Mitigating Flow Maldistribution in Data Center Two-Phase Cooling Systems with Flow Restrictors

Serdar Ozguc, Qingyang Wang, Akshith Narayanan, Richard W. Bonner III
Accelsius

The Need for Two-Phase Cooling in Data Centers

Power usage in data centers account for 1.8% of the overall electricity expenditure in the United States [1] and the cooling infrastructures make up 50% of the total energy consumption of the data centers [2]. Power consumption translates to high operational costs and carbon footprint which can be potentially reduced by implementing higher efficiency thermal management systems. Moreover, the growing power densities in data centers, needed to address the demand for high-performance computing, are beginning to push the thermal limits of conventional air-cooled systems.

Liquid-cooled solutions were shown to significantly reduce the overall power consumption relative to air-cooling [3] and can handle increased power densities due to liquid's higher thermal conductivity and heat capacity. Single-phase liquid-cooling with

Serdar Ozguc
Dr. Serdar Ozguc earned his Ph.D. in Mechanical Engineering from Purdue University, where his research focused on the design and optimization of single- and two-phase microchannel heat sinks. He is currently a Senior R&D Thermal Engineer at Accelsius, where he specializes in the development of advanced two-phase direct-to-chip liquid cooling CDUs for data center applications.

Qingyang Wang
Dr. Qingyang Wang is a Senior R&D Thermal Engineer at Accelsius, where he is developing two-phase direct-to-chip cooling technologies for data centers. Prior to this role, he was an Associate Professor at North China Electric Power University. Dr. Wang obtained his Ph.D. in Mechanical Engineering from the University of California, San Diego, in 2020, and his B.S. in Thermal Engineering from Tsinghua University in 2015. With over 40 peer-reviewed publications in leading journals and conferences, he is recognized for his expertise in multiphase flow and heat transfer, thermal management, and thermal characterization.

Akshith Narayanan
Akshith Narayanan is a recent Masters Graduate in mechanical engineering from Georgia Institute of Technology. He completed his degree with a thesis working on an investigation of near-junction flow boiling of high-power electric vehicle inverters. He joined Accelsius, an emerging two phase direct to chip liquid cooling company looking to provide an elegant two-phase solution for data center high power chips. He has years of experience in heat transfer, fluid mechanics with specific expertise in two phase boiling.

Richard W. Bonner III
Dr. Bonner is a distinguished heat transfer researcher with 20 years in thermal product development, specializing in two-phase cooling. He has authored over fifty papers and holds five U.S. patents, having designed cooling products for 125+ clients in various sectors. He is a former AIChE Transport and Energy Processes Division Director and holds a B.S., M.S., and Ph.D. in Chemical Engineering from Lehigh University.

water-based coolants can achieve good thermal performance due to water's favorable thermal characteristics. However, a minor leak in a water-cooled system can cause a catastrophic electrical failure, due to water's unfavorable electrical characteristics. Two-phase liquid-cooling with refrigerants circumvents this problem where high heat transfer performance is achieved through liquid-to-vapor phase change of the dielectric fluid. Heat fluxes on the order of 1 kW/cm^2 have been dissipated using two-phase cooling [4].

Flow Maldistribution in Two-Phase Cooling

Pumped two-phase cooling has been studied extensively in the literature [5]. However, implementation and testing of two-phase cooling in data center applications are limited. The highly parallel architecture of liquid-cooling loops in server racks can suffer from coolant flow maldistribution between heat generating components. This problem is more prominent in two-phase flows because the difference of pressure drop between liquid and vapor flows can lead to instability. High heat loads on cold plates result in increased vapor generation, leading to a rise in pressure drop. This, in turn, diverts the coolant through the cold plates with lower heat loads, potentially causing device overheating.

Flow restrictors can be adopted upstream of boiling [6] to suppress maldistribution by increasing the liquid line pressure drop relative to the vapor line. However, careful design of restrictors is crucial to effectively mitigate maldistribution without inducing excessive pressure drops that reduce the overall flow rate of the system.

