22nd International Conference on Sustainable Energy Technologies
Kuala Lumpur, Malaysia, 13 to 15 August 2026 · ISBN 978-1-995850-00-9 · ISSN 2978-4387

CONFERENCE PAPER · LOW-CARBON BUILDINGS

Parametric Characterisation and Operating Boundaries of Air- and Liquid-Cooling Architectures in High-Density Data Centres

Xinyi Wang[1], Lu Wang[1], Haojiang Ran[1], Tiexiao Xu[1], Bo Zhang[1], Zhen Li[1]*

[1] Tsinghua University, Beijing, China

RECEIVED 21 May 2026 · REVISED 25 June 2026 · ACCEPTED 9 July 2026
PRESENTED 13 to 15 August 2026 · PUBLISHED 23 November 2026

ABSTRACT

The rapid growth of artificial intelligence has continuously increased chip and rack power densities, making thermal management an increasingly important factor in both the cooling capability and energy efficiency of data centres. Improving heat transfer efficiency can reduce the temperature difference required for heat rejection and potentially decrease the dependence on energy-intensive mechanical cooling. However, comparisons among air cooling, cold-plate cooling, and immersion cooling are often affected by simultaneous differences in coolant properties, package structures, heat-sink dimensions, and cooling-source temperatures. In this study, a unified chip-to-facility thermal resistance framework was developed to compare rack-level air cooling, single-phase lidded cold-plate cooling, single-phase Direct-to-Die cold-plate cooling, and single-phase immersion cooling. Under unified geometric conditions with a common fin height of 20 mm, the total thermal resistance of the liquid-cooling architectures was only 24.2%–37.5% of that of air cooling. The dominant thermal resistance also shifted substantially among different architectures. Air cooling was mainly limited by the heat sink, heat exchanger, and airflow warming and mixing processes, whereas thermal interface materials and heat spreading within the base plate became increasingly important for liquid cooling. Architecture-specific enhancement strategies were therefore introduced by improving thermal interface materials, heat-spreading structures, and heat-sink dimensions. At the adopted coolant-velocity modelling upper bounds and a facility-water inlet temperature of 10 °C, the second-generation configurations achieved maximum single-chip cooling capacities of 1.74, 4.43, 6.09, and 6.34 kW for air cooling, lidded cold-plate cooling, Direct-to-Die cold-plate cooling, and immersion cooling, respectively. The results demonstrated that different cooling architectures required different enhancement pathways. Reducing thermal resistance not only increased the achievable cooling capacity but also provided greater flexibility for integrating free-cooling sources and reducing reliance on mechanical refrigeration, supporting the development of more sustainable high-density data centres.

KEYWORDS Data centre cooling, Thermal resistance analysis, Liquid cooling, Cooling architectures, Heat dissipation capability

1. Introduction

The rapid growth of digitalisation, artificial intelligence (AI), and high-performance computing has driven an unprecedented expansion of data centre infrastructure worldwide (International Energy Agency 2025). Data centres consumed approximately 415 TWh of electricity worldwide in 2024, equivalent to around 1.5% of global electricity consumption (International Energy Agency 2025). The latest IEA update estimated that data-centre electricity consumption increased to approximately 485 TWh in 2025 and could reach around 950 TWh by 2030, accounting for approximately 3% of global electricity demand (International Energy Agency 2026). Cooling systems represent a significant

* Corresponding author. Email address: lizh@tsinghua.edu.cn (Zhen Li)

© 2026 The Author(s).
Published by Caravel Press.

doi.org/10.65582/set.2026.022

fraction of total data-centre energy demand, particularly in conventional air-cooled facilities (Zhao *et al.* 2023). Therefore, improving the energy efficiency of data-centre cooling systems has become a critical pathway towards sustainable computing infrastructure (Zhang *et al.* 2024).

The continuous increase in computing density, especially driven by advanced AI accelerators and high-performance processors, has further intensified thermal management challenges (Li *et al.* 2024). Conventional air cooling has benefited from its technological maturity, reliability, and compatibility with existing facilities; however, its limited heat transport capability restricts further increases in chip and rack power density (Azarifar *et al.* 2024). Consequently, various liquid-cooling technologies, including cold-plate cooling and immersion cooling, have received growing attention due to their superior heat transfer capability and potential for reducing cooling energy consumption (Kong *et al.* 2024; Huang *et al.* 2024). These technologies provide alternative heat-transfer pathways by replacing air-based heat removal with direct liquid heat transport or dielectric-fluid-based cooling.

However, comparisons among different cooling architectures are often based on nominal product specifications or reported system-level performance indicators, such as maximum cooling capacity or Power Usage Effectiveness (PUE) (Chang *et al.* 2024; Cai and Gou 2026). Such comparisons frequently combine variations in coolant properties, package structures, heat sink geometries, and cooling-source temperatures (Alkrush *et al.* 2024), making it difficult to determine the fundamental reasons behind the superiority of one architecture over another. More importantly, these comparisons cannot clearly identify which components should be prioritised for further optimisation. As highlighted in previous studies, a unified evaluation framework is necessary to distinguish the intrinsic advantages of cooling media and heat-transfer architectures from those introduced by technology-specific design optimisations (Park *et al.* 2026).

In this study, a unified chip-to-facility thermal analysis framework is established to systematically compare representative data centre cooling terminals, including air cooling, cold-plate liquid cooling, and immersion cooling. The analysis focuses on three fundamental questions: (1) where the dominant thermal resistances are located within different cooling pathways; (2) how the preferred enhancement strategies vary among cooling architectures; and (3) how the heat dissipation capability evolves at both single-chip and rack levels under different operating conditions. By providing a physics-based comparison under consistent boundary conditions, this work aims to reveal the inherent trade-offs among heat-transfer capability, optimisation potential, and deployment constraints, thereby providing guidance for future sustainable data centre cooling system design and technology selection.

