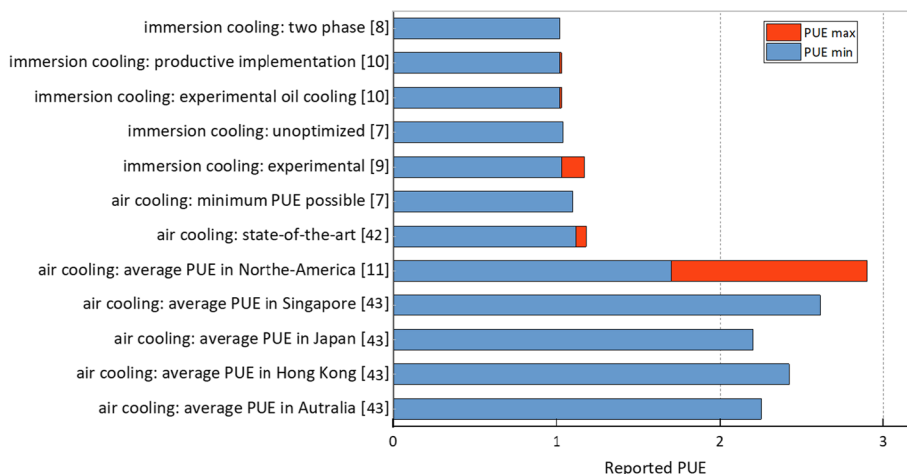


Fig. 3 PUE values based on the cooling technology

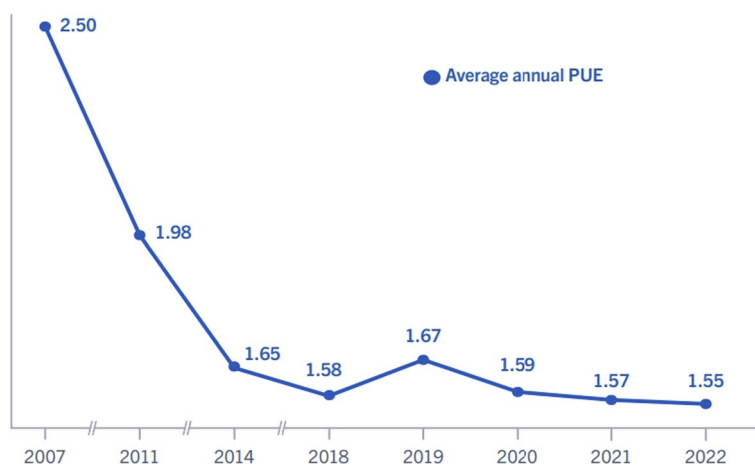


Fig. 4 The annual average PUE for data centers worldwide (Davis et al. 2022)

at maximum cooling capacity without efficiency in mind. For further illustration, Fig. 3 show a plot of the values from Table 1 with respect to the type of technology used for cooling.

While PUE is a useful metric for evaluating the efficiency of a data center over time, it is not equally suitable for comparing data centers with one another. Furthermore, it does not give direct information about the computing efficiency of the data center. In Reddy et al. (2017), we reviewed metrics for data centers and explored refinements and alternatives to PUE. In first approximation, PUE just offers a reasonable indicator of data center efficiency. From a computing perspective, it is also desirable to have a value representing the computing power in relation to energy consumption. Computing performance is traditionally measured by running a set of operations (a benchmark) on a device and measuring the completion time. An example of such a performance benchmark is measuring the number of floating point operations per second.

To compare the efficiency of different systems or facilities, one measures the computing performance of running a benchmark in the average/peak power of the system.

This results in a metric of floating-point operations per second per watt (FLOPS/W), a metric capable of showing the performance in relation to the energy consumption.

