System-Level Analysis of IT-to-Facility Heat Exchanger Performance In Two-Phase Direct-to-Chip Cooled AI Data Centers

Trevor A. Whitaker, Qingyang Wang, Serdar Ozguc, Akshith Narayanan, Jacob D. Moore, Lucas Beran, Richard W. Bonner
Accelsius, Austin, TX
Email: twhitaker@accelsius.com

Abstract— **The rapid rise in power density and thermal design power (TDP), driven largely by artificial intelligence (AI) data center workloads, necessitates advanced thermal management solutions at both the IT and facility levels. Within the IT infrastructure, two-phase direct-to-chip (2PDTC) cooling is gaining increasing interest, offering lower pumping power, improved thermal performance, decreased mechanical cooling requirements, and extended maintenance intervals. However, it is not yet clear how evolving data center architectures impact 2PDTC cooling loops, particularly at the condenser, where careful thermal design is essential to maintain overall system performance. This work explores how facility cooling impacts the performance of a pumped-refrigerant R515b IT loop with three different facility-side fluids: air, PG25, and R515b. A comparative analysis of component- and system-level performance is conducted for each configuration in the context of state-of-the-art AI data centers. Parametric evaluations of condensers are performed using analytical models, with primary loop inlet temperatures varied from 10–50 °C and heat exchanger duties from 100–1000 kW. Industry-standard brazed plate heat exchangers in counterflow are modeled for liquid and refrigerant combinations, while liquid-to-air radiators in crossflow are considered for air-cooled facility loops. Analytical models are validated against empirical data from a commercially available 250 kW modular two-phase coolant distribution unit using each of the three available facility fluids. Finally, a system-level operational expenditure (OpEx) estimate is presented for the different heat exchanger configurations. The results demonstrate that IT cooling design should be influenced not only by white space thermal dissipation requirements but also by the characteristics of the grey space facility fluid. These findings highlight the importance of integrated thermal-system design approaches to achieve the efficiency, scalability, and performance targets of next-generation high-power computing.**

Keywords—data centers, thermal management, two-phase, heat exchangers

I. INTRODUCTION

Recent advances in artificial intelligence (AI) algorithms, particularly large language models (LLMs) like those employed by ChatGPT and Claude, have spurred a significant increase in demand for high-performance processors and graphics processing units (GPUs). Invariably, as processor power increases, so too does thermal design power (TDP). Current generation Nvidia Blackwell GPUs have nearly doubled the TDP of the previous generation Hopper chips and future generations are expected to continue this trend, with Rubin Ultra (2027) expected to have over five times higher TDP than Hopper [1]. Extreme TDP and miniaturized chip architecture result in significant increases in the heat fluxes managed by technology cooling systems (TCS).

These constantly-increasing thermal demands present challenges for existing TCS technologies, and have motivated the development of more efficient and effective cooling strategies for the AI datacenter industry [2, 3, 4, 5]. Air cooling, the predominant heat transfer medium at the chip level for older compute chips, is quickly becoming obsolete [6]. Chip manufacturers are requiring liquid cooling for high-flux components as soon as next generation. The two most seriously considered heat removal mechanisms for TCSs in future server builds are single-phase direct to chip (1PDTC) and two-phase direct to chip (2PDTC). 1PDTC relies on forced convection to produce large heat transfer coefficients, while 2PDTC utilizes a fluid's enthalpy of vaporization to absorb heat from the chip. All practical applications of both mechanisms operate within microchannel cold plates to maximize their performance [7, 8, 9].

The benefits, tradeoffs, and unique characteristics of each method is presently a highly active area of investigation within the industry [10, 11, 12, 13, 14]. However, discussions of the efficacy of the TCS are typically bounded by the TCS itself, and fail to examine the effect of the TCS cooling mechanism on the entire data center facility cooling system (FCS). Importantly, the behavior of the heat exchanger that acts as an interface between TCS and FCS can change dramatically depending on the combination of fluids used. In order to fully account for thermal performance of the data center as a system (rather than the server rack as a system), the physics of the IT-to-Facility heat exchanger (ITFHX) must be resolved.