2. Methodology

2.1. Thermal architectures and heat transfer pathways

A unified chip-to-facility thermal analysis framework was established to systematically compare the heat transfer characteristics of different data centre cooling architectures. Figure 1 illustrates the investigated cooling architectures, including air cooling, single-phase lidded cold-plate cooling, single-phase Direct-to-Die cold-plate cooling, and single-phase immersion cooling (Kalantarpour and Vafai 2024). To ensure a consistent comparison, identical chip heat sources, server layouts, rack configurations, and external cooling-source conditions were applied to all architectures, while only the heat transfer pathways between the chip and coolant were varied.

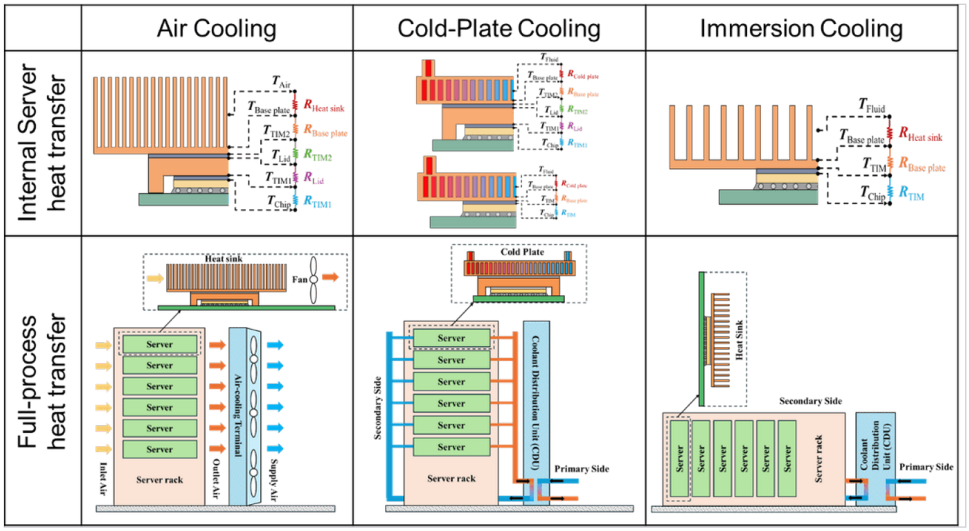

Figure 1. Schematic illustration of heat transfer pathways for different cooling architectures.

At the chip level, the generated heat was transferred through a series of thermal resistances before being removed by the cooling medium. For conventional lidded structures, heat passed sequentially through the first thermal interface material (TIM1), metallic lid, second thermal interface material (TIM2), heat sink base plate, and heat sink fins, followed by convection to the coolant. For Direct-to-Die cold-plate cooling, the lid and TIM2 layer were removed and the cold plate was connected to the chip through TIM1. For immersion cooling, the lid and TIM2 thicknesses were set to zero as a study-specific modelling assumption rather than as a universal feature of immersion-cooled hardware.

The main differences among the cooling architectures were associated with the terminal heat transfer processes. Air cooling used air as the cooling medium, and the airflow was driven by fans to remove heat from the heat sink surface. Cold-plate cooling adopted water as the working fluid, which was circulated through cooling channels by pumps to enhance heat transfer. Immersion cooling directly immersed the server in Daikin SS110 dielectric coolant, and the coolant circulation was driven by pumps to remove heat from the chip-level heat transfer structures.

After absorbing heat from the server, the secondary cooling medium transferred heat to the primary cooling water through a terminal secondary-fluid-to-primary-water heat exchanger. The cooled medium was then recirculated to complete the cooling cycle. Based on these unified heat transfer pathways, thermal resistance networks were constructed for different architectures, enabling a physics-based comparison of their heat transfer characteristics and optimisation potential.

2.2. Unified boundary conditions

To ensure a fair comparison among different cooling architectures, unified boundary conditions were applied to minimise the influence of differences in chip configuration, heat sink geometry, server layout, and cooling-source conditions. The evaluation framework considered the complete heat-transfer pathway from the chip to the facility cooling loop, including package-level thermal resistance, heat spreading structures, coolant flow conditions, and external cooling-source temperatures. The detailed parameters are summarised in Table 1 (Wang *et al.* 2024).

Table 1. Unified boundary conditions for thermal comparison of different cooling architectures.

Category	Parameter	Value
Chip heat source	Chip size	30 mm × 30 mm
	Maximum chip-surface hotspot temperature	85 °C
Package structure	TIM thickness	0.1 mm
	Lid thickness	1 mm (0 mm for Direct-to-Die and immersion cooling)
	TIM2 thickness	0.1 mm (0 mm for Direct-to-Die and immersion cooling)
Heat sink structure	Base plate size	200 mm × 100 mm × 3 mm
	Base plate material	Copper, 400 W/(m·K)
	Fin spacing	2 mm
	Fin thickness	0.1 mm
Server and rack configuration	GPU configuration	8 GPUs per server
	Rack specification	42U rack with 40U available space
Cooling medium	Air cooling	Air
	Cold-plate cooling	Water
	Immersion cooling	Daikin SS110 dielectric coolant
Flow limitation	Air-velocity modelling upper bound	5 m/s
	Liquid-velocity modelling upper bound	0.5 m/s
Cooling source	Facility-water inlet temperature	10 °C / 45 °C

The same chip size, allowable maximum chip-surface hotspot temperature, heat sink geometry, and server-level configuration were maintained for all architectures. For Direct-to-Die cold-plate cooling, the lid and TIM2 layer were removed and their thicknesses were set to zero (Schultz and Parida 2024). For immersion cooling, the same zero-thickness treatment was retained as a study-specific modelling assumption.

The coolant velocity limits were determined considering both thermal performance and practical operation requirements. An air-side superficial velocity of 5 m/s was adopted as the modelling upper bound, considering fan power consumption, airflow resistance, and noise constraints in data-centre environments (Chacón *et al.* 2024). A liquid-side superficial velocity of 0.5 m/s was adopted as the modelling upper bound to account for pumping power, pressure drop, and flow stability rather than as a universal equipment limit (Yang *et al.* 2025). Two facility-water inlet temperatures of 10 °C and 45 °C were selected to represent a low-temperature chilled-water condition and a high-temperature liquid-cooling/free-cooling condition, respectively (Wang *et al.* 2026).