In a data center, in addition to the servers for running the operations, we have non-computational equipment such as the cooling system which consumes energy. Therefore, we consider an alternative metric for computing efficiency as follows.

$$\eta = \frac{\text{computing performance}}{\text{power}} (\text{FLOPS/W}) \quad (2)$$

The computing performance and the power of IT equipment are relevant for the calculation of FLOPS/W; however, they do not change the overall efficiency while changing the cooling solution. This is because changing the cooling system does not influence the computing performance and the IT equipment's power consumption. Typically, when modifying the cooling systems, the IT energy remains constant while the total energy changes. To calculate the improvement in FLOPS/W by changing the cooling method, the knowledge of PUE and FLOPS/W can be combined for an improved computing efficiency metric, presented as Eq. (3).

$$\eta_{\text{data center}} = \frac{\frac{b_0}{b_1}}{\text{IT power}} \cdot \frac{1}{\text{PUE}} \quad (3)$$

where b_0 and b_1 stand for the number of benchmark operations and benchmark time, respectively. The inverse of PUE represents the part of the energy that has been used by IT equipment. For example, a data center with a PUE of 1.5 uses two-third of its energy for IT equipment. By changing the cooling method, the value of b_0 , b_1 , and IT power remain the same. This means that overall computing efficiency depends directly on the fraction of power used for IT equipment. Furthermore, percentage differences in FLOPS/W can be calculated having only the PUE.

Knowing this, the expected improvement in the data center's overall computing efficiency by switching from air cooling to immersion cooling can be calculated by comparing $1/\text{PUE}$ values. For example, when migrating from best practice air cooling ($\text{PUE} = 1.12$) to immersion cooling ($\text{PUE} = 1.02$), the overall computing efficiency is increased by 9.8%. While this is a significant increase that reduces the operational expenditures, migrating from standard air cooling ($\text{PUE} = 2$) to immersion cooling ($\text{PUE} = 1.02$) increases the computing efficiency by 96.1%. Therefore, such a conversion would reduce the energy consumption by half. The increase and decrease in computing efficiency when changing cooling method is shown in Table 2.

Let us present a numerical example based on one of the most powerful server CPUs currently on the market, the AMD EPYC 7742. This processor runs with 225-watt maximum power consumption and achieves around 3.48 TeraFLOPS (Trader 2019). Therefore, the processor can calculate 15.5 GigaFLOPS per watt at its peak performance:

$$\eta_{\text{EPYC7742}} = \frac{3.48 \text{ TFLOPS}}{225\text{W}} = 15.5 (\text{GFLOPS/W}) \quad (4)$$

The central processing unit that calculates the FLOPS, is not the only power consuming component of a server. Therefore, to calculate the server computing efficiency, one

Table 2 Efficiency improvement by switching cooling technique

| From/To | Standard Air cooling | State of the art Air cooling | Immersion Cooling |
|------------------------------|----------------------|------------------------------|-------------------|
| Standard air cooling | – | 78.6% | 96.1% |
| State of the art air cooling | – 44% | – | 9.8% |
| Immersion cooling | – 49% | – 8.9% | – |

needs to know the proportion of power consumed by the processor in relation to the server's total power. The computing efficiency of a server is calculated by:

$$\eta_{\text{server}} = \eta_{\text{processor}} \cdot p_{\text{processor}} \quad (5)$$

where $p_{\text{processor}}$ stands for the proportion of power consumed by the processor in relation to the total server consumption. Various works have reported different power breakdowns for the components of a server. In accordance with the results presented in Gill and Buyya (2018), we assume that 50% of the server's power is consumed by its processor. From Eq. 5, the computing efficiency of the server amounts to 7.73 GFLOPS/W.

To calculate the computing efficiency of the IT equipment, we consider the power of all IT equipment including storage and network facility. In this way, the computing efficiency of IT equipment is calculated as:

$$\eta_{\text{IT equipment}} = \eta_{\text{server}} \cdot p_{\text{server}} \quad (6)$$

where p_{server} stands for the proportion of the server's power in relation to the total IT power. Indeed, we use p_{server} in Eq. 6, to take into account the energy consumption of non-computational IT equipment (storage and network) in calculating the computing efficiency of IT equipment. The contribution of the servers, storage, and network facility in the total power consumption of a data center has been reported in several works (Dayarathna et al. 2015; Shehabi et al. 2016). For our numerical example, we assume the average of reported values, that is 77%. Therefore, $\eta_{\text{IT equipment}}$ is 5.95 GFLOPS/W.