In this work an analytical investigation of ITFHX performance is conducted for the most likely combinations of fluids. While air is not of interest as a heat transfer medium in the TCS, many datacenters—particularly those being retrofitted for AI servers—rely on cold air as the primary (facility) coolant. A 25% propylene glycol-water mixture (PG25) is by far the most widely used single-phase coolant on both the IT and facility cooling loops. While several dielectric refrigerants are being actively investigated for 2PDTC systems, current generation HFO refrigerants such as R515b are the most attractive candidates. Due to performance enhancement exhibited by refrigerants on both sides of the ITFHX, it is likely that facility-side refrigerants will be investigated in the near future. Herein, ITFHX performance is presented for all combinations of fluids with air, PG25, and R515b on the facility side with PG25 and R515b on the IT side.

The analytical heat exchanger model presented here captures an understanding that is critically important to end-to-end cooling system performance and can lend insight into the design and retrofit of data centers for current and future-generation, high-power AI servers. Careful design of the entire data center thermal architecture is critical to efficient capital allocation and reduction of future retrofitting.

II. Methodology

A. Model Overview

The model used in the present work was built in Python and uses REFPROP 10.0 for thermophysical property evaluation. It was developed from the analytical ε-NTU heat exchanger model.

B. Stream Equations

Stream equations are developed from an overall energy balance on a single fluid stream. Their general form is

$$q = \dot{m}\left(c_p \Delta T + \chi i_{fg}\right), \tag{1}$$

where q is the heat load in W, $\dot{m}$ is the mass flow rate in kg s^{-1}, c_p is the isobaric specific heat capacity in J kg^{-1} K^{-1}, ΔT is the temperature rise across the heat exchanger in K, χ is the exit quality, and i_{fg} is the enthalpy of vaporization in J kg^{-1}.

For a single-phase stream, no vaporization occurs and thus χ and i_{fg} are zero. For a two-phase stream, ΔT represents the exit subcooling if the two-phase fluid is on the IT side and the inlet subcooling if the two-phase fluid is on the facility side.

C. Phase Contributions and Stream Discretization

Because the ε-NTU model cannot account for fluids that are both single-phase and two-phase in a single stream, the heat exchanger for any two-phase fluid having a nonzero inlet or exit subcooling must be discretized into single-phase and two-phase segments, analyzed in series. For a heat exchanger with phase change occurring on both facility and IT sides, the model must further discretize into three sections: single-phase in both streams (1p), single-phase in one stream and two-phase in the other (1p2p), and two-phase in both streams (2p). The inlet conditions to each heat exchanger segment are determined by calculating the exit conditions of the segment directly upstream.

Phase contributions are calculated for a single stream as follows:

$$q_{1p} = \dot{m} c_p \Delta T, \tag{2}$$

$$q_{1p2p} = \left| q_{1p,H} - q_{1p,C} \right|, \tag{3}$$

and

$$q_{2p} = \dot{m} \chi i_{fg}, \tag{4}$$

where subscripts H and C represent the hot and cold (IT and facility) streams, respectively.

In this way, each heat exchanger is divided into 3 segments, each accounting for some portion of the total heat exchanger load. For a heat exchanger with only one phase change fluid, q_{2p} becomes zero and for a heat exchanger with two single-phase fluids, both q_{2p} and q_{1p2p} are zero.

D. Heat Exchanger Equation

The total heat exchanger energy balance given by the ε-NTU method is

$$q = \varepsilon C_{min}\left(T_{H,i} - T_{C,i}\right), \tag{5}$$

where ε is the heat exchanger effectiveness, C_{min} is the heat capacity rate in W K^{-1}, and $T_{H,i}$ and $T_{C,i}$ are the hot and cold stream inlet temperatures in K, respectively. For 1p and 1p2p heat exchanger segments, this equation is solved with analytical calculations of effectiveness that are readily available for the counterflow and crossflow heat exchanger configurations used in this study. The overall heat transfer coefficient (UA) for all conditions is set to a value that will result in an approach temperature of 3 K to facilitate solving (5). Note that while approach temperature for single-phase heat exchangers is typically taken as

