On-Chip Cooling for AI Data Centers using Rack-Level Two-Phase Loop Thermosyphons: A Promising Approach for Low Thermal Resistance

Manohar Bongarala
Nokia Bell Labs
New Providence, NJ, USA
manohar.bongarala@nokia-bell-labs.com

Syed Faisal
Nokia Bell Labs
New Providence, NJ, USA
syed.faisal@nokia-bell-labs.com

Rishav Roy
Nokia Bell Labs
New Providence, NJ, USA
rishav.roy@nokia-bell-labs.com

David J. Apigo
Nokia Bell Labs
New Providence, NJ, USA
david. apigo@nokia-bell-labs.com

Sarwesh N. Parbat
Nokia Bell Labs
New Providence, NJ, USA
sarwesh.parbat@nokia-bell-labs.com

Yang Liu
Nokia Bell Labs
New Providence, NJ, USA
yang.3.liu@nokia-bell-labs.com

Todd Salamon
Nokia Bell Labs
New Providence, NJ, USA
todd.salamon@nokia-bell-labs.com

Mark Earnshaw
Nokia Bell Labs
New Providence, NJ, USA
mark.earnshaw@nokia-bell-labs.com

***Abstract*— As artificial intelligence workloads continue to increase rack power density and chip heat flux, liquid cooling is becoming essential for maintaining reliable thermal operation in data centers. Two-phase cooling provides high heat transfer coefficients using the latent heat of vaporization, enabling large heat loads to be removed with low chip-to-coolant temperature rise. Loop thermosyphons extend these advantages by using buoyancy-driven circulation, offering a path toward higher efficiency, improved reliability and reduced complexity due to the absence of rack-level pumps. However, their practical use in rack-level systems remains challenged by horizontal chip orientations, sufficiency in parallel flow distribution under non-uniform heat loads, and challenges to scaling to multiple sleds while dissipating high heat densities. This newsletter summarizes three experimental demonstrations that address these concerns across increasing levels of system complexity: a single horizontal evaporator, a single sled with two parallel evaporators, and a three-sled rack-level thermosyphon. Together, the results show low overall thermal resistance, stable operation under severe power imbalance including cases where one heater is turned off, and multi-kilowatt heat dissipation in a 1U-compatible architecture. These findings indicate that passive two-phase thermosyphons are a promising high heat-transfer capacity and high energy-efficiency cooling approach for future aggressive power density AI data centers.**

Keywords—two-phase cooling, loop thermosyphon, passive cooling, natural circulation, electronics cooling, server cooling, AI data centers, rack-level thermal management, R1233zd(E), parallel evaporators

I. Introduction

The rapid adoption of artificial intelligence has sharply increased the demand for compute capacity in both training and inference. These workloads are executed on densely integrated accelerators, traditionally GPUs, and increasingly custom AI ASICs, that must be placed in close proximity to reduce communication latency. Rack-scale AI systems now illustrate the high-heat-density thermal challenge: NVIDIA's GB200 NVL72, for example, integrates 72 Blackwell GPUs rated at approximately 1.2 kW each and 36 Grace CPUs rated at approximately 300 W each in a single rack, producing about 97 kW of direct chip heat dissipation and a total rack power of roughly 120 kW including the power from auxiliary components [1]. These heat loads must be removed while maintaining chip temperatures below reliability limits, commonly near 95-105 °C depending on device specification. Conventional air cooling with finned heat sinks becomes impractical at these heat densities, which is why direct liquid cooling is now being adopted widely.

