Journal of Enhanced Heat Transfer, 29(6):27–39 (2022)

EXPERIMENTAL STUDY DEMONSTRATING PUMPING POWER SAVINGS AT RACK LEVEL USING DYNAMIC COOLING

Pardeep Shahi,* Satyam Saini, Pratik Bansode, Rajesh Kasukurthy, & Dereje Agonafer

Mechanical and Aerospace Engineering Department, The University of Texas at Arlington, 500 W 1st St., Arlington, TX 76019, USA

*Address all correspondence to: Pardeep Shahi, Mechanical and Aerospace Engineering Department, The University of Texas at Arlington, 500 W 1st St., Arlington, TX 76019, USA, E-mail: pardeep.shahi@mavs.uta.edu

Original Manuscript Submitted: 3/17/2022; Final Draft Received: 3/29/2022

Data center energy demands continue to rise with the increasing utilization of high-performance computing applications. As a result, a corresponding increase in power densities at the package level has necessitated advancements in IT equipment thermal management technologies. Cold plate–based liquid cooling allows high heat flux dissipation owing to very high convective heat transfer coefficients. A disadvantage of traditional liquid cooling in data centers is that a constant flow rate is provisioned to servers irrespective of their workloads. The present research investigates, empirically, the concept of provisioning targeted coolant flow rate to a server with a radiator fan arrangement. An air-cooled open compute server was retrofitted with cold plates on and sealed completely to have no ambient interaction with the environment. The cold plates were provisioned with a targeted amount of coolant based on instantaneous CPU workload at different coolant inlet temperatures. The server design was able to keep the temperatures of all the components like dual in-line memory modules, platform controller hub, and CPUs below their critical temperature values. The results show that a maximum cooling percentage reduction of 24% is achieved at maximum flow rate and minimum CPU utilization. The results also show that cooling power savings are also reduced at higher CPU workloads and coolant inlet temperatures.

KEY WORDS: *data center, liquid cooling, electronics cooling, cold plate*

1. INTRODUCTION

A data center is a dedicated space that centralizes an organization's IT equipment and support infrastructure for computing, storing, and transmitting digital information when required. The size of a data center facility may vary from a few tens of servers, referred to as server rooms, to hundreds of thousand servers, known as enterprise or hyperscale data centers. Some of the largest network service providers in the world like Google, Facebook, Yahoo, Amazon, eBay, Microsoft, and Twitter own large-scale enterprise-class data centers. Typically, an enterprise-class data center consists of more than 100,000 servers with a facility size of more than 5000 sq. ft. and can consume a total power anywhere between 1 and 500 MW (Fernandes, 2015).

Data centers consume approximately 2% of the total electricity consumed across the United States (Fernandes et al., 2018; Coles and Greenberg, 2014). Cooling energy consumption costs

account for 30%–50% of the total IT power consumed in a typical data center (Bansode et al., 2019; Jennings, 2015). Most data centers are traditionally air-cooled, which has disadvantages like low heat carrying capacity, adding to energy consumption because of the chillers and compressors, hot spot formation, issues associated with aisle containment, and so on. Liquid cooling addresses these critical issues related to traditional air cooling of servers due to its ability to absorb higher thermal masses and lower coolant transport energy requirements. These cooling methods use water in dielectric form or other commercially used dielectric refrigerants as a cooling medium running through cold plates or rear door heat exchangers and demonstrate the benefits of data center liquid cooling strategy (Amoli et al., 2012; Ellsworth and Iyenger, 2009). Water can improve cooling efficiency and yield further cost savings using higher inlet fluid temperatures and utilizing waste heat for relevant applications within the facility (Zimmerman et al., 2012). A further enhancement in the thermal performance of liquid cooling strategies can also be achieved by the addition of high thermal conductivity nanoparticles to the coolant in direct and indirect liquid cooling methodologies (Niazmand et al., 2020; Shahi et al., 2020).

Copper cold plates have long been used as a method of cooling high-powered devices similar to the thermal conduction module (Chu et al., 1982). Further, moving the cooling source closer to the heat-generating component helps to reduce hot spots within the data center white space, thereby improving cooling effectiveness. ASHRAE TC 9.9 (ASHRAE, 2011) provides ample data to design guidelines, codes, and standards for implementing water-cooled data centers and mission-critical environments. Since the use of liquid cooling at the server and rack level is being adopted widely for high-performance computing applications, the committee has introduced liquid cooling classes as well with varying coolant temperature ranges as shown in Table 1. These temperature ranges of the facility supply-side water are defined based on the type of infrastructure cooling design that is to be used (Sahini, 2017).