$$T_{app} = T_{H,o} - T_{C,i}, \tag{6}$$

it is evaluated more generally in the present model as the minimum temperature difference between the hot and cold streams to properly account for phase change-related pinching. The UA and 3 K approach temperature assumption can be justified by the fact that heat exchanger size for current coolant distribution units (CDUs), which is directly proportional to UA, is relatively unconstrained in practical applications and commercially available CDUs often have greatly oversized heat exchangers to minimize approach temperatures. For this reason, and somewhat counterintuitively, heat exchanger performance in data center applications is not constrained by its geometry, and both UA and approach temperature can be selected arbitrarily within reasonable typical ranges. Additionally, fixing a constant approach temperature allows for better comparison between ITFHX performance with different fluids. One exception to this is in some configurations of heat exchangers using air as the facility fluid. In these cases, the heat exchanger effectiveness becomes the limiting factor for approach temperature and even with infinite UA, the approach will be greater than 3 K.

For 2p heat exchanger segments, (5) cannot be used to solve for inlet temperatures as εC_{min} is not defined in these cases. However, the present model assumes that all phase change occurs isothermally, a reasonable assumption for pure and azeotropic refrigerants with realistically small pressure drop. For these conditions, log mean temperature difference is equal to $T_{H,i} - T_{C,i}$ and the temperatures can be solved for simply as

$$q = UA\left(T_{H,i} - T_{C,i}\right). \tag{7}$$

E. Parametric Sweep and Solution Methodology

Each fluid pair in the ITFHX is tested for a range of heat loads (100, 250, 500, and 1000 kW) and facility inlet temperatures (10, 20, 30, 40, and 50 °C). For each condition, UA is iteratively selected to achieve 3 K approach temperature, χ is set to 0.7 for all two-phase conditions, and $\dot{m}$ of each stream is selected to achieve 10 K temperature rise for single-phase streams or 2 K subcooling for two-phase streams. The 10 K temperature rise is the industry standard, selected to manage system pressure drop and fluid velocity. For two-phase systems, subcooling at the heat exchanger exit is purely dependent on the difference in height between the heat exchanger and the reservoir, and is somewhat arbitrary [15].

With the above model inputs, the facility-side stream equation (1) can be solved directly. Next, phase contributions are calculated and streams are discretized into 1p, 1p2p, and 2p segments as appropriate. For each heat exchanger segment, the relevant heat exchanger equation (5) or (6) is solved, returning inlet temperatures to every segment. Finally, each segment on the IT side can be solved using the stream equation (1), after which the system is fully defined.

Note that for two-phase fluids, small errors in the temperature used for property lookup can lead to relatively large errors in the subcooling because the single-phase contribution is small relative to the two-phase contribution. Fluid properties are only calculated using the input temperature (primary inlet temperature), introducing error for large approach temperature differences or temperature rises. Because of this, one of the limitations of the model in its current state is ability to predict the subcooled temperature entering or exiting the heat exchanger. However, because IT cooling performance is predominantly dictated by the saturation temperature, instead of the inlet temperature, the effect of this limitation can be minimized by analyzing a condenser with very low subcooling. Future versions of the model will implement an iterative procedure to eliminate any error in temperatures used for property evaluation.

Moreover, the quantity of interest in this study is the effective temperature of the IT fluid at the cold plate, as this is a better indicator of cold plate performance for both single-phase and two-phase systems. In a two-phase cold plate, it is reasonable to assume the fluid is isothermal and the IT return temperature is equal to the saturation temperature in the cold plate. In single-phase systems, temperature rises proportional to the heat added along the cold plate and the average fluid temperature will be the effective temperature. Therefore, for two-phase systems without significant subcooling, IT return temperature is the quantity of interest and for single-phase systems, the average IT fluid temperature is the quantity of interest.