2.3. Thermal resistance network model

A thermal resistance network model was developed to describe the complete heat transfer pathway from the maximum chip-surface hotspot to the facility cooling loop. The temperature difference between the chip hotspot and the primary cooling water was decomposed into individual thermal resistance components as Equation 1.

$$T_{chip} - T_{water} = Q_{rack} R_{total} \tag{1}$$

where T_{chip} is the maximum chip-surface hotspot temperature, T_{water} is the primary cooling water temperature, Q_{rack} is the rack heat load, and R_{total} represents the overall thermal resistance between the chip and the facility cooling loop.

The total thermal resistance was expressed as Equation 2.

$$R_{total} = R_{TIM1} + R_{Lid} + R_{TIM2} + R_{Base} + R_{HS} + R_{HX} + R_{Mixing} \tag{2}$$

where R_{TIM1}, R_{Lid}, and R_{TIM2} represent the conduction resistances of the thermal interface material and package layers, respectively. R_{Base} represents the thermal resistance of the heat sink base plate, including both through-plane conduction resistance and spreading resistance caused by the mismatch between the chip heat source area and the heat dissipation area. R_{HS} represents the heat transfer resistance between the heat sink surface and the secondary cooling medium. R_{HX} represents the heat transfer resistance of the terminal secondary-fluid-to-primary-water heat exchanger, while R_{Mixing} represents the equivalent thermal resistance associated with fluid organisation, including airflow bypass and mixing in air cooling and sequential coolant heating where applicable.

The conduction resistance of solid layers was calculated as Equation 3.

$$R_{cond} = \delta/(kA) \tag{3}$$

where δ, k, and A represent the layer thickness, thermal conductivity, and effective heat transfer area, respectively. Since the chip heat source area was considerably smaller than the heat dissipation area, the spreading resistance within the heat sink base plate was additionally considered to describe the non-uniform heat spreading process.

The heat transfer resistance between the heat sink and coolant was calculated as Equation 4.

$$R_{HS} = 1/(h_{HS}A_{HS,eff}) \tag{4}$$

where h_{HS} is the convective heat transfer coefficient between the heat sink and coolant, and $A_{HS,\,eff}$ is the effective heat transfer area considering fin efficiency.

Similarly, the heat exchanger resistance was determined based on the heat transfer coefficient and effective heat transfer area as Equation 5.

$$R_{HX} = 1/(h_{HX}A_{HX}) \tag{5}$$

where h_{HX} and A_{HX} represent the heat transfer coefficient and effective heat transfer area of the terminal secondary-fluid-to-primary-water heat exchanger, respectively.

The mixing resistance was introduced to quantify the additional temperature penalty caused by non-uniform coolant temperature distribution and sequential heating effects. It was defined as Equation 6.

$$R_{Mixing} = (T_{mix} - T_{ref})/Q \tag{6}$$

where T_{mix} represents the mean secondary-fluid temperature at the server, and T_{ref} represents the corresponding mean secondary-fluid temperature at the cooling terminal. The difference includes airflow bypass and hot-return mixing in air-cooled systems and sequential coolant heating where applicable; no air-style bypass-mixing penalty was applied to the cold-plate or immersion-cooling models.

For Direct-to-Die cold-plate cooling, the lid and TIM2 layers were removed and their corresponding thermal resistances were set to zero, whereas TIM1 remained in the heat-transfer pathway. For immersion cooling, zero lid and TIM2 resistances were retained as the study-specific

modelling assumption described above. The developed model therefore provided a unified framework for decomposing the temperature difference between the chip and cooling source into physically meaningful resistance components, enabling the identification of dominant thermal limitations and optimisation pathways for different cooling architectures.

3. Results and discussion

3.1. Comparison of thermal characteristics among different cooling architectures

To eliminate the influence of heat-sink geometry on the comparison of cooling media, the thermal characteristics of different cooling architectures were first evaluated under unified geometric conditions. All architectures adopted identical chip dimensions, package structures, copper base plates, and plate-fin heat sinks with a fixed fin height of 20 mm. The cooling-water inlet temperature and chip heat load were also maintained at the same values. Therefore, this analysis mainly focused on the intrinsic differences introduced by cooling media and heat-transfer pathways.

Figure 2 presents the temperature–heat-transfer (T–Q) characteristics of different cooling architectures. The results show that liquid-cooling architectures achieved significantly lower chip temperature rises than air cooling under the same heat load. The maximum chip-surface hotspot temperature of the air-cooled architecture reached approximately 60 °C, whereas all liquid-cooled configurations maintained chip temperatures below 30 °C. Based on the thermal-resistance analysis, the total thermal resistance of liquid-cooling architectures was only 24.2%–37.5% of that of air cooling, demonstrating the superior heat transport capability of liquid media.