The efficacy of a cooling solution is governed by both the coolant's ability to carry heat at acceptable flow rate and pressure drop, and the ease of transferring heat from the chip to the coolant with minimal temperature rise. Single-phase liquid cooling, often using water or water-glycol mixtures, offers much higher volumetric heat capacity than air; for example, water absorbs about 42 kJ/liter across a 10 K temperature rise compared to only 12 J/liter for air. Two-phase cooling instead leverages latent heat, allowing heat absorption at nearly constant saturation temperature. Low-GWP, low-ODP refrigerants such as R1233zd(E) are attractive because their saturation pressure and temperature can be tuned for chip-cooling applications [2]. With a latent heat of about 190 kJ/kg near room-temperature saturation, R1233zd(E) can absorb roughly 4.5 times more energy by vaporization than water can absorb sensibly over a 10 K rise at the same mass flow rate. This heat-carrying advantage is complemented by lower chip-to-coolant thermal resistance, where thermal resistance is the temperature rise per unit heat load, reported in K/W. Heat flows through a series of resistances: spreading in silicon and package materials, conduction through thermal interface materials and lids, conduction through the cold plate or heat sink, and finally convection to the coolant. The final convective resistance is commonly characterized by the heat transfer coefficient, reported in W/m²-K, whose typical values increase from only 10–20 W/m²-K for forced air to about 100–20,000 W/m²-K for single-phase liquids and about 2,500–100,000 W/m²-K for boiling or condensation [3]. These higher coefficients arise because bubble agitation renews colder liquid near the wall while evaporation at the contact line and microlayer removes heat directly as latent heat [4], [5]. Although the high end of the heat transfer coefficients cited above are typically associated with water based two-phase systems, the local heat transfer coefficients for water-based single-phase cold plates and two-phase refrigerant-based cold plates can be comparable. Therefore, the advantage of two-phase refrigerant cooling is not simply a higher local heat-transfer coefficient than single-phase, but the combination of nearly isothermal latent-heat absorption, lower required mass flow for a given heat load, and the potential for passive circulation and higher-temperature heat rejection.

The higher heat transfer ability of two-phase cooling has a cascading effect in improving the energy efficiency of the cooling approach. By reducing the thermal resistance from the chip to the coolant, two-phase heat sinks also allow a higher coolant temperature. This higher coolant temperature enables simpler and low-energy systems such as dry coolers to reject heat directly to the ambient. These benefits, however, are realized only when the two-phase loop is designed to manage flow instabilities, contain significant dryout margins, and incorporate condenser-side heat rejection. In addition, evaporation results in volumetric expansion within the loop, providing a self-pumping mechanism that can circulate the working fluid without a mechanical pump. This passive circulation removes pump-related reliability concerns and costs, making two-phase cooling even more appealing.

In a two-phase loop thermosyphon, heat added at the evaporator generates vapor that creates a liquid-vapor mixture which naturally rises through a tube (riser) to the condenser located at a higher elevation (see Figure 1). Heat is then rejected at the condenser, where the vapor condenses back to liquid, and the denser liquid falls through another tube (downcomer) to

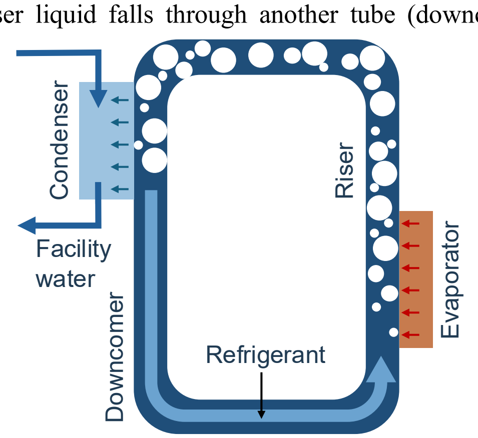

Fig. 1. Schematic of a two-phase loop thermosyphon showing the flow of the working fluid- vapor generated by the heat from the evaporator travels up along the riser and condenses to liquid in the condenser, which then assisted by gravity returns through the downcomer.

provide the hydrostatic head that sustains circulation. This architecture is attractive because it combines latent-heat transport with passive circulation, enabling high heat-transfer capacity, no pumping requirement, reduced maintenance needs, quieter operation, and improved reliability compared to actively pumped cooling loops.