III. Results

A. Experimental Validation

A limited set of experimental data were used to validate the analytical heat exchanger model. Tests were performed on a commercially available thermal simulation rack (Accelsius MR250 TSR) that consists of a R515b CDU and a 42U server rack with 6 kW load simulation sleds (Accelsius 6 kW LSS). Load simulation sleds consist of 5 independent copper thermal test vehicles (TTV), each capable of dissipating up to 1 kW at 27 W cm^{-2}. The TSR is capable of managing heat loads up to 250 kW using either PG25 or air as the primary coolant. Although the combination of R515b with PG25 and air primary coolants represents a limited sample of the conditions present in this work, the specific combinations available for testing are sufficient to validate the segmented heat exchanger model and its ability to predict both single-phase and two-phase behavior. The experimental matrix is shown in Table 1, and model inputs are set to match the inputs from the experiments. PG25 data are taken with constant heat load (250 kW) and varied primary inlet temperature, while data for air are collected at relatively constant inlet temperatures (~30 °C) with varied heating loads.

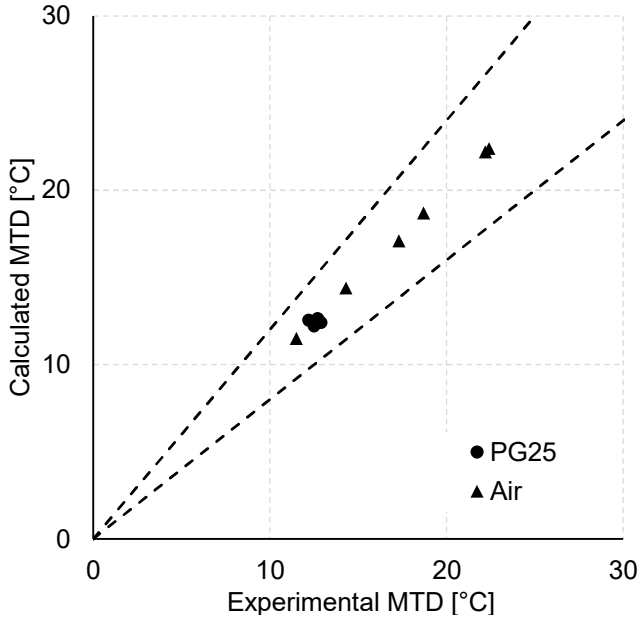

Fig. 1 Validation of model predictions for R515b with PG25 and air as primary coolants. Dashed bands correspond to relative errors of ±20%.

Table 1 Conditions tested for experimental validation of model.

PRIMARY FLUID	POWER [KW]	PRIMARY INLET TEMP. [°C]	SECONDARY INLET TEMP. [°C]	MTD [°C]
PG25	250	20.4	32.6	12.2
	250	30	42.7	12.7
	250	40	52.9	12.9
	250	45	57.5	12.5
AIR	76.8	27	38.5	11.5
	103.8	29.9	47.2	17.3
	124.8	32.5	54.7	22.2
	77.8	32.1	46.4	14.3
	101.2	32.5	51.2	18.7
	127.2	31.4	53.8	22.4

A comparison between experimental data and model predictions is shown in Fig. 1. Here the maximum temperature difference (MTD), defined as

$$MTD = T_{H,i} - T_{C,i}, \tag{8}$$

is plotted as opposed to an absolute temperature to better capture the error between model and experiments [15]. For absolute temperatures, error between the model and experiments is <1%. Data for PG25 as the primary coolant are represented by circles and data for primary air are represented by triangles. Because rack power is constant in the PG25 tests at 250 kW and only the cold supply temperature varies, these data are clustered around a relatively constant MTD of 12 °C. Primary air testing, conducted with constant air temperature and varied rack power between 77 and 127 kW, shows a wider range of MTD. The dashed lines represent error bands corresponding to ±20%, and the model shows very good agreement with experimental data.

B. Parametric Study

The ITFHX parametric sweep included the three primary coolants (air, PG25, R515b), two secondary coolants (PG25, R515b), heat loads (100, 250, 500, and 1000 kW), and cold inlet temperatures (10, 20, 30, 40, and 50 °C). All conditions use an exit quality of 0.7 for R515b and a target approach temperature of 3 °C. Fig. 2 shows trends in effective fluid temperature (mean TCS fluid temperature for single-phase fluids and condenser return temperature for R515b) as a function of facility fluid inlet temperature.