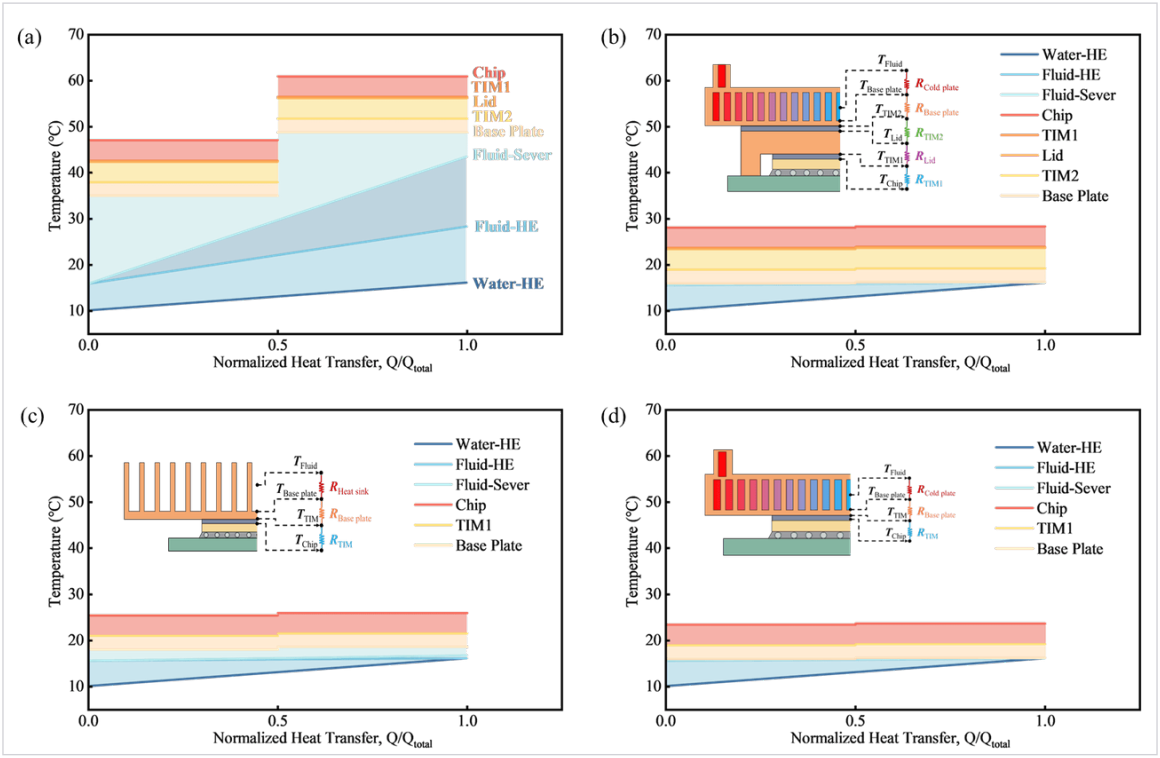

Figure 2. Temperature–heat-transfer (T–Q) characteristics of different cooling architectures under unified conditions: (a) rack-level air cooling, (b) single-phase lidded cold-plate cooling, (c) single-phase immersion cooling, and (d) single-phase Direct-to-Die cold-plate cooling.

3.2. Architecture-specific optimisation of thermal resistance pathways

Although the unified geometry comparison revealed the intrinsic thermal advantages of different cooling media, practical cooling products usually adopt substantially different hardware configurations. Air-cooled heat sinks generally require larger fin structures and sufficient airflow

passages, while liquid-cooling systems can achieve higher heat removal capability with more compact thermal structures. Therefore, architecture-specific hardware optimisation was further considered to investigate the dominant thermal limitations and corresponding improvement pathways of different cooling technologies.

For the air-cooled architecture, the optimised thermal resistance distribution is presented in Figure 3. The results indicated that the dominant thermal resistances remained concentrated in the heat sink, heat exchanger, and coolant mixing processes. The air-side heat transfer resistance accounted for the major temperature loss because air has relatively low thermal conductivity and volumetric heat capacity. Even with enlarged heat sink dimensions, the reduction in solid conduction resistance was limited compared with the remaining air-side resistance. Therefore, further improvement of air cooling mainly relies on increasing effective heat transfer area, improving airflow distribution, and reducing heat exchanger and mixing losses.

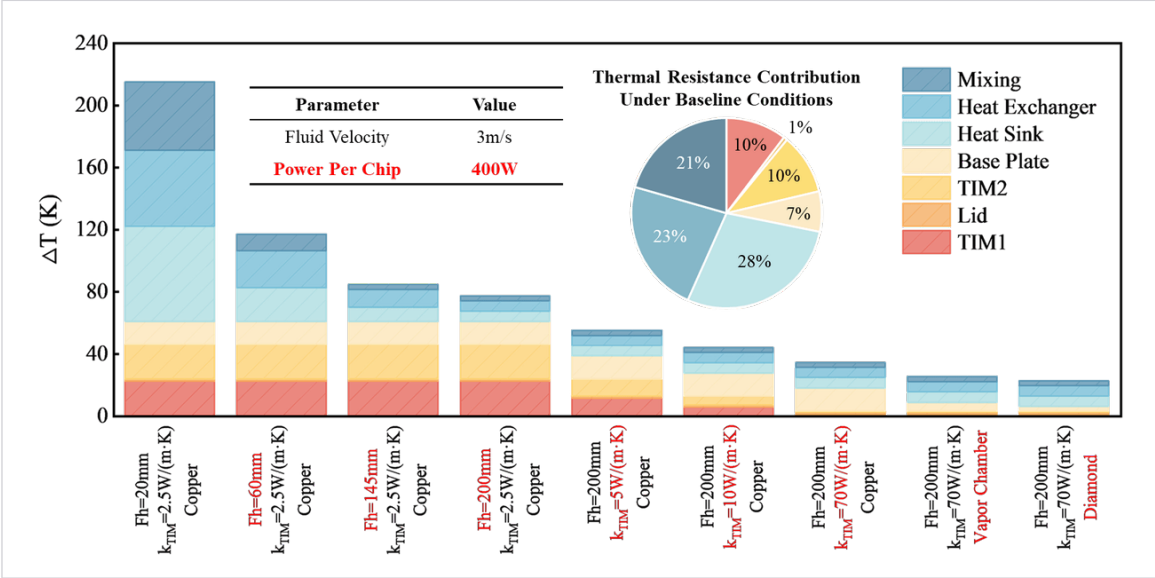

Figure 3. Temperature-difference distribution of the air-cooling architecture with different heat-transfer enhancement strategies.

For the lidded cold-plate cooling architecture, Figure 4 shows that the thermal resistance distribution was significantly different from air cooling. With water as the coolant, the heat sink and coolant-side resistances were substantially reduced, and the dominant limitations shifted towards the package-level thermal pathway. The TIM layers and lid structure became important resistance components because heat still needed to pass through multiple solid interfaces before reaching the coolant. These results indicated that further increasing liquid-side heat transfer capability alone would provide limited improvement, while reducing interface resistance became a more effective optimisation strategy.