Thermosyphons have therefore been examined across a wide range of heat-transport applications, including solar collectors [6], power electronics [7], nuclear reactor cooling [8], and electronics thermal management [9], [10]. The nuclear reactor cooling literature provides a substantial foundation for understanding natural-circulation behavior, but the public literature on thermosyphon-based on-chip electronics cooling remains comparatively limited, even as commercial loop-thermosyphon solutions are already being offered. Prior demonstrations have shown the feasibility of single-chip cooling [9], [11], [12], [13] and, in some cases, multi-chip [14] or cascaded thermosyphon arrangements [15], [16]. However, open reports on a rack-level architecture that passively supplies liquid to multiple horizontally-mounted sleds within one shared thermosyphon loop are still sparse. This distinction is important because horizontal sled deployment further complicates circulation because vapor removal and liquid replenishment must occur reliably without the vertical orientation that typically facilitates gravity-driven unidirectional flow in loop. Further, rack integration introduces system-level issues that are not fully captured by single-device studies, including flow stability, flow maldistribution among parallel branches, and implications of heat sources operating at different power levels. When there are parallel two-phase flow paths, the total pressure drop across each branch must be nearly the same, but the flow rate and exit vapor quality in each branch can differ. Because pressure drop decreases with increase in mass flow rate in two-phase region, a small perturbation in one branch can grow into flow maldistribution, where one evaporator receives most of the liquid while another receives very little. This behavior is commonly associated with two-phase Ledinegg-type instability, and the resulting maldistribution have been documented in boiling channels [17]. The consequence for electronics cooling is severe as a flow-starved evaporator can experience reduced liquid replenishment, resulting in local dryout and rapid wall-temperature rise. Therefore, parallel thermosyphon architectures must be evaluated not only for total heat-transport capacity but also for the stability and uniformity of liquid delivery under unequal heat loads. This gap in the literature is important to address because the industry is moving rapidly toward liquid-cooled, high-density racks, while the public research on passive two-phase rack loops is still emerging.

In this newsletter, we summarize recent findings from a sequence of loop-thermosyphon prototypes scaled from a single horizontal evaporator to one-sled and three-sled server configurations [18], [19], [20]. The discussion is organized around the practical questions that determine whether passive two-phase thermosyphons can become a credible rack-level cooling technology for artificial intelligence data centers: whether stable natural circulation can be maintained with horizontal heat sinks, whether parallel evaporators can tolerate large chip-to-chip power imbalance without dryout, liquid starvation, or severe flow maldistribution, and whether multi-

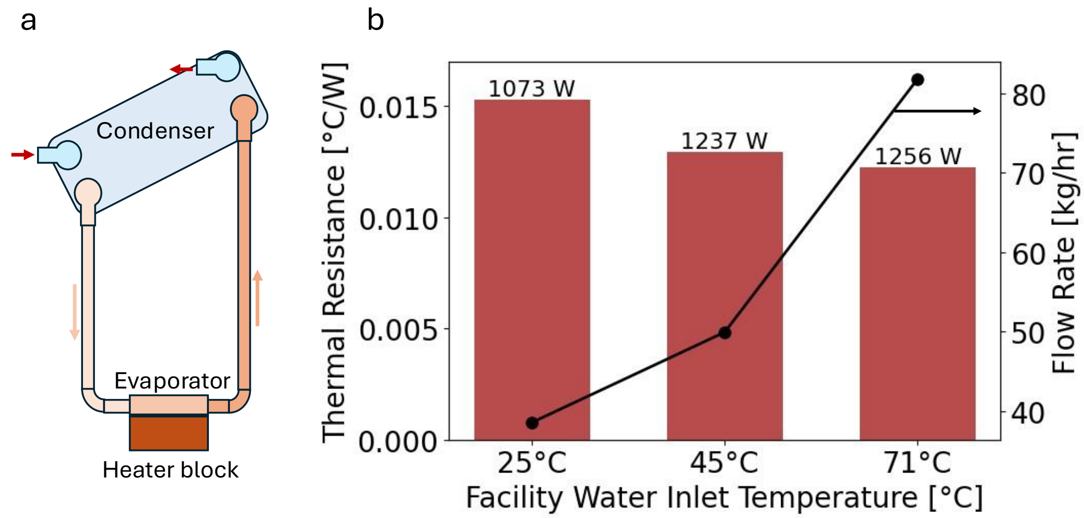

Fig. 2. (a) Schematic of the single-evaporator thermosyphon. (b) Thermal resistance from evaporator to facility water inlet temperature, shown as bar plots; and natural-circulation flow rate in the thermosyphon, shown as line plots, at three coolant inlet temperatures. The resistance presented here includes the entire stack from chip to the facility water inlet involving the TIM, evaporator, thermosyphon and condenser.

sled operation can achieve kilowatt-scale heat dissipation within a 1U-compatible server form factor.

II. Prototype Demonstration Overview

To evaluate the feasibility of thermosyphon-based cooling at the rack level, the following sections present experimental results from three prototype demonstrations with progressively increasing system complexity. These experiments progress from a 76 mm × 76 mm heat source, to a single-sled with 76 mm × 76 mm and 40 mm × 40 mm heat sources connected fluidically in parallel, and finally to a three-sled configuration in which each sled contains the same pair of heat sources.