A large portion of the power consumed in typical air-cooled data centers can be attributed to the chiller plant, to provide chilled water to the data center, and by computer room air conditioning units and computer room air handlers, to cool the data center white space. Significant savings in energy consumption and associated costs are possible by reducing the dependence on vapor compression–based cooling systems (Rubenstein et al., 2010; David et al., 2013). This can be achieved by utilizing a cooling strategy that uses coolant at temperatures above external

TABLE 1: ASHRAE liquid cooling classes

Classes	Typical infrastructure design		Facility supply water temperature (C)	IT equipment availability
	Main cooling equipment	Supplemental cooling equipment		
W1	Chiller/cooling tower	Water-side economizer chiller	2–17	Available
W2			2–27	
W3	Cooling tower	Chiller	2–32	Not generally available, dependent on future demand
W4	Water-side economizer (with dry cooler or cooling tower)	Nothing	2–45	
W5	Building system heating	Cooling tower	> 45	Specialized systems

ambient conditions, between 27.5°C and 45°C, to cool primary heat-generating components like CPUs and GPUs. The remaining processing components such as hard-disk drives (HDDs), platform controller hub (PCH), and other auxiliary components can be cooled using ambient air (ASHRAE, 2021). The issue of operating IT equipment at elevated temperatures for power savings can be alleviated by resorting to warm-water cooling. This method also eliminates the requirement of mechanical chillers and compressors as used in air cooling and uses a water-side economizer to supplement cooling. A comprehensive literature review of liquid cooling technologies was conducted for liquid cooling thermal management and the advantages of warm water cooling (Copeland, 2005). According to ASHRAE liquid cooling guidelines, the W4 envelope, as shown in Table 1, allows the cooling system to operate within a wide range of 2°C–45°C in the case of direct contact liquid cooling, which improves energy efficiencies (Patterson, 2011).

In a typical data center, not all the servers work at the same CPU utilization all the time. For liquid-cooled data centers, rack or row-based centralized pumps in coolant distribution units typically pump a constant coolant flow rate to all the servers irrespective of the server utilization. While the flow rate redundancy is understood for air cooling, where the cooling latency factor due to fluctuating workloads is higher owing to the low heat removal capacity of air, this is not true for direct-to-chip liquid cooling. CPU utilization-based coolant provisioning has been assessed at the cold plate level using numerical methods. In this study, a jet impingement cold plate was proposed, composed of four different jets that dynamically controlled the coolant flow rate to different parts of the cold plate based on the case temperature of the multichip module being cooled (Fernandes et al., 2015). Dynamic cooling has also been studied for air-cooled-type data centers where the floor tile dampers are actively modulated to direct the airflow to servers with higher workloads (Khalili et al., 2019). Other approaches to save power and dynamically manage CPU thermal loads include workload dispatching and thermally aware workload routes inside CPUs (Xu et al., 2014).

An attempt, thus, can be made to save this redundant pumping power at the rack and data center level by dynamic or target delivery of the coolant to the servers based on their workload utilization. This would mean that servers operating at higher CPU utilization (dissipating higher heat) would be provisioned with a higher coolant flow rate and those operating at lower CPU workload will receive a lower flow rate. By doing so, the redundant flow rate, and hence pump power per rack or row, can be saved. These savings can be significant for hyperscale data centers where many such pumps are constantly operating. The present investigation combines the advantages of warm water liquid cooling and hybrid cooling and explores the effect of dynamic or target coolant delivery in a server. A third-generation open-compute air-cooled Winterfell server was retrofitted with cold plates in a closed loop with a two fan–radiator arrangement. Twenty percent ethylene glycol and dielectric water mixture were used as the coolant. Experiments were conducted by varying inlet temperature, CPU workload utilization, and flow rate in the cooling loop. The results obtained from this study will pave the path for a future investigation, where a self-sensing flow control device installed in each server will be used for dynamic control of the coolant flow rate at the rack level.