Panel a) shows effective temperature using air as the primary coolant. In this configuration, the effective TCS fluid temperature is slightly lower for PG25 than it is for R515b, delivering fluid to the cold plate with an effective temperature approximately 3 °C cooler than R515b for the same air inlet temperature. Even though the maximum cross-flow effectiveness for two single-phase fluids is relatively low (~0.64) and the target approach temperature cannot be achieved, ITFHX performance is slightly better for PG25 than R515b. Note that as approach temperature is fixed, the size of the ITFHX will be significantly larger for facility air than for PG25 or R515b.

Similarly in panel b), showing performance relative to PG25 as the primary fluid, a PG25 TCS will deliver fluid with lower effective temperature to the cold plate by approximately 5 °C. Although performance at the ITFHX is slightly worse for two-phase fluids with a single-phase primary fluid, it is again important to consider entire system performance when making design decisions. For example, two-phase fluids can typically produce much higher heat transfer coefficients within both the cold plate and ITFHX than their single-phase counterparts [15, 16]. Even when supplying warmer fluid to the cold plate, the improved cold plate performance can lead to enhanced total system performance.

Just as the fully single-phase heat exchanger (single-phase fluids on facility and IT sides) tends to outperform a mixed-phase heat exchanger with refrigerant on the IT side, a fully two-phase heat exchanger shows improved performance over a mixed-phase heat exchanger with PG25 on the IT side. Panel c) shows this trend, with an R515b TCS delivering up to 3.5 °C

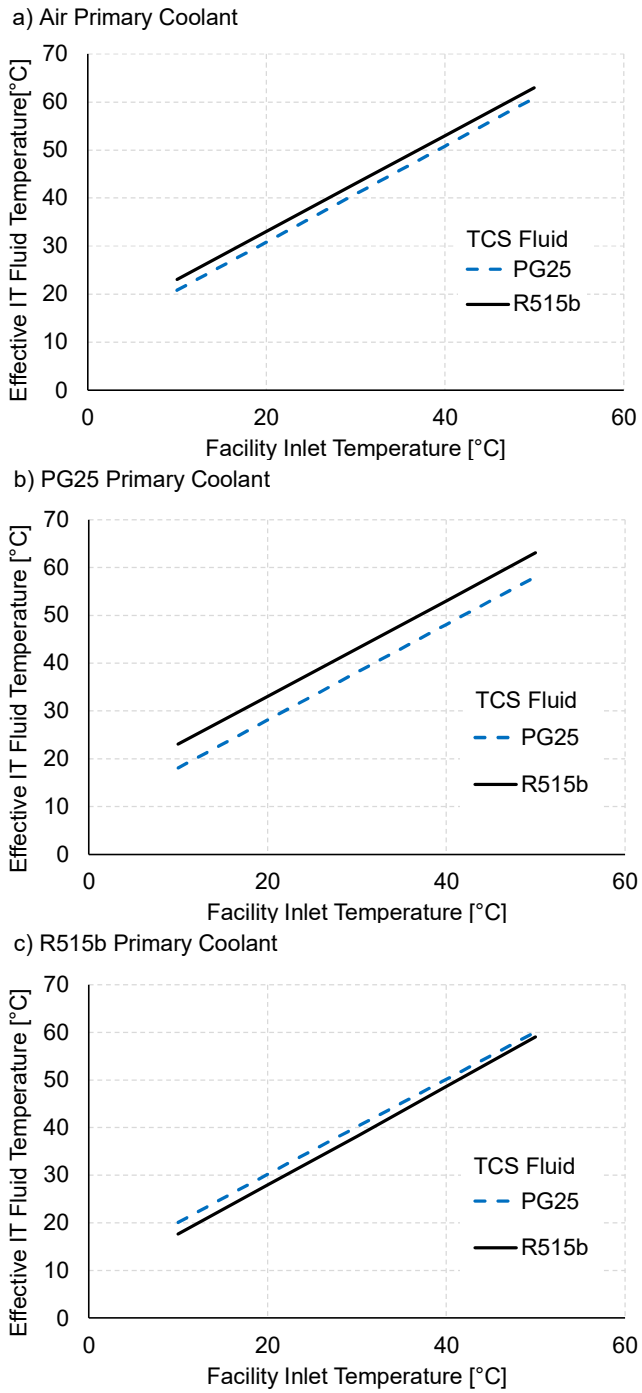

Fig. 2 Required mass flow rate as a function of ITFHX heat load for the three candidate fluids