Finally, the computing efficiency of the data center is calculated by:

$$\eta_{\text{data center}} = \eta_{\text{IT equipment}} \cdot \frac{1}{\text{PUE}} \quad (7)$$

By Eq. 7, the computing efficiency of a data center cooled by standard air cooling, state-of-the-art air cooling, and immersion cooling is calculated as 2.98, 5.32, and 5.84, respectively. The computing efficiency of an immersion-cooled data center is almost 9% higher than a state-of-the-art air-cooled hyperscale data center. For the average data center, the switch to immersion cooling offers even more improvement. Operators can decrease the power consumption by about 50% without any decrease in computing performance.

Computing and power density

Computing density is defined as the amount of computation that a system can offer in relation to its size. The metric used to determine the computing density of a data center is FLOPS/m³. This metric considers everything: from the servers to the power management system and up to the cooling equipment. Comparing air-cooled and

immersion-cooled in these terms also shows significant differences. The required space for immersion cooling has been reported to be about one third of traditional air cooling method (Matsuoka et al. 2017).

The increase in computing density by immersion cooling comes from various factors. In immersion cooling, there is no space needed in between the servers and racks for airflow. Tubs can stand right next to each other, and the only limitation is the accessibility of the service crew. In addition, there is significantly less cooling equipment needed (only piping), and no raised floors or air vents are required (Matsuoka et al. 2017). According to the literature, with immersion cooling the density increases at both the rack and facility levels (Tuma 2010). Furthermore, many air-cooled data centers trade their computing density for efficiency and reliability (Miller 2014). The cooling units of these centers are often bigger than they need to be.

Power density is another important factor for data center operators. Power density must be low in air-cooled data centers, otherwise, either fan speeds need to be increased or air temperature needs to be lowered, where both solutions lower the energy efficiency. According to the research, rack-level power density of air-cooled data centres is about 0.018–0.028 kW/l (Kheirabadi and Groulx 2016), while the density of immersion-cooled data centres is between 0.045 kW/l (Matsuoka et al. 2017) and 0.23 kW/l (Gess et al. 2014). Even more, 4 kW/l has been reported to be possible in immersion-cooled data centers with enough coolant flow and a compact IT design (Tuma 2010, 2010). Table 3 presents the data gathered from various references on the power density of air-cooled and immersion-cooled data centers. Since power densities are presented in kW per rack for air-cooled data centers and for immersion cooling there is no typical rack size, the values need to be converted to kW/l. This metric is the most meaningful one when the whole facility is considered.

As the reported numbers show, the power density with immersion cooling is almost six times higher than air cooling. Therefore, power density can be another strong motivation for data center operators to consider immersion cooling for their future projects, specifically for running their high-performance workloads. Going forward, computing is driving devices into higher power consumption and CPUs and GPUs capable of consuming more power are expected to be manufactured. Given this trend, air cooling can not provide sufficient cooling capacity, while maintaining high power densities.

Table 3 Power densities reported in the literature

| Source | Air cooling Density | Immersion cooling Density | Normalized to kW/l |
|------------------------------|---------------------|---------------------------|--------------------|
| (Matsuoka et al. 2017) | | 14 kW/bathtub | 0.045 |
| (Gess et al. 2014) | | 400 W/bathtub | 0.23 |
| (An et al. 2018) | | 250 kW/rack | 0.14 |
| (Tuma 2010) | | 4 kW/l | 4 |
| (Kheirabadi and Groulx 2016) | 33–50 kW/rack | | 0.018–0.028 |
| (Smolaks 2019) | 40 kW/rack | | 0.022 |

Cost

An important factor to consider when adopting immersion cooling is capital and operational expenditures. The capital expenditures include the sealed chassis for immersing the IT equipment, dielectric fluids, as well as pumps and tubing. Similar to air-cooled data centers, the operational expenditures include electricity, staff, network connection fees, as well as supporting and maintaining the IT equipment.