Across the prototypes, R1233zd(E) was used as the working fluid. Heat was supplied through copper heater blocks, which were in turn attached to finned copper microchannel heat sinks serving as evaporators. The heat was eventually rejected through a liquid-to-liquid brazed-plate condenser cooled by a water-glycol secondary loop (facility water loop). The thermosyphon architecture consisted of horizontal evaporators connected between a downcomer and a riser, with boiling in the evaporators generating the two-phase mixture that drives natural circulation toward the elevated condenser. Thermocouples were used to measure temperatures at the heater blocks, within the two-phase loop, and on the secondary side of the condenser, while ultrasonic flow meters were used to measure the flow rate within the loop. Detailed description of the experimental setups, measurement techniques and data processing are presented in respective papers [18], [19], [20].

III. Single-Evaporator Demonstration: Establishing Low Thermal Resistance And Stable Horizontal Operation

A single-evaporator prototype was first built to test the simplest form of rack-level thermosyphon operation and to determine whether stable circulation and low thermal-resistance metrics could be achieved with a horizontal evaporator. The schematic of this baseline loop is shown in Figure 2 a. Experiments were conducted by varying the condenser secondary-side conditions, specifically the water-glycol coolant flow rate and the coolant inlet temperature set by the chiller. A horizontally-oriented copper microchannel heat sink with a 76 mm × 76 mm heated footprint and an 18 mm height was used to ensure compatibility with the vertical envelope of a 1U server sled. The loop was charged with R1233zd(E) at a 29.6% volumetric fill ratio and dissipated more than 1.5 kW without mechanical pumping. The asymmetric inlet and outlet piping, the plenum geometry, and the condenser orientation were selected to encourage startup and sustain unidirectional circulation over a broad range of heat inputs, approximately 200 W to 1500 W.

Figure 2 b summarizes the thermal-fluid performance at three condenser secondary-side inlet (facility water inlet) temperatures. Increasing the condenser secondary-side inlet temperature from 25 °C to 71 °C increased the loop flow rate from about 38 kg/hr to 82 kg/hr at comparable high heat loads and reduced the measured chip-to-condenser-inlet thermal resistance from approximately 0.015 K/W to 0.012 K/W. This behavior indicates that elevated heat-rejection temperature can be beneficial for this working fluid and loop architecture because the reduced liquid and vapor viscosities at higher saturation temperature lower the loop frictional pressure drop and strengthen natural circulation. Across the experimental matrix, the maximum heat dissipated was 1533 W over the 76 mm × 76 mm footprint, corresponding to a heat flux of 26.54 W/cm² and an overall thermal resistance of 0.0129 K/W.

IV. Single-Sled Demonstration: Flow Distribution Under Skewed Power Profiles

The second demonstration introduced two heat sources within a single half-width server sled by placing a 76 mm × 76 mm evaporator and a 40 mm × 40 mm evaporator on a new