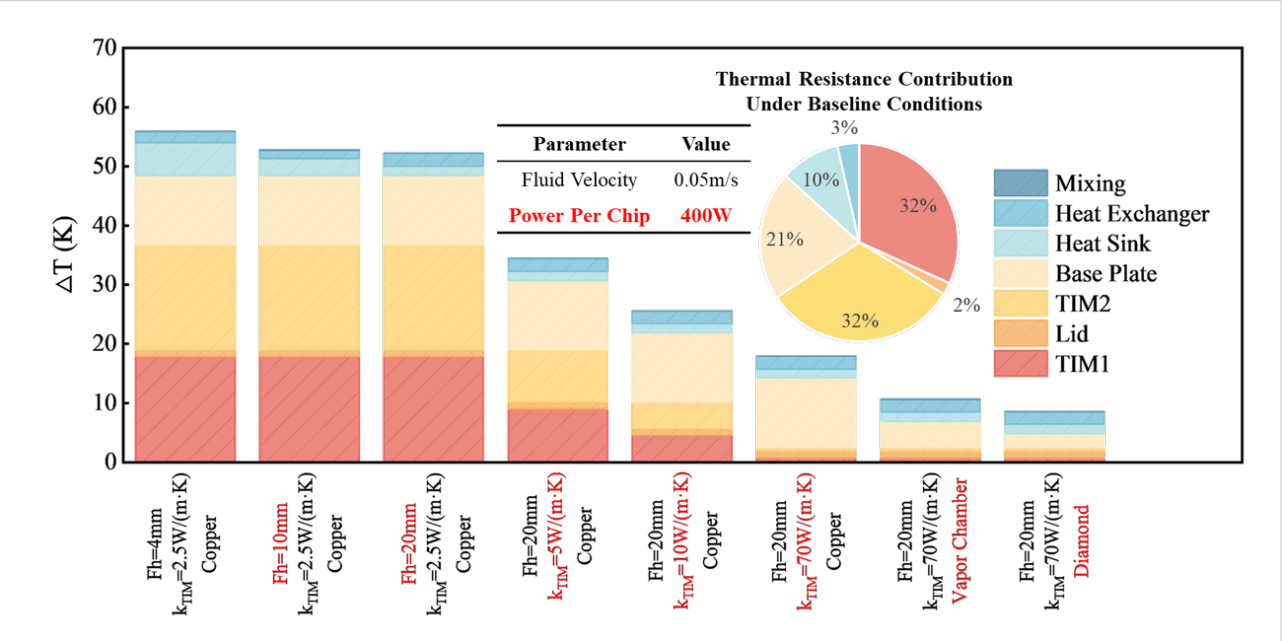

Figure 4. Temperature-difference distribution of single-phase lidded cold-plate cooling with different heat-transfer enhancement strategies.

For Direct-to-Die cold-plate cooling, Figure 5 demonstrates that removing the lid and TIM2 layer further reduced the package-level thermal resistance. However, the remaining interface resistance and heat spreading resistance within the base plate became the primary limitations. This result provided the physical basis for the development of direct-contact cold-plate structures, where the lid and conventional thermal interface layers were further removed by directly attaching the cooling structure to the bare die surface. By eliminating additional interface resistance, direct-contact cooling can further shorten the heat transfer pathway and improve heat dissipation capability for high-density chips.

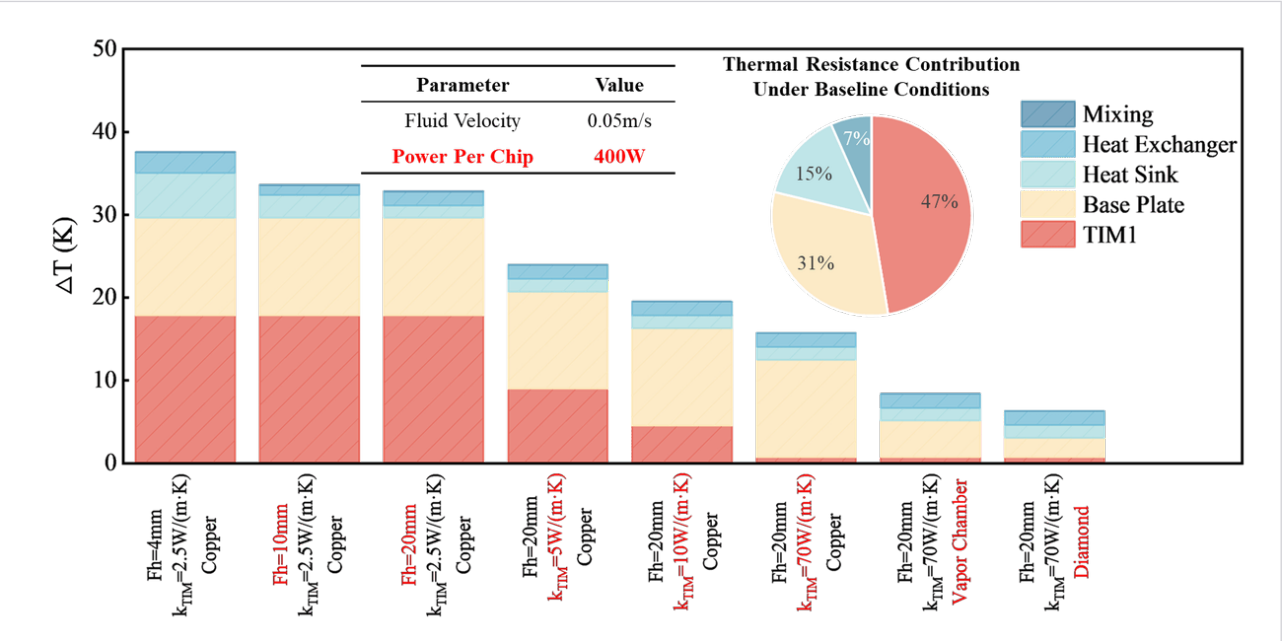

Figure 5. Temperature-difference distribution of single-phase Direct-to-Die cold-plate cooling with different heat-transfer enhancement strategies.

For immersion cooling, Figure 6 illustrates the thermal resistance characteristics of the optimised immersion architecture. Compared with air cooling, immersion cooling eliminated the air-side heat transfer limitation and significantly reduced the heat sink-related resistance. However, due to the lower thermal conductivity and heat capacity of dielectric coolant compared with water, the coolant-side resistance remained non-negligible. Therefore, further improvement of immersion cooling requires simultaneous optimisation of coolant properties, flow distribution, and chip-level heat spreading structures.