According to Bunger et al. (2020), for air-cooled data centers with a power density of 10 kW per rack, the capital expenditures are \$7.02 per watt. For a liquid-cooled data center with a similar power density, the cost is reduced slightly to \$6.98 per watt. The benefit of liquid cooling is also that much higher power densities can be achieved. Assuming a power density of 40 kW per rack, the expenditures are further reduced to \$6.02 per Watt. Furthermore, when it comes to the operational expenditures, a reduction between 9–20% is expected with regards to the energy cost due to the absence of fans. This is in agreement with the results of Neudorfer et al. who state that a 5–10% reduction of the IT energy consumption is expected due to the avoidance of internal fans (Neudorfer et al. 2016). The need for fans is also absent for the entire facility, reducing the total energy costs by an amount in the range of 15–25%. Another point to consider is that the dielectric fluids present in the sealed chassis can be used virtually indefinitely, assuming some form of filtration is present.

Day et al. highlight that when building a new data center and optimizing it for liquid cooling from the ground up, capital expenditure savings can be achieved over air-cooled data centers (Day et al. 2019). On the contrary, retrofitting an air-cooled data center with liquid cooling can result in higher costs. An important exception is retrofitting an air-cooled data center with limited floor space and power capacity. In this case, the increased power density possible with liquid cooling, as well as a reduction in energy consumption, can address both the space and power limitations with one solution.

System maintenance

The operational costs of a data center include system maintenance. The cooling method influences the components' environmental conditions and consequently affects the number of and time to failures and overall equipment lifetime and reliability.

The number of maintenance requests caused by failures has been reported to be almost 6.6% higher for an immersion-cooled data center compared to a traditional air-cooled counterpart (Coles and Herrlin 2016). The higher number of maintenance requests associated with immersion cooling impose additional operating costs. In addition, higher failure rates can degrade the components' lifetime, but immersion cooling can compensate for this degradation with lower junction temperatures (Jalili et al. 2021). Besides the number of maintenance requests, the maintenance procedure is more challenging with immersion cooling. In this case, immersed IT equipment is removed by opening the lid and lifting the equipment out of the tank. This can lead to liquid evaporation.

According to Villa (2020) and Alibaba (2018), the maintenance overheads and reliability concerns, as well as the leaks and spills, are one of the main contributors to the low adoption of immersion cooling. Indeed, the enclosure that the racks are immersed in must be sealed perfectly to avoid liquid evaporation/losses. Complete enclosure sealing would mitigate the problem but is not practical. The high number of maintenance

requests, compared to the air cooling system, cause access issues and adds operating costs related to compensating the fluid losses (Varma 2019).

The operating costs imposed by the liquid loss have been evaluated in Coles and Herrlin (2016). The cost of the lost liquid divided by the cost of the IT equipment's energy usage has been reported to be 4.68 for a specific implementation. This number shows that the maintenance overhead is a significant drawback associated with immersion cooling. It can even mitigate the energy efficiency improvement of immersion cooling. However, it should be noted that the implementation characteristics including the liquid price and the electricity costs affect the reported value. In Coles and Herrlin (2016), these values are set at 75 \$/liter and 0.09 \$/kWh, respectively. One-phase immersion cooling is reported to have fewer maintenance needs than two-phase immersion cooling (Varma 2019).

However, there are several reliability improvement potentials with immersion cooling. With immersion cooling, the servers operate continuously at full power without risk of failure due to overheating and temperature swings. In addition, the negative effects of dust and low air quality are eliminated by immersion cooling. Moreover, by removing fans, the total number of maintenance requests decreases. These facts may lead to better reliability, however, as pointed out in Shah et al. (2016), comprehensive studies of reliability and availability are required.

Finally, considering the higher maintenance costs due to liquid loss and increased number of IT failures as well as not having enough information on the real implementations of immersion cooling in real data centers, maintenance challenges and reliability concerns are considered as the downside of this method when compared to traditional and well established air cooling.