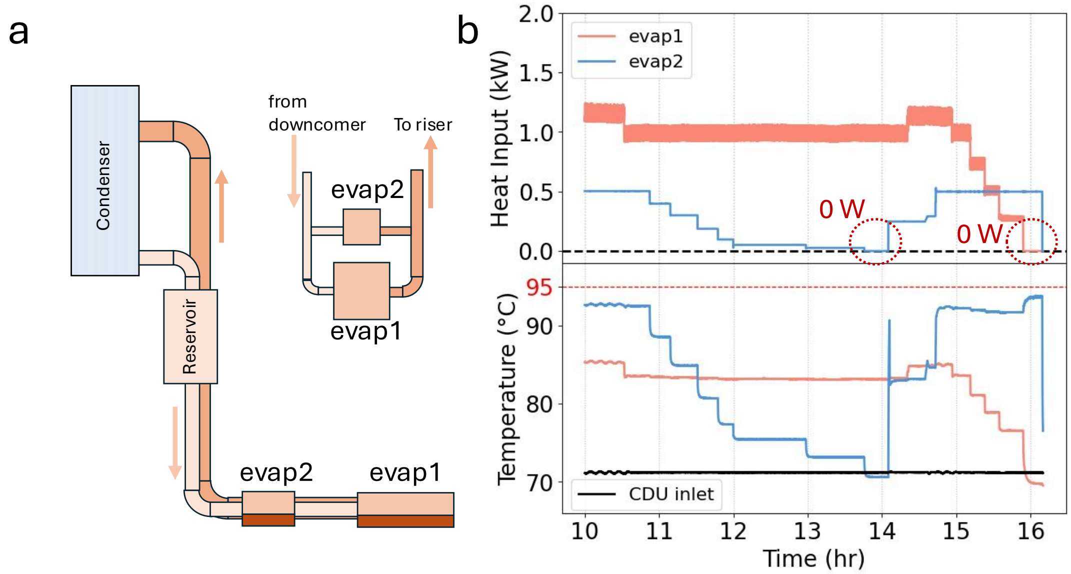

Fig. 3. (a) Schematic of the single-sled thermosyphon with two parallel evaporators connected between common inlet and outlet headers. A reservoir was installed at the top of the downcomer to provide consistent pressure head. (b) Thermal response of the large (evap1) and small (evap2) evaporators under skewed heat input ratios at a facility-water inlet temperature of 71 °C.

thermosyphon loop with larger piping. This arrangement reflects a practical server environment, where devices with different functions and power ratings can be mounted on the same board. The larger and the smaller evaporators, designed for heat inputs up to 1000 W and 500 W, respectively, were intentionally operated on different power scales. The sled schematic is shown in Figure 3 a. Rather than connecting the evaporators in series, where they would be in the same hydraulic path, the two evaporators were arranged in parallel. In a series configuration, the same flow would pass sequentially through each evaporator, eliminating the risk of flow maldistribution among branches. In the parallel configuration used here, both evaporator inlets were connected to the downcomer, and both outlets were connected to the riser. While series connection offers a conservative approach by avoiding the flow maldistribution, it also comes at the expense of increased pressure drop. Therefore, both evaporator inlets were connected to the downcomer, and both outlets were connected to the riser. This parallel configuration reduces sled-level pressure drop and enables higher total heat dissipation, but it also introduces the possibility of flow maldistribution, liquid starvation, and dryout.

To examine this risk, the experiment was conducted as a steady-state power-ratio study. At each operating condition, the heat inputs were changed in steps and the system was allowed to reach steady state before data were recorded. At a CDU inlet temperature of 71 °C, the heat input to the large and small evaporators initially started from 1 kW and 500 W, respectively. The heat input to the smaller evaporator was then reduced stepwise to zero while the larger evaporator dissipated 1 kW. After this sequence, the power to the smaller evaporator was increased in steps back to 500 W. Thereafter, the small evaporator power input was held fixed at 500 W and the larger evaporator was subsequently ramped down to zero in steps. This procedure created a broad set of skewed heat-input ratios, including cases in which one evaporator was driven at high power while the other was operated at much lower power or completely turned off. Figure 3 b highlights the thermal response of the system, showing the temperatures of the larger (evap1) and small (evap2) evaporators subject to the various heat loads. Even under asymmetric heat-input conditions, including cases where one heater block was completely off, the thermosyphon continued to supply sufficient liquid to both evaporators and maintained stable operation. The maximum measured heater-block temperatures remained below 95 °C despite the high 71 °C facility-water temperature, demonstrating strong cooling performance for the rack-level thermosyphon stack. Under total heat loads up to 1.5 kW per sled with R1233zd(E), the pumpless loop generally supplied liquid in proportion to the heat dissipation and avoided dryout under elevated-temperature operation. This result is important because real server sleds commonly contain chips operating at different power levels; the experiment therefore shows that a parallel evaporator thermosyphon can tolerate large power imbalance without losing stable cooling performance.

V. Three-Sled Demonstration: Scaling Toward Rack-Level Cooling

The third experiment scaled the half-width sled architecture from one sled to three vertically stacked sleds connected to the same condenser, downcomer, and common riser as the single-sled experiment. Each sled contained one 76 mm × 76 mm evaporator and one 40 mm × 40 mm evaporator connected in parallel between the downcomer and the riser. The large-to-small evaporator power ratio was fixed at 2:1, and each half-width sled was operated at powers up to ~1.7 kW preserving the per-sled loading of the single-sled experiment while increasing the total system heat load. The three-sled assembly was designed around a 1U-compatible half-width sled envelope, with the top