Rack-Level Two-Phase Cooling System

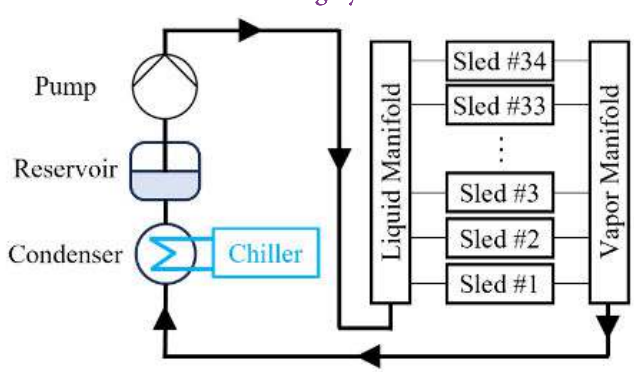

Figure 1: Flow diagram of the investigated server-rack level two-phase flow loop for data center cooling

Pumped two-phase flow loop for a data center rack consisting of 34 server sleds with heat generating components was investigated. A flow diagram of the loop is shown in *Figure 1*. The heat dissipation from each sled is between 0-2 kW and can vary among sleds. Refrigerant R-1233zd(E) was used as the two-phase coolant in the main loop. Pumped coolant entered the sleds as a liquid, absorbed the generated heat, and left as a liquid/vapor mixture with thermodynamic qualities of 0-85%. The sleds were connected to the loop via a liquid manifold upstream and a vapor manifold downstream. The liquid and vapor manifolds are large pipes (1" and 2" inner diameter respectively) connected to the sleds with additional tubing.

Hydrodynamic Maldistribution Model

A numerical model of the liquid manifold, vapor manifold, and the server sleds was developed. In this model, two-phase flow is simplified using the homogenous flow assumption, wherein the vapor bubbles and the surrounding liquid move at the same velocity. The liquid and vapor manifolds are discretized along the length, and the conservation equations for mass, momentum, and energy are applied. The frictional losses in the manifolds are estimated using friction factor correlations for fully developed laminar [7] and turbulent flows [8] in circular channels alongside with a mixture viscosity correlation [9].

Pressure drop between the liquid and vapor manifolds ($\Delta P_{manifold}$) at a discretized position j along the length depends on the pressure drop of the sled components in between (ΔP_{sled}), which include hoses, tubes, fittings, cold plates, etc. and the pressure drop across the flow restrictors ($\Delta P_{restrictor}$) placed upstream of the cold plates as shown in *equation 1*.

$$\Delta P_{manifold} = \Delta P_{sled} + \Delta P_{restrictor} \qquad (1)$$

Sled Pressure Drop and Thermal Resistance

Two-phase pressure drop of an individual sled was experimentally measured to derive an empirical correlation. A flow diagram of the experimental flow loop is shown in *Figure 2*. The investigated server sled consists of two heat generating components that were emulated using two thermal test vehicles (TTV) with heat spreaders and cold plates. Pressure drop and temperature data were collected over ranges of flow rates (5-26 g/s) and total heat inputs (0.2-2 kW), resulting in thermodynamic qualities between 0 and 100%[1]. A total of 104 data points are collected.

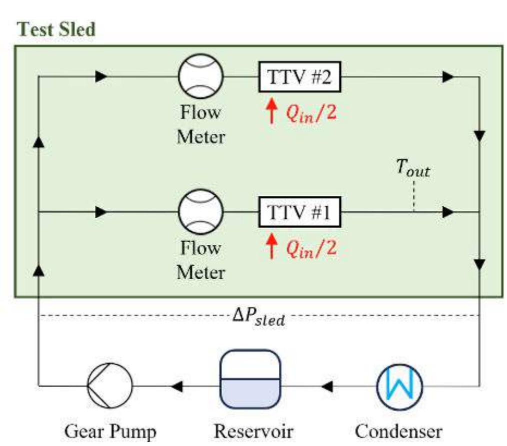

Figure 2: Flow diagram of the experimental flow loop for testing pressure drop across a server-sled

[1] Since the thermodynamic quality is difference between fluid and saturated liquid enthalpies divided by the heat of vaporization: $x = (h - h_f)/h_{fg}$, with superheating, x can be >100%.