cooler effective fluid temperatures to the cold plate than a PG25 TCS. The thermodynamic justification for this trend can be understood by examining the rate of entropy generation within the heat exchanger as a function of heat capacity ratio,

$$C_r = \frac{C_{min}}{C_{max}}. \tag{9}$$

For a given effectiveness, heat exchangers with a heat capacity ratio approaching unity will reduce the rate of irreversible entropy production through more uniform temperature profiles and thus lower overall temperature differences. The rate of entropy generation in a heat exchanger with $C_{min} = C_{Cold}$ (as is the case for all heat exchanger arrangements presented here) can be expressed as [17]

$$\frac{\dot{S}_{gen}}{C_C} = \frac{1}{C_r}\ln\left\{1 - \varepsilon C_r\left(1 - \frac{T_{C,i}}{T_{H,i}}\right)\right\} + \ln\left\{1 + \varepsilon\left(\frac{T_{H,i}}{T_{C,i}} - 1\right)\right\}. \quad (10)$$

Therefore, decreasing the temperature difference between the hot and cold streams by increasing the heat capacity ratio will reduce entropy generation and improve the performance of the heat exchanger. Mixed-phase heat exchangers exhibit $C_r = 0$, and will therefore inherently experience high entropy generation.

With the additional benefits provided by two-phase TCSs such as enhanced two-phase heat transfer coefficients, it is likely that significant performance gains can be realized with the combination of facility and IT refrigerant in the ITFHX. Moreover, because i_{fg} for two-phase fluids in practical data center applications is large relative to $c_p\Delta T$ for single-phase fluids, required mass flow rates are typically much lower in two-phase TCSs.

Fig. 3 shows the relationship between required mass flow rate and ITFHX heat load. This relationship will be identical on either the TCS or FCS side of the heat exchanger. With an exit quality of 0.7 for R515b the required mass flow rate is approximately 40% of the required mass flow rate for PG25 and 10% of that for air. With this understanding, it becomes clear that the use of refrigerant in both the FCS and TCS will minimize the mass flow rate of both fluid streams, supply the lowest effective fluid temperature to the cold plate through entropy generation minimization in the ITFHX, and generate the highest heat transfer coefficients within the cold plate.

C. OpEx Implications

The primary contribution to data center operational expense (OpEx) imparted by TCS and FCS fluid selection is the alteration of required facility coolant temperature. Recent studies show that each degree increase in FCS temperature results in ~4% OpEx savings annually and improved power usage effectiveness (PUE) due to increased available "free-cooling" hours [18]. Reduction of the effective TCS fluid temperature presented in section III.B. can be reframed as potential to increase FCS system temperature. Therefore, balanced (i.e. heat exchangers with the same fluid on both primary and secondary sides) PG25 and R515b ITFHXs can operate at 5 °C and 3.5 °C higher FCS temperatures, respectively, than their unbalanced counterparts at the same effective TCS fluid temperature. These increases in FCS temperatures correspond to 20% and 14% reductions in annual energy costs as a result of TCS fluid selection alone. While in-rack infrastructure for two-phase cooling systems typically carries higher capital expenditure (CapEx), total data center CapEx is similar between single-phase and two-phase cooling. Moreover, in a comparative analysis of data center total cost of ownership (TCO), OpEx was estimated to account for approximately 30% of the TCO [19]. Therefore, for the balance heat exchanger conditions analyzed here, the increase in facility water temperature enabled by a reduction of entropy production in the heat exchanger could result in TCO savings of 4-6% over a 5 year period.