2. EXPERIMENTAL SETUP

2.1 IT Equipment Used

The server used in this study is a Winterfell hybrid cooled server with Intel v2.0 motherboard, which has two CPUs and up to 16 dual in-line memory modules (DIMMs) of installable memory

as shown in Fig. 1 (Miller, 2012). Each CPU has a rated thermal design power (TDP) of 115 W. The server is retrofitted with active cold plates on two CPU units that indirectly cool the CPUs and a fan-assisted radiator that recirculates air within the server, cooling the rest of the heat-dissipating auxiliary components as shown in Fig. 2. The fans in the server circulate ambient air within the server responsible for cooling auxiliary components like DIMMs, PCH, voltage regulator device, and HDDs. The fans are four-wire pulse width modulation controlled with a 975–10,000 rpm range.

The baffle shown in Fig. 2 forces the airflow through the DIMMS present on the left-hand side as the airflow comes from the fans. In essence, the server becomes mainly liquid cooled as the coolant entering the chassis passes through the radiator and two cold plates in series before leaving to the enclosed coolant reservoir. Each cold plate has a distributed pump (circled) integrated into it as shown in Fig. 2. The flow is driven across the system utilizing these pumps. The rated power consumption and speed of the distributed pump are 3.6 W and 6250 rpm. For the current study, the chassis was populated with eight DIMM cards of 2 GB each dissipating a TDP of approximately 4 W. The physical experimental setup with its components is shown in Fig. 3.

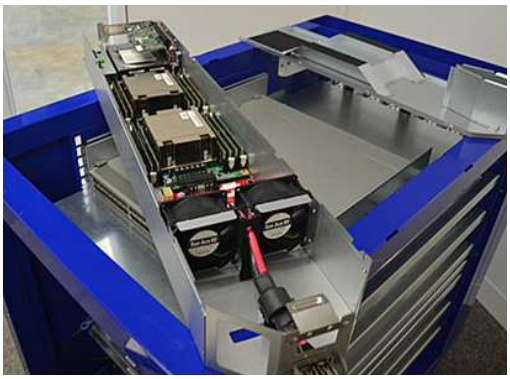

FIG. 1: Air-cooled version of the original open compute server used in the study

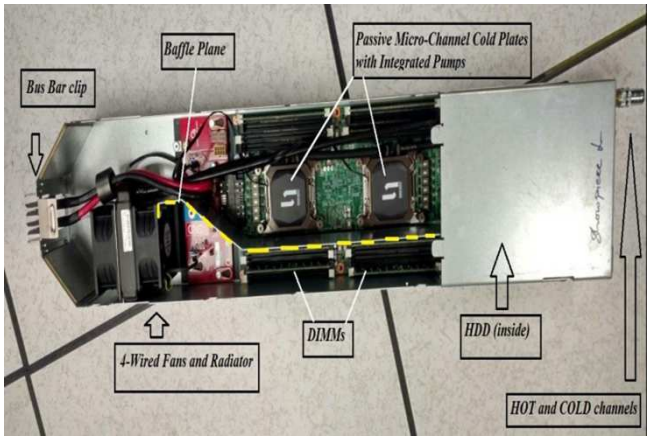

FIG. 2: Modified air-cooled server retrofitted with cold plates and a baffle (shown with yellow dashed line) to separate cold and hot air paths within the server

FIG. 3: Experimental setup showing the server cooling loop, pump control circuitry, and instrumentation for data acquisition

The primary cooling loop in the setup consisted of a coolant inlet from a custom-made reservoir toward the server as shown in the setup schematic in Fig. 4. Twenty percent inhibited ethylene glycol solution with deionized water was used as the coolant in current experiments due to ease of availability. A summary of relevant temperature-dependent thermophysical properties of the coolant used is shown in Table 2. It should be noted, however, that while ethylene glycols offer better thermal properties, their higher toxicity as compared to propylene glycol and inherent biocide behavior of propylene glycol are a few reasons they are not used widely in the industry. A liquid-to-liquid heat exchanger was used to fix and maintain the desired temperature setpoints. An assured automation DM-P series digital low-flow meter was used to measure the flow rate of the coolant, which was calibrated as per the field use with an accuracy of $\pm$ 1%. A commercially available thermal interface material was used to achieve a good thermal contact ($k = 11$ W/mK) between the CPU heat spreader and cold plate base.