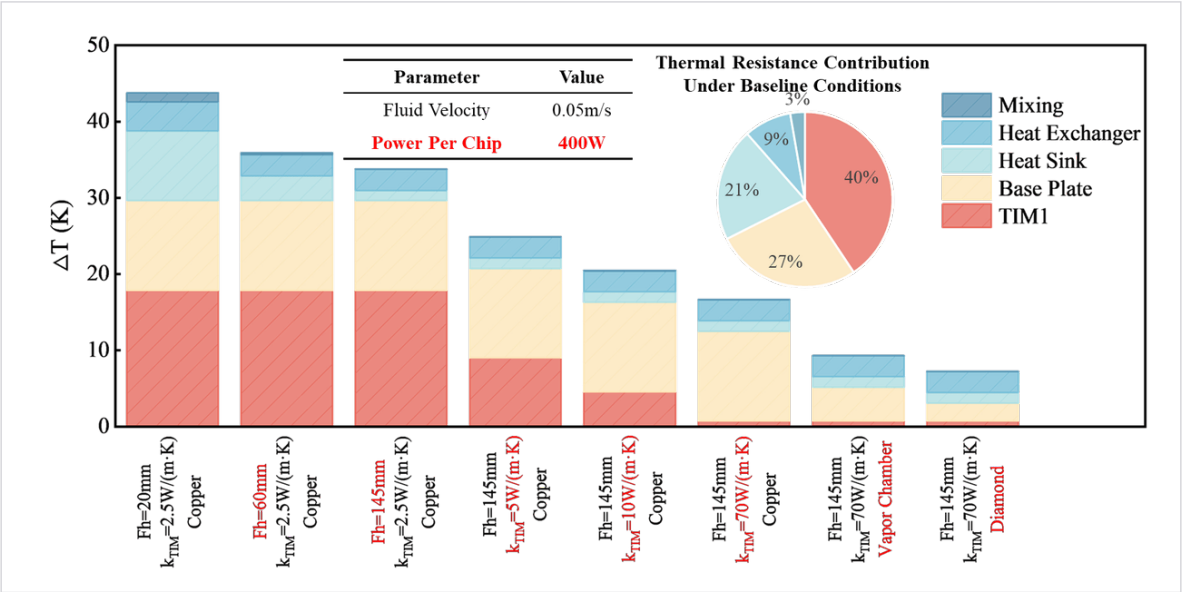

Figure 6. Temperature-difference distribution of single-phase immersion cooling with different heat-transfer enhancement strategies.

Overall, the dominant thermal resistance shifted from fluid-side heat transfer limitations in air cooling to interface and heat spreading limitations in liquid cooling. This transition determined the different optimisation priorities among cooling architectures: air cooling requires enhanced airflow and heat rejection capability, whereas liquid cooling requires advanced interface materials, direct-contact structures, and improved heat spreading technologies.

3.3. Optimisation strategies and cooling capacity enhancement

Based on the thermal resistance analysis presented above, different cooling architectures exhibited distinct optimisation priorities. Therefore, architecture-specific heat-transfer enhancement strategies were introduced rather than applying identical improvements to all cooling methods. As summarised in Table 2, the optimisation strategy mainly focused on three aspects: reducing interface resistance through advanced thermal interface materials, improving heat spreading capability using enhanced base-plate structures, and adjusting heat sink dimensions according to the characteristics of each cooling medium.

Table 2. Heat-transfer enhancement strategies for different cooling architectures.

Enhancement level	Thermal interface material	Base plate or heat-spreading structure	Architecture-specific heat sink design
Baseline configuration	Conventional thermal grease, k_{TIM} = 2.5 W/(m·K)	Copper base plate, $k_x = k_y$ = 400 W/(m·K)	Air cooling: 60 mm; cold-plate cooling: 4 mm; immersion cooling: 20 mm
First-generation enhancement	Phase-change-class TIM, k_{TIM} = 10 W/(m·K)	Vapor chamber, $k_{in\text{-}plane}$ = 20000 W/(m·K) and $k_{through\text{-}plane}$ = 80 W/(m·K)	Air cooling: 200 mm; cold-plate cooling: 20 mm; immersion cooling: 145 mm
Second-generation enhancement	Liquid-metal-class TIM, k_{TIM} = 70 W/(m·K)	Diamond base plate, $k_x = k_y$ = 2000 W/(m·K)	Air cooling: 200 mm; cold-plate cooling: 20 mm; immersion cooling: 145 mm

For air cooling, increasing heat sink height was considered because air-side heat transfer remained the dominant limitation. For liquid cooling architectures, reducing interface resistance and enhancing heat spreading capability became more important due to the reduced coolant-side resistance. These optimisation strategies were subsequently applied to evaluate the achievable cooling capacity of different architectures.

The cooling capacities achieved by different architectures after applying the proposed enhancement strategies are summarised in Table 3. The adopted coolant-velocity modelling upper bounds and a 10 °C facility-water inlet temperature were selected as the operating boundary because they represent favourable conditions for high-density data centre cooling.

Table 3. Maximum heat dissipation capability of different cooling architectures under 10 °C facility-water inlet temperature.

Cooling Architecture	Baseline configuration		First-generation enhancement		Second-generation enhancement	
	Single chip (W)	Single rack (kW)	Single chip (W)	Single rack (kW)	Single chip (W)	Single rack (kW)
Air cooling	330	53	1099	53	1736	83
SP Lidded cold plate	541	173	1755	562	4427	1417
SP Direct-to-Die cold plate	852	273	2686	860	6090	1949
SP Immersion cooling	760	243	2740	219	6340	507

The results indicated that the enhancement strategies significantly improved the achievable heat dissipation capability of all architectures. Among them, immersion cooling achieved the highest single-chip cooling capability, reaching 6.34 kW under the second-generation enhancement condition, followed by Direct-to-Die cold-plate cooling with 6.09 kW. At the rack level, Direct-to-Die cold-plate cooling achieved the highest heat dissipation capacity, reaching approximately 1.95 MW.