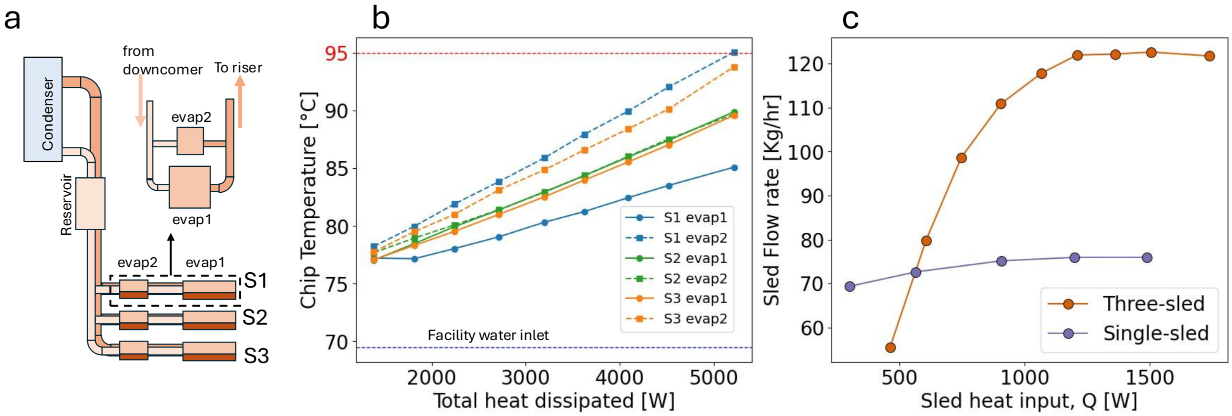

Fig. 4. (a) Schematic of the three-sled rack-level thermosyphon with three vertically stacked half-width sleds connected to a common condenser, downcomer, and riser. (b) Heater-block temperatures of all large and small evaporators as a function of total heat dissipated by the cooling solution. (c) Per sled flow rate as a function of sled heat input for both three and single sled thermosyphons. The per sled flow rate in three sled case is obtained by assuming equal distribution of the total measured flow rate across all the sleds.

sled carrying additional pressure and temperature instrumentation; therefore, power-density estimates exclude the extra volume associated with those sensors. (see Figure 4 a for an illustrative schematic of the three sled configuration)

Figure 4 b plots the heater-block temperatures for all evaporators as a function of the total heat dissipated by the cooling solution. At the system level, the three-sled thermosyphon dissipated 5.2 kW while maintaining a mean chip-to-facility water inlet thermal resistance of 0.016 K/W across the large evaporators and an effective rack-level power density of 3.5 kW per 1U sled (extrapolated by 2× half-width density). All the temperatures of the heater blocks remained under 95 °C with peak temperature of 95 °C for the smaller evaporator in sled 1 at 576 W. These results are significant because the three-sled evaporator scaled the approximate heat load per sled as the single-sled configuration, with three times the total heat rejection, through a unified passive loop. The thermal data also indicate that the evaporators remained well below their operating limits: the heater-block temperatures increased gradually with heat input, and no sharp temperature excursions characteristic of dryout were observed within the tested power range.

The fluidic response in Figure 4 c was particularly important for assessing scalability. The measured per-sled flow rate in the three-sled configuration reached approximately 122 kg/hr per sled, substantially higher than the approximately 76 kg/hr measured in the single-sled experiment under comparable per-sled heat input. This trend was contrary to the intuitive expectation that adding sleds would increase frictional pressure drop due to more vapor, and suppress natural circulation. Instead, the measured pressure-drop and flow-rate trends suggest that the larger vapor generation rate in the three-sled case increased the vapor area fraction in the common riser. This resulted in reduction of effective mixture density which lowered the riser hydrostatic pressure drop penalty and strengthened the buoyancy head available to drive circulation. At higher heat inputs, the flow rate did not show a strong reduction, which further suggests that the loop had not approached a circulation-limited operating condition within the tested range.

Although 3.5 kW per 1U sled is already a high heat density relative to traditional server platforms, emerging AI/ML systems require even higher cooling capacity. The present three-sled data therefore should be interpreted as a successful scaling demonstration rather than a maximum-capacity test. The absence of dryout signatures, the sustained increase in circulation with heat input, and the previously demonstrated capability of the same class of evaporator to dissipate more than 2.6 kW using a multi-chip thermal test vehicle [21] suggest that additional headroom remains for higher-power operation.