A multi-variable, second-order polynomial fit was used to generate an empirical correlation for the pressure drop data. The resulting correlation for sled pressure drop (ΔP_{sled}) matches quite well with the measured data, with most of the predictions having less than 25% error. *Figure 3* shows thermal resistance (normalized by the lowest measured resistance) at exit thermodynamic qualities between 0-100% for all flow rates. The thermal resistance is lowest at 54% thermodynamic quality and sharply increases near 0% and 100%, as expected. Thermal performance improves from single-phase liquid cooling, with near zero thermodynamic qualities, to higher thermodynamic qualities as flow boiling enhances convection. However, thermal resistance increases near x_{sled} = 100% because the generated vapor starts interfering with heat transfer. The two-phase loop shown in *Figure 1* was designed to operate at an exit thermodynamic quality of 70% to provide a margin of safety for reliability. However, individual sleds are expected to have higher qualities due to flow maldistribution. An upper limit of x_{target} = 85% was chosen to ensure acceptable cooling performance in the sleds.

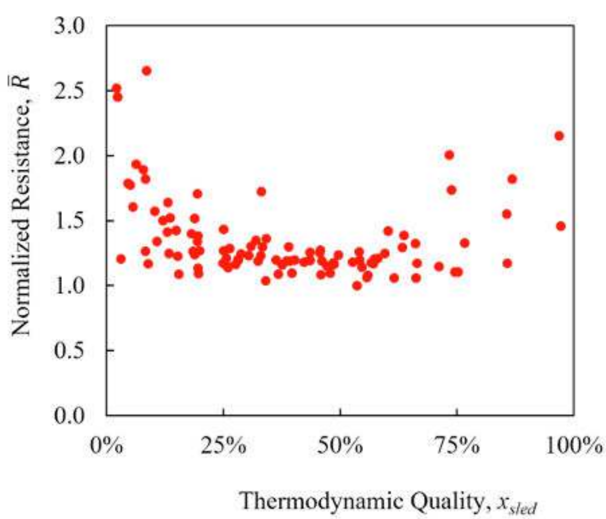

Figure 3: Experimental thermal resistance vs. exit thermodynamic quality for all flow rates tested

Restrictor Pressure Drop

Two restrictors were placed in each sled, positioned upstream of the cold plates. *Equation 2* was used to model the normalized pressure drop across restrictors ($\Delta\bar{P}_{restrictor}$) with respect to the normalized mass flow rate ($\bar{m}$).

$$\Delta\bar{P}_{restrictor} = \alpha\bar{m}^{\beta} \tag{2}$$

Term α is the flow resistance factor and β is the flow scaling exponent. Pressure drop and mass flow rate are normalized so that the flow resistance factor (α) is the ratio of restrictor pressure drop to sled pressure drop under maximum heat load (2 kW). Values of α and β are dictated by the geometry of the restrictor and coolant properties:

- Linear scaling (β = 1) occurs in a long tube with a small diameter where viscous shear dominates the pressure drop.
- For orifice restrictors, pressure drop scales quadratically with mass flow rate (β = 2) due to a momentum dominated flow, and the value of α depends on the orifice size.
- Higher order scaling can be achieved through moving or flexible parts in the restrictors. For example, commercial flow regulators incorporate flexible polymers that constrict the flow area with increasing pressure differential.

Nonuniform Heating

The sleds might be under different loads during operation. Therefore, the two-phase flow was evaluated using the hypothetical heating profile of *equation 3*.

$$Q_{in} = 2\ kW \cdot \left(\frac{y}{H}\right)^2 \tag{3}$$

In this equation, y is the distance from the bottom sled and H is the total height between 34 sleds, which are stacked on top of each other. The sled at the bottom receives no heat while the sled at the top receives 2 kW.

Flow Maldistribution Without Restrictors

The two-phase flow loop was analyzed without flow restrictors to serve as a benchmark. The inlet flow rate to the liquid manifold was selected such that each sled should have an exit quality of 70% in the absence of maldistribution. *Figure 4* shows the resulting exit thermodynamic qualities for each sled, with higher exit thermodynamic quality in the sleds near the top. In addition to the gravitational effects, the flow maldistribution is exacerbated by the nonuniform heating. Since pressure drop increases with vapor generation, the sleds with higher heat input near the top receive less flow. The maximum predicted thermodynamic quality is 176%, which indicates that the cold plate is under dry-out that leads to a high thermal resistance that is unacceptable for data center cooling.