The improvement mechanisms were consistent with the thermal resistance analysis. Air cooling benefited mainly from increased heat transfer area, while liquid cooling benefited from reduced interface resistance and enhanced heat spreading capability.

To evaluate the adaptability of different cooling architectures under sustainable high-temperature cooling scenarios, the cooling capacity was further evaluated with a 45 °C facility-water inlet temperature, as summarised in Table 4.

Table 4. Maximum heat dissipation capability of different cooling architectures under 45 °C facility - water inlet temperature.

Cooling Architecture	Baseline configuration		First-generation enhancement		Second-generation enhancement	
	Single chip (W)	Single rack (kW)	Single chip (W)	Single rack (kW)	Single chip (W)	Single rack (kW)
Air cooling	170	27	567	27	899	43
SP Lidded cold plate	268	86	877	281	2231	714
SP Direct-to-Die cold plate	425	136	1348	431	3075	984
SP Immersion cooling	383	122	1378	110	3216	257

Compared with the 10 °C condition, all architectures exhibited reduced heat dissipation capability due to the decreased available temperature difference. However, liquid-cooling architectures maintained higher cooling capacity because of their lower thermal resistance and more efficient heat transport pathways. Direct-to-Die cold-plate cooling continued to demonstrate the highest rack-level cooling capability, indicating its strong potential for future high-temperature and chiller-free data centre applications.

4. Conclusion

A unified thermal resistance framework was established to compare four representative data centre cooling architectures under consistent chip, server, rack, and cooling-source boundaries. The main conclusions were as follows.

First, liquid cooling significantly reduced the overall thermal resistance compared with air cooling. Under the unified 20 mm fin-height condition and identical chip heat load, the total thermal resistance of the investigated liquid-cooling architectures was only 24.2%–37.5% of that of rack-level air cooling. This demonstrated that replacing air with liquid as the principal heat transport medium substantially reduced the temperature penalty between the chip and cooling source.

Second, the dominant thermal resistance shifted as the cooling architecture evolved from air to liquid cooling. Air cooling remained primarily limited by the heat sink, terminal heat exchanger, and airflow warming and mixing processes, whereas thermal interface materials and heat spreading within the base plate became increasingly important after the fluid-side resistance was reduced in liquid cooling. Consequently, air cooling benefited primarily from larger heat-transfer surfaces and improved airflow management, while liquid cooling required greater emphasis on interface materials, shortened package pathways, and advanced heat-spreading structures.

Third, architecture-specific enhancement strategies substantially extended both chip-level and rack-level cooling boundaries. At a facility-water inlet temperature of 10 °C, the second-generation configurations achieved maximum single-chip capacities of 1.74, 4.43, 6.09, and 6.34 kW for air cooling, lidded cold-plate cooling, Direct-to-Die cold-plate cooling, and immersion cooling, respectively. The corresponding rack-level capacities were approximately 83, 1417, 1949, and 507 kW. The different rankings at chip and rack levels demonstrated that cooling performance should not be evaluated solely from the maximum chip heat load, because heat-sink dimensions and server packing density also affected the achievable rack capacity.

Finally, cooling-source temperature strongly influenced the achievable thermal boundary. When the facility-water inlet temperature was increased to 45 °C, the second-generation single-chip capacities decreased to approximately 0.90, 2.23, 3.08, and 3.22 kW for air cooling, lidded cold-plate cooling, Direct-to-Die cold-plate cooling, and immersion cooling, respectively, while the corresponding rack-level capacities decreased to approximately 43, 714, 984, and 257 kW. These results highlighted the importance of low-temperature cooling sources. For air-cooled systems in particular, access to low-temperature natural heat sinks, such as deep reservoir or lake water and deep seawater, could provide additional thermal margin and extend the applicable power-density range. Liquid cooling maintained substantially higher cooling capability under the 45 °C condition, indicating stronger adaptability to high-temperature cooling and reduced dependence on mechanical refrigeration; nevertheless, lower cooling-source temperatures further increased its achievable heat dissipation capacity. Therefore, the coordinated selection of cooling architecture, thermal enhancement strategy, and available natural cooling source is important for the development of high-density and energy-efficient data centres.

The present model employed idealised steady-state boundary conditions and simplified effective heat-transfer coefficients. Transient workload variations, detailed manifold flow maldistribution, contact-pressure-dependent interface resistance, component reliability, and facility-level climate variations were not explicitly considered. These effects should be incorporated in future experimental validation and system-level optimisation.

Declarations

Author contributions. Xinyi Wang: Writing (original draft), Writing (review and editing), Data curation, Formal analysis, Investigation, Methodology, Conceptualization, Visualization, Validation, Project administration. Lu Wang: Investigation, Formal analysis, Project administration, Software. Haojiang Ran: Investigation, Data curation. Tiexiao Xu: Data curation, Methodology. Bo Zhang: Data curation, Methodology. Zhen Li: Conceptualization, Methodology, Supervision, Project administration, Resources, Software, Funding acquisition.

Funding. This study was supported by funding from Daikin Industries, Ltd.

Awards. DAIKIN Industries

Conflicts of interest. The authors declare no competing interests.

Data availability. The data that support the findings of this study are available from the corresponding author upon reasonable request.