VI. Conclusions and Outlook

Two-phase cooling is emerging as a promising path for next-generation AI hardware because latent heat transport and high boiling heat-transfer coefficients can remove larger heat loads at lower chip-to-coolant temperature differences than conventional air or single-phase liquid cooling. Passive two-phase thermosyphons are especially attractive because in addition to thermal-performance advantage they avoid pump-related reliability and energy-consumption concerns. The unanswered challenge is not whether thermosyphons can cool a single device, but whether they can scale to the horizontal, parallel, and nonuniform operating conditions expected in server sleds and rack-level systems.

The demonstrations summarized here address several of those scaling questions. The single-evaporator loop established low overall thermal resistance and stable natural circulation in a horizontal, 1U-compatible evaporator. The single-sled experiment then showed that two parallel evaporators could tolerate highly skewed heat-input ratios, including cases where one heat source was turned off, without liquid starvation or dryout. Finally, the three-sled prototype extended the same passive loop architecture across multiple sleds, maintaining stable operation, low thermal resistance, and increasing circulation as total vapor generation increased.

Together, these results indicate that a unified rack-level thermosyphon can passively distribute liquid across multiple parallel evaporators and sleds while preserving stable circulation and strong thermal performance. This makes passive two-phase thermosyphon cooling a promising route toward reliable, high-density, and energy-efficient thermal management for future AI data centers, with further work needed to quantify ultimate power limits, long-duration reliability, and full rack integration under realistic workload transients.

Acknowledgment

The information, data, or work presented herein was funded in part by the Advanced Research Projects Agency-Energy (ARPA-E), U.S. Department of Energy, under Award Number DE-AR0001597. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof. The authors thank Dr. Peter de Bock, Dr. Jungho Kim and Dr. Thomas Bress for their feedback on the project. The authors also thank Prof. Yogendra Joshi for providing comments on the newsletter.

References

[1] "GB200 Datasheet." Accessed: Jun. 19, 2026. [Online]. Available: https://www.supermicro.com/datasheet/datasheet_SuperCluster_GB200_NVL72.pdf

[2] R. Akasaka and E. W. Lemmon, "An International Standard Formulation for trans-1-Chloro-3,3,3-trifluoroprop-1-ene [R1233zd(E)] Covering Temperatures from the Triple-Point Temperature to 450 K and Pressures up to 100 MPa," *J. Phys. Chem. Ref. Data*, vol. 51, no. 2, p. 023101, May 2022, doi: 10.1063/5.0083026.

[3] T. L. Bergman and F. P. Incropera, Eds., *Fundamentals of heat and mass transfer*, 7th ed. Hoboken, NJ: Wiley, 2011.

[4] V. P. Carey, *Liquid-Vapor Phase-Change Phenomena: An Introduction to the Thermophysics of Vaporization and Condensation Processes in Heat Transfer Equipment*, 3rd ed. Boca Raton: CRC Press, 2020. doi: 10.1201/9780429082221.

[5] M. Bongarala, J. A. Weibel, and S. V. Garimella, "A method to partition boiling heat transfer mechanisms using synchronous through-substrate high-speed visual and infrared measurements," *Int. J. Heat Mass Transf.*, vol. 226, p. 125516, Jul. 2024, doi: 10.1016/j.ijheatmasstransfer.2024.125516.

[6] L. Serconek Fuso, L. H. R. Cisterna, and M. B. H. Mantelli, "Experimental study of two phase loop thermosyphons for hybrid solar systems," *Energy Convers. Manag.*, vol. 293, p. 117550, Oct. 2023, doi: 10.1016/j.enconman.2023.117550.

[7] F. Cataldo and J. R. Thome, "Experimental evaluation of the thermal performances of a thermosyphon cooling system rejecting heat by natural and forced convection," *Appl. Therm. Eng.*, vol. 127, pp. 1404–1415, Dec. 2017, doi: 10.1016/j.applthermaleng.2017.08.166.

[8] Y. Zvirin, "A review of natural circulation loops in pressurized water reactors and other systems," *Nucl. Eng. Des.*, vol. 67, no. 2, pp. 203–225, Jan. 1982, doi: 10.1016/0029-5493(82)90142-X.