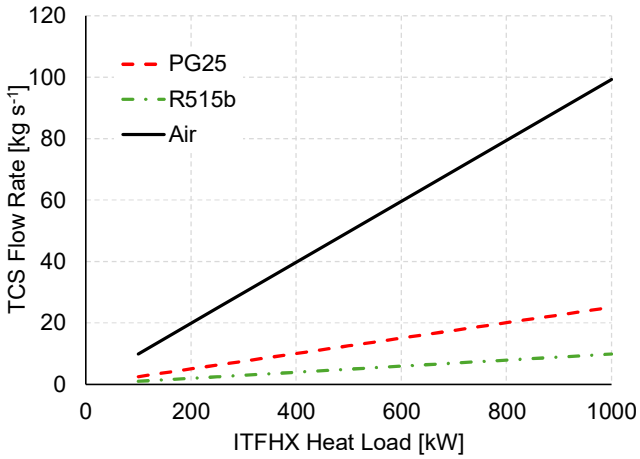

Fig. 3 Effective IT fluid temperature for PG25 and R515b as a function of facility inlet temperature for primary coolants a) air, b) PG25, and c) R515b.

IV. Conclusions

The present study investigated IT-to-Facility heat exchanger (ITFHX) performance using different combinations of fluids in the technology cooling system (TCS) and facility cooling system (FCS). Varied FCS supply temperatures and heat exchanger loads were investigated with PG25 and R515b in the TCS and air, PG25, and R515b in the FCS. Calculation of the effective TCS temperature (average TCS temperature for single-phase fluids and TCS return temperature for two-phase fluids) revealed that ITFHX performance is maximized when the fluid streams on either side of the heat exchanger are balanced ($C_r = 1$) and entropy generation is minimized. For the conditions tested, unbalanced heat exchangers can result in effective TCS temperatures up to 5 °C higher than their balanced counterparts. In systems using air as the primary coolant, low maximum heat exchanger effectiveness for fully single-phase heat exchangers results in higher effective temperatures, even with infinite overall heat transfer coefficient (UA). The best performing ITFHX configurations are balanced with either PG25 or R515b in both streams, which can reduce OpEx costs for the data center by up to 20% compared with the unbalanced counterpart. These configurations supply the same effective temperature to the cold plate. However, superior two-phase heat transfer coefficients at the cold plate and significant (60%) reduction in mass flow rate for the R515b ITFHX compared to PG25 makes the fully refrigerant-based cooling system the best performance and lowest operational cost configuration.

The heat exchanger model developed in this work is limited primarily by small errors in property evaluation resulting from a single-pass solution approach. Future improvements to the model will implement an iterative solution of the heat exchanger to eliminate these errors, along with an inclusion of models for cold plate and facility heat rejection to fully capture end-to-end performance of the thermal system. The findings presented in

this work help to enable system-level analysis of high-performance AI data center thermal stackups which are critical for datacenter efficiency and preparation for future-generation thermal demands.

REFERENCES

[1] A. Shilov, "Future AI processors said to consume up to 15,360 watts of power — massive power draw will demand exotic immersion and embedded cooling tech," 17 June 2025. [Online]. Available: https://www.tomshardware.com/pc-components/cooling/future-ai-processors-said-to-consume-up-to-15-360w-massive-power-draw-will-demand-exotic-immersion-and-embedded-cooling-tech. [Accessed 2 January 2026].

[2] A. A. Alkrush, M. S. Salem, O. Abdelrehim and A. Hegazi, "Data centers cooling: A critical review of techniques, challenges, and energy saving solutions," *International Journal of Refrigeration,* vol. 160, pp. 246-262, 2024.

[3] E. Masanet, "Recalibrating global data center energy-use estimates," 28 Feb 2020. [Online]. Available: 10.1126/science.aba3758. [Accessed 2 January 2026].

[4] A. Shehabi, S. J. Smith, A. Hubbard, A. Newkirk, N. Lei, M. A. B. Siddik, B. Holecek, J. Koomey, E. Masanet and D. Sartor, "2024 United State Data Center Energy Usage Report," Lawrence Berkeley National Laboratory, Berkeley, CA, 2024.