Conclusions

Air cooling is the traditional solution for addressing the heat dissipation problem of data centers. High energy consumption and low cooling capacity—and consequently limited power density—are the main challenges associated with air cooling. Immersion cooling is emerging as a novel method with many advantages in terms of efficiency, density, and cost. In the present work, we provide a quantitative comparison of these two approaches and overview the results presented in the literature regarding the two alternatives. While most data centers around the world are relying on air cooling, immersion cooling is recognized as a potential alternative and several cloud providers have already constructed their immersion-cooled data centers. The key findings of the present work are listed next.

First, based on the PUEs reported in the literature, we conclude that the immersion cooling method consumes less energy and has a higher computing efficiency. Our analysis shows that a typical immersion-cooled data center consumes almost 50% less energy compared to its air-cooled counterpart. Second, immersion-cooled data centers allow for much more compact designs, more than three times the density of their air-cooled counterparts. In liquid immersion-cooled data centers, there is no trade-off between efficiency and density. Conversely, air-cooled data centers can only be either-or. Immersion-cooled data centers can be placed in ordinary spaces having lower requirements and the need for additional equipment. Third, while research on immersion cooling is

mostly targeting efficiency, the aspect of power density should not be overlooked. The increase of computing power in a specific volume is even more important than the efficiency improvement. The most conservative figures for immersion cooling are about double density in kW/l on rack-level compared to the maximum possible in air-cooled data centers. Resources estimating 4 kW/l illustrate what is potentially possible with IT equipment optimized for liquid immersion cooling. This high-density capability makes immersion cooling the first-choice solution for running high-performance workloads. In addition, it allows manufacturing the CPUs and GPUs with higher frequency and higher power consumption. Fourth, the capital expenditure for an air-cooled data center with a power density of 10 kW per rack is about 4% higher compared to its immersion-cooled counterpart. In addition, a 9–20% operational expenditures reduction is expected with immersion cooling. Fifth and final, the main downside of immersion cooling is the challenges related to the maintenance and reliability concerns. This is a reasonable explanation for why even with significant advantages in efficiency and density, not many companies have switched their cooling solution to immersion cooling, yet.

In general, immersion cooling appears to be a solution that should be strongly considered when designing a new data center and we expect its adoption to grow. At the same time, retrofitting any air-cooled data center with liquid cooling does not seem convenient in most cases, if feasible at all. Considering the reliability concerns and maintenance challenges and uncertainties, immersion cooling might not be useful for small and average power densities. On the contrary, immersion cooling appears to be the best solution for high power densities.

Abbreviations

| | |
|---------|---|
| PUE | Power usage efficiency |
| CRAC | Computer room air conditioning |
| CRAH | Computer room air handling |
| HACS | Hot aisle containment system |
| CACS | Cold aisle containment system |
| 1PIC | Single-phase immersion cooling |
| 2PIC | Two-phase immersion cooling |
| FLOPS/W | Floating-point operations per second per watt |

Acknowledgements

Not applicable.

Author contributions

YB and MA conceived of the presented idea. KH, BS, and YB reviewed the references and developed the analytical methods. KH and BS prepared the initial manuscript and MA helped to improve the writing and the structure. All authors discussed the results and contributed to the final manuscript.

Funding

The presented research is funded by the Netherlands Organisation for Scientific Research (NWO) in the framework of the Indo-Dutch Science Industry Collaboration programme with project NextGenSmart DC (629.002.102).

Availability of data and materials

Not applicable.

Declarations

Ethics approval and consent to participate

Not applicable.

Consent for publication

Not applicable.

Competing interests

The authors declare that they have no competing interests.