[9] A. Chauhan and S. G. Kandlikar, "Characterization of a dual taper thermosiphon loop for CPU cooling in data centers," *Appl. Therm. Eng.*, vol. 146, pp. 450–458, Jan. 2019, doi: 10.1016/j.applthermaleng.2018.10.010.

[10] G. S. R. B. Armas, G. Rouaze, J. B. Marcinichen, J. R. Thome, and L. W. Zhang, "Air-Cooled Loop Thermosyphon Cooling System for High Heat Load CPUs—Part II: Experimental Results and Validation," *IEEE Trans. Compon. Packag. Manuf. Technol.*, vol. 11, no. 10, pp. 1687–1694, Oct. 2021, doi: 10.1109/TCPMT.2021.3112949.

[11] R. L. Amalfi, J. B. Marcinichen, J. R. Thome, and F. Cataldo, "Design of Passive Two-Phase Thermosyphons for Server Cooling," presented at the ASME 2019 International Technical Conference and Exhibition on Packaging and Integration of Electronic and Photonic Microsystems, American Society of Mechanical Engineers Digital Collection, Dec. 2019. doi: 10.1115/IPACK2019-6386.

[12] A. Pal, Y. K. Joshi, M. H. Beitelmal, C. D. Patel, and T. M. Wenger, "Design and performance evaluation of a compact thermosyphon," *IEEE Trans. Compon. Packag. Technol.*, vol. 25, no. 4, pp. 601–607, Dec. 2002, doi: 10.1109/TCAPT.2002.807997.

[13] L. Yuan, Y. K. Joshi, and W. Nakayama, "Effect of Condenser Location and Tubing Length on the Performance of a Compact Two-Phase Thermosyphon," presented at the ASME 2001 International Mechanical Engineering Congress and Exposition, American Society of Mechanical Engineers Digital Collection, Nov. 2021, pp. 291–299. doi: 10.1115/IMECE2001/HTD-24406.

[14] N. Lamaison, C. L. Ong, J. B. Marcinichen, and J. R. Thome, "Two-phase mini-thermosyphon electronics cooling: Dynamic modeling, experimental validation and application to 2U servers," *Appl. Therm. Eng.*, vol. 110, pp. 481–494, Jan. 2017, doi: 10.1016/j.applthermaleng.2016.08.198.

[15] X. Chen, T. Ding, L. Wang, and Z. Li, "Experimental research and energy effectiveness analysis on chip-level two-stage loop thermosyphon system for data centre free cooling," *Appl. Therm. Eng.*, vol. 245, p. 122767, May 2024, doi: 10.1016/j.applthermaleng.2024.122767.

[16] F. Cataldo, R. L. Amalfi, J. B. Marcinichen, and J. R. Thome, "Implementation of Passive Two-Phase Cooling to an Entire Server Rack," in *2020 19th IEEE Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm)*, Orlando, FL, USA: IEEE, Jul. 2020, pp. 396–401. doi: 10.1109/ITherm45881.2020.9190327.

[17] L. E. O'Neill and I. Mudawar, "Review of two-phase flow instabilities in macro- and micro-channel systems," *Int. J. Heat Mass Transf.*, vol. 157, p. 119738, Aug. 2020, doi: 10.1016/j.ijheatmasstransfer.2020.119738.

[18] M. Bongarala *et al.*, "Thermofluidic Performance of a Two-Phase Loop Thermosyphon for Server Cooling: Effects of Condenser Secondary Side," in *2025 24th IEEE Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm)*, May 2025, pp. 1–9. doi: 10.1109/ITherm55376.2025.11235702.

[19] S. Faisal *et al.*, "Experimental Characterization of a Two-Phase Loop Thermoyshon with Two Parallel Evaporators," in *2026 25th IEEE Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm)*, 2026.

[20] M. Bongarala *et al.*, "Experimental Study of Single- and Multi-Sled Two-phase Loop Thermosyphons for High-Power Server Cooling," in *18th International Heat Transfer Conference*, 2026.

[21] Y. Liu *et al.*, "Experimental Demonstration of High-Power Thermal Test Vehicle using Two-Phase Cooling for AI Datacenters, 5G RAN, and Edge Compute Nodes," in *2025 IEEE 75th Electronic Components and Technology Conference (ECTC)*, Dallas, TX, USA: IEEE, May 2025, pp. 2074–2080. doi: 10.1109/ECTC51687.2025.00355.