2.2 Methodology

The distributed pumps were controlled using a DC power supply to maintain two extreme cases of minimum and maximum attainable flow rates based on the pump operating voltage. Table 3 provides a summary of various testing parameters that were used in the present investigation and

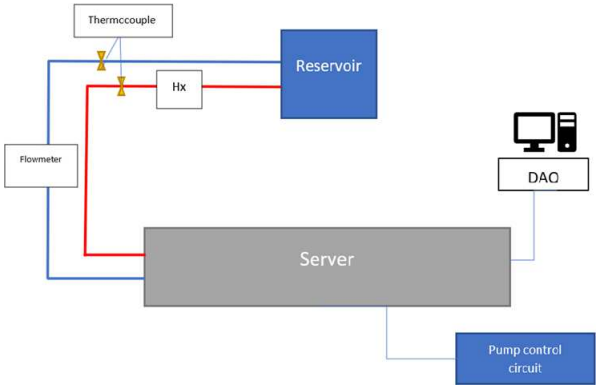

FIG. 4: Schematic of the experimental setup

TABLE 2: 20% Ethylene glycol thermophysical properties at the experimental temperature setpoints

Property	Value at 25°C	Value at 35°C	Value at 45°C
Specific heat (kJ/kgK)	3.82	3.84	3.86
Thermal conductivity (W/mK)	0.505	0.514	0.523
Density (kg/m^3)	1032.32	1028.32	1024.59
Viscosity (Pa·s)	1.41×10^{-3}	1.17×10^{-3}	9.73×10^{-4}

TABLE 3: Experimental parameters

S. No.	CPU+MEM utilization	Inlet temperature	Flow rate
1	10%	25°C	0.2 lpm
2	50%	35°C	0.6 lpm
3	100%	45°C	—

their respective values. K-type thermocouples were used to continuously monitor the inlet and outlet temperatures of the coolant using an Agilent Bench link datalogger. A WattsUp power meter was used to continuously monitor the power consumption change in the server during the experimental run, and the pump power was directly obtained from the DC power supply. To stress the server to the desired workload, an open-source stress tool for windows, Prime95, was used. It tests the computer for stability issues by stressing the CPU to its maximum limit by running Lucas-Lehmer iteration indefinitely and terminates a stress test only when it encounters an error and then informs the user that the system may be unstable. Prime95's stress test feature can be configured to test various system components by changing the fast Fourier transform size. CoreTemp was used for monitoring real-time CPU core temperature when the server utilization varies. CoreTemp is completely motherboard independent and the temperature readings are a good estimate of the temperature values as the software is collecting the data directly from a digital thermal sensor (DTS), located near the hottest part of each core. This sensor does not rely on an external circuit located to report temperature; all core temperature values are stored in the processor. Typical maximum case temperatures corresponding to 90°C DTS temperature have a value of 75°C–78°C.

PCH and DIMM temperatures were monitored using two open-source sensor software. A total of 18 experimental runs were conducted by varying CPU utilization and flow rates for the three inlet temperatures chosen. After attaining the desired temperature in the cooling loop, the CPUs were stressed and data logging was started to record core temperatures once the coolant temperature in the loop reached a steady state. The temperatures were recorded for 45 minutes and the maximum value of core, PCH, and DIMM temperatures were reported for each run.

2.3 Error Analysis

Each sensor in the study was calibrated using standard calibrating equipment and procedures. The K-type thermocouples used to measure the fluid temperature were calibrated using a Fluke 7109A portable calibration bath between a temperature of 0°C and 100°C using a two-point calibration method. Table 4 shows the error calculation quantified from the calibration process for pressure sensors and thermistors. It was observed that the pressure sensors were very precise

TABLE 4: Summary of error analysis for the sensors used in the experiments

Reference temperature (°C)	Measured temperature (°C)	% Error
10	9.8	2
90	89.3	0.8
Reference flow rate at Coriolis (lpm)	**Measured flow rate at sensor (lpm)**	**Error %**
0.2	0.207	3.5
0.5	0.519	3.8
0.7	0.721	3
1	1.03	3

after factory calibration and did not need additional calibration. A two-point calibration method was used for the thermocouples by calculating the error in the temperature reading. The calibration equation obtained was directly used as input in the DAQ software as gain and offset values. To calibrate the ultrasonic flow sensors, a Coriolis mass flow meter was used by placing the flow sensor in the same closed loop along with the Coriolis flow meter. A Coriolis mass flow meter was used for calibration as its readings are universal since it measures the mass flow rate rather than just the flow rate with a typical accuracy for liquids between 0.05% and 0.1%. The results for calibration and error analysis for both the flow meter and the thermocouples used are shown in Table 4.