Received: 1 July 2023 Accepted: 31 July 2023

Published online: 09 August 2023

References

- Alibaba (2018) Immersion cooling for green computing. <https://www.opencompute.org/files/Immersion-Cooling-for-Green-Computing-V1.0.pdf>
- An X, Arora M, Huang W, Brantley WC, Greathouse JL (2018) 3d numerical analysis of two-phase immersion cooling for electronic components. In: 2018 17th IEEE Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm), pp. 609–614. IEEE
- Asperitas (2021) Immersion cooling solutions for datacenters. <https://www.asperitas.com/>
- Bunger R, Torell W, Avelar V (2020) Capital cost analysis is of immersive liquid-cooled vs. air-cooled large data centers. Schneider Electric White Paper 282
- Chandrasekaran S, Gess J, Bhavnani S (2017) Effect of subcooling, flow rate and surface characteristics on flow boiling performance of high performance liquid cooled immersion server model. In: 2017 16th IEEE Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm), pp. 905–912. IEEE
- Coles H, Herrlin M (2016) Immersion cooling of electronics in data installations. Technical Report LBNL-1005666, Berkeley National Laboratories
- Davis J, Bizo D, Lawrence A, Rogers O, Smolaks M, Simon L, Donnellan D (2022) Uptime institute global data center survey 2022. Uptime Institute Intelligence
- Day T, Lin P, Bunger R (2019) Liquid cooling technologies for data centers and edge applications. Schneider Electric White Paper 265
- Dayarathna M, Wen Y, Fan R (2015) Data center energy consumption modeling: a survey. *IEEE Commun Surv Tutor* 18(1):732–794
- Dennard RH, Gaensslen FH, Yu H-N, Rideout VL, Bassous E, LeBlanc AR (1974) Design of ion-implanted mosfet's with very small physical dimensions. *IEEE J Solid State Circuits* 9(5):256–268
- Dunlap K, Rasmussen N (2006) The advantages of row and rack-oriented cooling architectures for data centers. *American Power Conversion*
- Ebrahimi K, Jones GF, Fleischer AS (2014) A review of data center cooling technology, operating conditions and the corresponding low-grade waste heat recovery opportunities. *Renew Sustain Energy Rev* 31:622–638
- Eiland R, Fernandes J, Vallejo M, Agonafer D, Mulay V (2014) Flow rate and inlet temperature considerations for direct immersion of a single server in mineral oil. In: Fourteenth Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm), pp. 706–714. IEEE
- Evans T (2012) The different technologies for cooling data centers. APC white paper 59
- Fan Y, Winkel C, Kulkarni D, Tian W (2018) Analytical design methodology for liquid based cooling solution for high tdp cpus. In: 2018 17th IEEE Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm), pp. 582–586. IEEE
- Geng H (2014) *Data center handbook*. Wiley, New Jersey
- Gess J, Bhavnani S, Ramakrishnan B, Johnson RW, Harris D, Knight R, Hamilton M, Ellis C (2014) Investigation and characterization of a high performance, small form factor, modular liquid immersion cooled server model. In: 2014 Semiconductor Thermal Measurement and Management Symposium (SEMI-THERM), pp. 8–16. IEEE
- Gill SS, Buyya R (2018) A taxonomy and future directions for sustainable cloud computing: 360 degree view. *ACM Comput Surv (CSUR)* 51(5):1–33
- Haghshenas K, Taheri S, Goudarzi M, Mohammadi S (2020) Infrastructure aware heterogeneous-workloads scheduling for data center energy cost minimization. *IEEE Trans Cloud Comput* 10(2):972–983
- Jalili M, Manousakis I, Goiri I, Misra PA, Raniwala A, Alissa H, Ramakrishnan B, Tuma P, Belady C, Fontoura M, et al (2021) Cost-efficient overclocking in immersion-cooled datacenters. In: *Proceedings of the International Symposium on Computer Architecture (ISCA'21)*
- Jones N (2018) The information factories. *Nature* 561(7722):163–6
- Kanbur BB, Wu C, Fan S, Tong W, Duan F (2020) Two-phase liquid-immersion data center cooling system: experimental performance and thermoeconomic analysis. *Int J Refrig* 118:290–301
- Kheirabadi AC, Groulx D (2016) Cooling of server electronics: a design review of existing technology. *Appl Therm Eng* 105:622–638
- Kuncoro I, Pambudi N, Biddinika M, Widiastuti I, Hijriawan M, Wibowo K (2019) Immersion cooling as the next technology for data center cooling: a review. *J Phys Conf Ser* 1402:044057
- Intel (2017) New 2nd generation Intel Xeon salable processor. <https://www.intel.com/content/dam/www/public/us/en/documents/guides/2nd-gen-xeon-sp-transition-guide-final.pdf>
- Li Z, Kandlikar SG (2015) Current status and future trends in data-center cooling technologies. *Heat Transf Eng* 36(6):523–538
- Matsuoka M, Matsuda K, Kubo H (2017) Liquid immersion cooling technology with natural convection in data center. In: 2017 IEEE 6th International Conference on Cloud Networking (CloudNet), pp. 1–7. IEEE
- McNevin A (2013) APAC data center survey reveals high PUE figures across the region. <https://www.datacenterdynamics.com/>
- Miller R (2014) Inside SuperNAP 8: switch's tier iv data fortress. <https://www.datacenterknowledge.com/archives/2014/02/11/inside-supernap-8-switchs-tier-iv-data-fortress/>
- Nadjahi C, Louahlia H, Lemasson S (2018) A review of thermal management and innovative cooling strategies for data center. *Sustain Comput Inform Syst* 19:14–28
- Neudorfer J, Ellsworth M, Kulkarni D, Zien H (2016) Liquid cooling technology update. The Green Grid White Paper 70