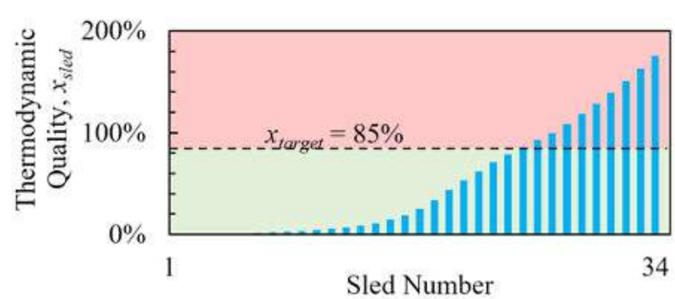

Figure 4: Predicted exit thermodynamic quality distribution under nonuniform heating (*Equation 3*) without flow restrictors

Flow Maldistribution with Orifice Restrictors: An orifice restrictor (β = 2) was first investigated to suppress maldistribution. An insufficiently low α value cannot suppress the maldistribution while an unnecessarily high value will induce a high pressure drop penalty. *Figure 5* shows the thermodynamic quality distribution for $\alpha = 2.0$. The maximum thermodynamic quality is 85% under nonuniform heating. Therefore, all the sleds operate under the maximum thermodynamic quality limit. Flow resistance factor (α) is the ratio of restrictor to sled pressure drop. For the orifice restrictor, the pressure drop across the orifice needs to be roughly twice as much as the pressure drop across the server sled components.

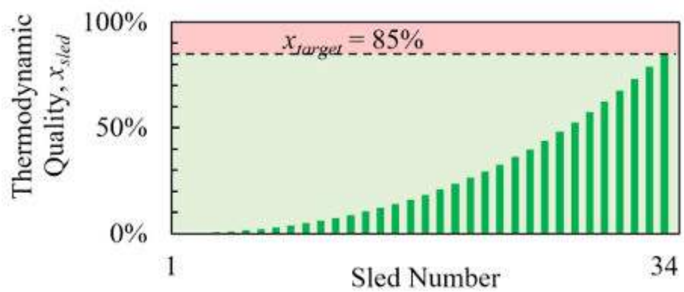

Figure 5: Predicted exit thermodynamic quality distribution under nonuniform heating (*Equation 3*) with orifice restrictors

Advantage of Higher Order Restrictors

The two-phase flow loop was analyzed for a range of β values. For each case, the value of α required to achieve a maximum 85% thermodynamic quality under nonuniform heating was found. Resulting α and β pairs are shown in *Figure 6*. The value of α, and hence the restrictor pressure drop, decreases with increasing β. The increasingly concave pressure drop-mass flow rate response severely punishes maldistribution, pressure drop is reduced at and below the desired flow rate. The value of α converges to 0 as $\beta \to \infty$. An ideal flow regulator can suppress the maldistribution without inducing additional pressure drop to the system at the desired flow rate and maximum heat input.

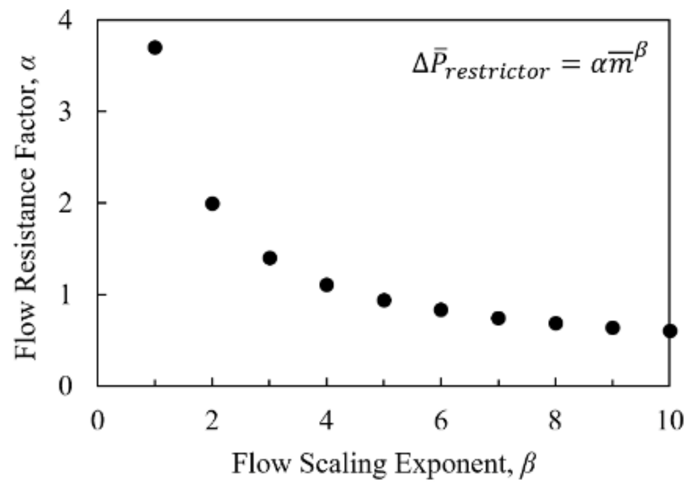

Figure 6: Flow resistance factors to obtain a maximum vapor exit quality of 85% at various flow scaling exponent values