[5] S. Chen, "Data centres will use twice as much energy by 2030--driven by AI," 10 April 2025. [Online]. Available: https://www.nature.com/articles/d41586-025-01113-z. [Accessed 2 January 2026].

[6] S. Krishnan, S. V. Garimella, G. M. Chrysler and a. R. V. Mahajan, "Towards a Thermal Moore's Law," *IEEE Transactions on Advanced Packaging,* vol. 30, no. 3, 2007.

[7] J. R. Thome, "Boiling in microchannels: a review of experiment and theory," *International Journal of Heat and Fluid Flow,* vol. 25, no. 2, pp. 128-139, 2004.

[8] S. G. Kandlikar, S. Garimella, D. Li, S. Colin and M. R. King, Heat Transfer and Fluid Flow in Minichannels and Microchannels, Elsevier, 2006.

[9] Q. Wang, S. Ozguc and R. W. Bonner, "A Practical Metric for Cold Plate Thermal Performance in Two-Phase Direct-to-Chip Cooling," in *2025 41st Semiconductor Thermal Measurement, Modeling & Management Symposium (SEMI-THERM)*, San Jose, CA, USA, 2025.

[10] X. Yuan, X. Zhou, Y. Pan, R. Kosonen, H. Cai, Y. Gao and Y. Wang, "Phase change cooling in data centers: A review," *Energy and Buildings,* vol. 236, 2021.

[11] I. Mudawar, "Recent Advances in High-Flux, Two-Phase Thermal Management," *Journal of Thermal Science and Engineering Applications,* vol. 5, no. 2, 2013.

[12] C. H. Hoang, S. Rangarajan, Y. Manaserh, M. Tradat and G. Mohsenian, "A Review of Recent Developments in Pumped Two-Phase Cooling Technologies for Electronic Devices," *IEEE Transactions on Components, Packaging, and Manufacturing Technology,* vol. 11, no. 10, pp. 1565-1582, 2021.

[13] Q. Wang and R. W. Bonner, "High-Performance Two-Phase Cooling under Different Cold Plate Orientations," in *2025 OCP EMEA Summit Future Technologies Symposium*, Dublin, Ireland, 2025.

[14] S. Ozguc, Q. Wang, A. Narayanan, J. Moore and R. W. Bonner, "Design Optimization of Manifold Integrated Skived Cold Plates for Two-Phase Flow-Boiling," in *2025 41st Semiconductor Thermal Measurement, Modeling & Management Symposium (SEMI-THERM)*, San Jose, CA, USA, 2025.

[15] Q. Wang, S. Ozguc and R. W. Bonner, "Systematic Analysis of Thermal Resistances in data Center Two-Phase Direct-to-Chip Cooling," in *Proceedings of the ASME 2025 International Technical Conference and Exhibition on Packaging and Integration of Electronic and Photonic Microsystems*, Anaheim, CA, USA, 2025.

[16] Q. Wang, D. P. Kulkarni, R. W. Bonner and J. C. Gulick, "Universal Direct-to-Chip Cold Plates for Single- and Two-Phase Cooling," 17 October 2024. [Online]. Available: https://accelsius.com/wp-content/uploads/A-Server-Level-Test-System-A2-03-1-049.pdf. [Accessed 5 January 2025].

[17] G. Prakash Narayan, J. H. Lienhard and S. M. Zubair, "Entropy generation minimization of combined heat and mass transfer devices," *International Journal of Thermal Sciences,* vol. 49, no. 10, pp. 2057-2066, 2010.

[18] Vertiv Group Corp., "Enhanced Efficiency in Data Center with Elevated Return Air Temperature," 2020. [Online]. Available: https://www.vertiv.com/en-asia/about/news-and-insights/articles/white-papers/enhanced-efficiency-in-data-center-with-elevated-return-air-temperature/. [Accessed 5 January 2025].

[19] Accelsius LLC, "Jacobs 10MW Data Center Reference Designs," 2026. [Online]. Available: https://accelsius.com/wp-content/uploads/Jacobs-TCO-Infographic-V121725.pdf. [Accessed 27 February 2026].