- Ngoko Y, Saintherant N, Cerin C, Trystram D (2018) How future buildings could redefine distributed computing. In: 2018 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW), pp. 1232–1240. IEEE
- Ni J, Bai X (2017) A review of air conditioning energy performance in data centers. *Renew Sustain Energy Rev* 67:625–640
- Niemann J, Brown K, Avelar V (2011) Impact of hot and cold aisle containment on data center temperature and efficiency. Schneider Electric Data Center Science Center, White Paper 135:1–14
- Patankar SV (2010) Airflow and cooling in a data center. *J Heat Transfer* 132(7):073001
- Ramdas S, Rajmane P, Chauhan T, Misrak A, Agonafer D (2019) Impact of immersion cooling on thermo-mechanical properties of pcb's and reliability of electronic packages. In: International Electronic Packaging Technical Conference and Exhibition, vol. 59322. American Society of Mechanical Engineers
- Rasmussen N (2017) The different types of air distribution for it environments. Schneider Electric
- Reddy VD, Setz B, Rao GSV, Gangadharan G, Aiello M (2017) Metrics for sustainable data centers. *IEEE Trans Sustain Comput* 2(3):290–303
- Shah JM (2018) Characterizing contamination to expand ashrae envelope in airside economization and thermal and reliability in immersion cooling of data centers. PhD thesis, The University of Texas at Arlington
- Shah JM, Eiland R, Siddarth A, Agonafer D (2016) Effects of mineral oil immersion cooling on it equipment reliability and reliability enhancements to data center operations. In: 2016 15th IEEE Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm), pp. 316–325. IEEE
- Shehabi A, Smith S, Sartor D, Brown R, Herrlin M, Koomey J, Masanet E, Horner N, Azevedo I, Lintner W (2016) United states data center energy usage report
- Shinde PA, Bansode PV, Saini S, Kasukurthy R, Chauhan T, Shah JM, Agonafer D (2019) Experimental analysis for optimization of thermal performance of a server in single phase immersion cooling. In: International Electronic Packaging Technical Conference and Exhibition, vol. 59322. American Society of Mechanical Engineers
- Smolaks M (2019) Power density—the real benchmark of a data centre. <https://virtusdatacentres.com/item/389-power-density-the-real-benchmark-of-a-data-centre>
- Son J, Dastjerdi AV, Calheiros RN, Buyya R (2017) Sla-aware and energy-efficient dynamic overbooking in sdn-based cloud data centers. *IEEE Trans. Sustain. Comput.* 2(2):76–89
- Submer (2021) Datacenters that make sense. <https://submer.com/>
- Sun Y, Agostini NB, Dong S, Kaeli D (2019) Summarizing cpu and gpu design trends with product data. arXiv preprint <http://arxiv.org/abs/1911.11313>
- Trader T (2019) AMD Launches Epyc Rome, First 7nm CPU. <https://www.hpcwire.com/2019/08/08/amd-launches-epyc-rome-first-7nm-cpu/>
- Tuma P (2010) Open bath immersion cooling in data centers: a new twist on an old idea. *Electronics Cooling* 2010:10
- Tuma PE (2010) The merits of open bath immersion cooling of datacom equipment. In: 2010 26th Annual IEEE Semiconductor Thermal Measurement and Management Symposium (SEMI-THERM), IEEE, pp 123–131
- Two-phase immersion cooling a revolution in data center efficiency (2015). 3M™ Novec™ Engineered Fluids
- Varma D (2019) Two-phase versus single-phase immersion cooling. <https://www.grcooling.com/wp-content/uploads/2020/03/grc-blog-library-tech-comparison-%E2%80%94-two-vs-single-phase-immersion-cooling.pdf>
- Villa H (2020) Liquid cooling vs. immersion cooling deployment. <https://blog.rittal.us/liquid-cooling-vs-immersion-cooling-deployment>
- Voices of the Industry (2020) Data centers feeling the heat! The history and future of data center cooling. <https://datacenterfrontier.com/history-future-data-center-cooling/>
- Watson B, Venkiteswaran VK (2017) Universal cooling of data centres: A cfd analysis. *Energy Procedia* 142:2711–2720. <https://doi.org/10.1016/j.egypro.2017.12.215>. Proceedings of the 9th International Conference on Applied Energy
- Weston S (2021) Microsoft is submerging servers in boiling liquid to prevent teams outages. <https://www.itpro.co.uk/server-storage/datacentr/359129/microsoft-submerges-servers-in-boiling-liquid-to-prevent-teams-outages>
- Zhong Y (2019) A large scale deployment experience using immersion cooling in datacenter. Alibaba Group: Open Compute Project Summit