Conclusions

Lower pressure drop restrictors in a pumped two-phase loop can enable a higher flow rate across the system to dissipate more heat. Therefore, restrictors with higher scaling exponents (β), such as flow regulators, are preferable. However, there are practical challenges that make their implementation difficult. First, commercially available flow regulators are significantly more costly compared to off-the-shelf orifice restrictors. Second, commonly adopted orifice restrictors suppress two-phase backflow by inducing high pressure drop upstream, which is not addressed by flow regulators Third, flow regulators need to be highly tailored for the system, which requires a good understanding of the entire pressure drop response of the flow loop under different conditions. If the available flow rate is underestimated, a flow regulator would completely block the additional flow, thereby compromising the added cooling capacity. If the available flow is overestimated, a flow regulator would not effectively suppress maldistribution. An orifice restrictor is more robust to uncertainty, providing reliable suppression to maldistribution when the system characteristics are not fully characterized.

References

[1] M. A. B. Siddik, A. Shehabi, and L. Marston, "The environmental footprint of data centers in the United States," *Environ. Res. Lett.*, vol. 16, no. 6, p. 064017, Jun. 2021, doi: 10.1088/1748-9326/abfba1.

[2] M. Dayarathna, Y. Wen, and R. Fan, "Data Center Energy Consumption Modeling: A Survey," *IEEE Commun. Surv. Tutor.*, vol. 18, no. 1, pp. 732–794, 2016, doi: 10.1109/COMST.2015.2481183.

[3] F. Rebarber, "Quantifying the Impact on PUE and Energy Consumption When Introducing Liquid Cooling Into an Air-cooled Data Center." Accessed: Jan. 01, 2024. [Online]. Available: https://www.vertiv.com/en-emea/about/news-and-insights/articles/blog-posts/quantifying-data-center-pue-when-introducing-liquid-cooling/#:~:text=Impact%20of%20the%20Introduction%20of,compared%20to%20100%25%20air%20cooling

[4] R. K. Mandel, D. G. Bae, and M. M. Ohadi, "Embedded Two-Phase Cooling of High Flux Electronics Via Press-Fit and Bonded FEEDS Coolers," *J. Electron. Packag.*, vol. 140, no. 3, p. 031003, Sep. 2018, doi: 10.1115/1.4039264.

[5] C. H. Hoang et al., “A Review of Recent Developments in Pumped Two-Phase Cooling Technologies for Electronic Devices,” *IEEE Trans. Compon. Packag. Manuf. Technol.*, vol. 11, no. 10, pp. 1565–1582, Oct. 2021, doi: 10.1109/TCPMT.2021.3117572.
[6] W. Chen, T. M. Conboy, G. W. Daines, and D. W. Fogg, “A Robust Two-Phase Pumped Loop with Multiple Evaporators and Multiple Radiators for Spacecraft Applications,” presented at the 47th International Conference on Environmental Systems, Charleston, South Carolina, Jul. 2017.
[7] *F. P. Incropera, D. P. DeWitt, T. L. Bergman, A. S. Lavine, “Fundamentals of Heat and Mass Transfer,” Wiley, 2011.*
[8] V. Mileikovskyi and T. Tkachenko, “Precise Explicit Approximations of the Colebrook-White Equation for Engineering Systems,” in *Proceedings of EcoComfort 2020*, vol. 100, Z. Blikharskyy, Ed., in Lecture Notes in Civil Engineering, vol. 100. , Cham: Springer International Publishing, 2021, pp. 303–310. doi: 10.1007/978-3-030-57340-9_37.
[9] “A, Cicchitti, C. Lombardi, M. Silvestri, G. Soldaini, R. Zavattarelli, ‘Two-phase cooling experiments – pressure drop, heat transfer and burnout measurements,’ Energia Nucleare, 7 (1960) 407-425.”.