References

Alkrush, Ahmed A.; Salem, Mohamed S.; Abdelrehim, O.; Hegazi, A.A. (2024) 'Data centers cooling: A critical review of techniques, challenges, and energy saving solutions', *International Journal of Refrigeration*, 160, pp. 246-262. Available at: https://doi.org/10.1016/j.ijrefrig.2024.02.007

Azarifar, Mohammad; Arik, Mehmet; Chang, Je-Young (2024) 'Liquid cooling of data centers: A necessity facing challenges', *Applied Thermal Engineering*, 247, pp. 123112. Available at: https://doi.org/10.1016/j.applthermaleng.2024.123112

Cai, Senhong; Gou, Zhonghua (2026) 'Towards energy-efficient data centers: A comprehensive review of passive and active cooling strategies', *Energy and Built Environment*, 7(1), pp. 206-226. Available at: https://doi.org/10.1016/j.enbenv.2024.08.009

Chacón, Andrés M.N.; Dominguez, Dany S.; Oliveira, Francisco B.S. (2024) 'Enhancing heat transfer processes in 2U-Rack Servers: A study case employing thermal simulation and the Taguchi method', *Case Studies in Thermal Engineering*, 54, pp. 104006. Available at: https://doi.org/10.1016/j.csite.2024.104006

Chang, Qiankun; Huang, Yuanfeng; Liu, Kaiyan; Xu, Xin; Zhao, Yaohua; Pan, Song (2024) 'Optimization Control Strategies and Evaluation Metrics of Cooling Systems in Data Centers: A Review', *Sustainability*, 16(16), pp. 7222. Available at: https://doi.org/10.3390/su16167222

Huang, Yongping; Liu, Bin; Xu, Shijie; Bao, Chujin; Zhong, Yangfan; Zhang, Chengbin (2024) 'Experimental study on the immersion liquid cooling performance of high-power data center servers', *Energy*, 297, pp. 131195. Available at: https://doi.org/10.1016/j.energy.2024.131195

International Energy Agency (2025) *Energy and AI*. Paris: IEA. Available at: https://www.iea.org/reports/energy-and-ai

International Energy Agency (2026) *Key Questions on Energy and AI*. Paris: IEA. Available at: https://www.iea.org/reports/key-questions-on-energy-and-ai

Kalantarpour, Raha; Vafai, Kambiz (2024) 'Advancements in data center thermal management', *Advances in Heat Transfer*. Elsevier, pp. 39-80. Available at: https://doi.org/10.1016/bs.aiht.2024.07.002

Kong, Rui; Zhang, Hainan; Tang, Mingsheng; Zou, Huiming; Tian, Changqing; Ding, Tao (2024) 'Enhancing data center cooling efficiency and ability: A comprehensive review of direct liquid cooling technologies', *Energy*, 308, pp. 132846. Available at: https://doi.org/10.1016/j.energy.2024.132846

Li, Ziyong; Luo, Hailiang; Jiang, Yuguang; Liu, Haichao; Xu, Lian; Cao, Kunyuan; Wu, Hongjie; Gao, Peng; Liu, Hong (2024) 'Comprehensive review and future prospects on chip-scale thermal management: Core of data center's thermal management', *Applied Thermal Engineering*, 251, pp. 123612. Available at: https://doi.org/10.1016/j.applthermaleng.2024.123612

Park, Woo Young; Lee, Donggun; Ganesan, Vishwanath; Du, Xuzhi; Pinkus, Ilan; Garimella, Vivek S.; Aguirre, Trevor G.; Nawaz, Kashif; King, William P.; Miljkovic, Nenad (2026) 'Holistic energy analysis method for thermal management architectures of data centers', *International Journal of Heat and Mass Transfer*, 268, pp. 128971. Available at: https://doi.org/10.1016/j.ijheatmasstransfer.2026.128971

Schultz, Mark D.; Parida, Pritish R. (2024) 'Compliant Direct Attach Liquid Cooling', *2024 23rd IEEE Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm)*. IEEE, pp. 1-6. Available at: https://doi.org/10.1109/itherm55375.2024.10709396

Wang, Huijuan; Yuan, Xuejun; Zhang, Kun; Lang, Xujin; Chen, Hua; Yu, Huimin; Li, Shengtao (2024) 'Performance evaluation and optimization of data center servers using single-phase immersion cooling', *International Journal of Heat and Mass Transfer*, 221, pp. 125057. Available at: https://doi.org/10.1016/j.ijheatmasstransfer.2023.125057

Wang, Zi-Xing; Xue, Ke; Chen, Jun-Yu; Li, Nan; Tao, Wen-Quan (2026) 'Numerical investigation on a high-temperature data center cooled by combined liquid-cooling and free-cooling - a comprehensive case study', *Applied Energy*, 404, pp. 127154. Available at: https://doi.org/10.1016/j.apenergy.2025.127154

Yang, Zhengchao; Yao, Qiufei; Wang, Yu; Gu, Junlong; Yu, Zhichen; Li, Qipeng; Sun, Xiaoyi; Yang, Xuejing (2025) 'Optimization of Heat Transfer and Flow Performance of Microchannel Liquid-Cooled Plate Based on Orthogonal Test', *Buildings*, 15(6), pp. 905. Available at: https://doi.org/10.3390/buildings15060905

Zhang, Yingbo; Li, Hangxin; Wang, Shengwei (2024) 'Life-cycle optimal design and energy benefits of centralized cooling systems for data centers concerning progressive loading', *Renewable Energy*, 230, pp. 120847. Available at: https://doi.org/10.1016/j.renene.2024.120847

Zhao, Ruolan; Du, Yahui; Yang, Xiaochen; Zhou, Zhihua; Wang, Wufan; Yang, Xueqing (2023) 'A critical review on the thermal management of data center for local hotspot elimination', *Energy and Buildings*, 297, pp. 113486. Available at: https://doi.org/10.1016/j.enbuild.2023.113486

HOW TO CITE THIS PAPER

Wang, X., Wang, L., Ran, H., Xu, T., Zhang, B. and Li, Z. (2026). Parametric Characterisation and Operating Boundaries of Air- and Liquid-Cooling Architectures in High-Density Data Centres. In: *Proceedings of the 22nd International Conference on Sustainable Energy Technologies (SET2026)*, Kuala Lumpur, Malaysia, 13 to 15 August 2026. Caravel Press, pp. 258–270. ISBN 978-1-995850-00-9. ISSN 2978-4387. https://doi.org/10.65582/set.2026.022

COPYRIGHT AND LICENCE

© 2026 The Author(s). This is an open access paper distributed under the terms of the Creative Commons Attribution (CC BY 4.0) licence, which permits unrestricted use, distribution and reproduction in any medium, provided the original work is properly cited.

WSSET World Society of Sustainable Energy Technologies Organised by the World Society of Sustainable Energy Technologies. Published by Caravel Press.