Publisher's Note

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Kawsar Haghshenas is an assistant professor in Distributed Systems at the University of Groningen. Before that, she was a post-doctoral researcher in the Smart Energy Systems (SES) laboratory at the University of Stuttgart. She received her Ph.D. degree in Computer Science from University of Tehran in 2020, M.Sc. degree in Computer Engineering from Sharif University of Technology in 2014, and B.Sc. degree in Computer Engineering from K.N.Toosi University of Technology in 2012. She was visiting the Embedded Systems Laboratory (ESL) at EPFL University from 2017 to 2018. Her research interests include system level energy optimization, machine learning based resource management, and workload scheduling in the area of datacenters and cloud computing.

Brian Setz is the Lead Software Developer at the Digital Lab of the Computing Science department at the University of Groningen. He holds a PhD in Computing Science from the University of Stuttgart. His research interests include: Internet of Things, Green Computing, Cloud Computing, and Energy Efficient Data Centers.

Yannis Blosch studied his B.Sc. degree in Computer Science at the Institute of Architecture of Application Systems (IAAS) of the University of Stuttgart.

Marco Aiello is professor of Computer Science and Head of the Service Computing Department at the University of Stuttgart, Germany. He is also an elected member of the European Academy of Sciences and Arts and honorary professor of Distributed Systems at the University of Groningen, The Netherlands, where he was a faculty member from 2006 till 2018. He holds a PhD in Logic from the University of Amsterdam, the Habilitation in Applied Informatics from TU Wien, and a master degree in Engineering from La Sapienza University of Rome. His research interests are in Service Computing, Smart Grids, Smart Energy Systems, and Data Centers.

Submit your manuscript to a SpringerOpen[®] journal and benefit from:

- ▶ Convenient online submission
- ▶ Rigorous peer review
- ▶ Open access: articles freely available online
- ▶ High visibility within the field
- ▶ Retaining the copyright to your article

Submit your next manuscript at ▶ [springeropen.com](https://www.springeropen.com)