3. RESULTS

The postprocessing and data analysis in the current research work included reporting core temperatures, critical component temperatures, and calculating the cooling power consumption as a percentage of total IT power consumed by the server. Figure 5 shows the variation of maximum core temperatures with varying inlet temperatures between 25°C and 45°C and at three different CPU utilizations of 10%, 50%, and 100%. The maximum core temperature was reported at 45°C inlet temperature, 100% workload, and the minimum flow rate was 79°C, which is well below the maximum junction temperature specified by the manufacturer for the processor. It was also noted that the core temperatures for 50% CPU utilization at 0.2 lpm and 100% CPU utilization at 0.6 lpm attained similar values for all inlet temperatures, as can be seen in Fig. 5. A similar case was observed for the cooling percentage curve as a function of the total IT power where the curve almost overlaps the aforementioned parameters. The given data is sufficient to conclude that even the minimum attainable flow rate should be sufficient to maintain safe junction temperatures. Another important observation made from Fig. 5 is that the core temperature at any utilization varies by 2°C–4°C only, although the lpm is increased by three times. This further bolsters the claim that a target flow delivery based on server workload can yield energy savings with a slight increment in core temperatures. For low CPU workloads, where the core temperatures may not be that high, the reliability deterioration may be negligible due to an increase in the CPU core temperature.

Figure 6 shows the variation of total cooling power as a percentage of IT power for varying flow rates, CPU utilization, and inlet temperatures. The cooling percentage (sum of pump power

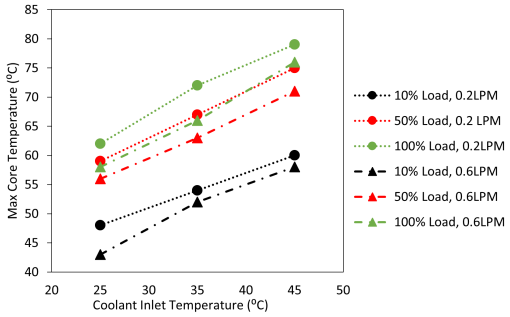

FIG. 5: Maximum core temperatures at various CPU utilization and coolant inlet temperatures

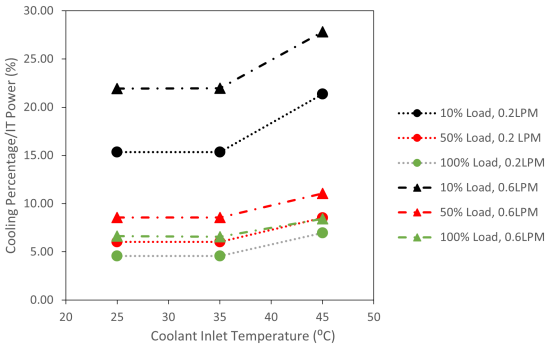

FIG. 6: Variation of the ratio of cooling power percentage and percentage of IT power for different inlet temperatures

and fan power) is defined as the percentage of power consumed for cooling the server when compared to total server power consumption. It can be seen that only 7% of total power is consumed by the fans in the extreme case of maximum inlet temperature and CPU utilization. A notable conclusion drawn from Fig. 6 is that the cooling percentage stays constant between 25°C and 35°C for all CPU utilization and both flow rates. This implies that there is no change in the fan RPM even at 35°C. Due to a lack of data points, it can't be substantially stated at what temperature the exact hike in cooling power occurs. The constant part of the graph in Fig. 6 also implies that it would be safe to operate the servers at higher coolant inlet temperatures without reaching critical temperatures even at minimum flow rates.

The data for variation in PCH and DIMM temperatures were obtained using two open-source sensor data software available for Windows OS as shown in Figs. 7 and 8. The maximum DIMM temperature obtained was around 62°C and the maximum PCH temperature reported was 69°C at minimum flow rates and maximum CPU utilization. All other values of component temperatures lay between these maximum points, represented by the shaded area on both graphs, so were not reported for convenience. The maximum difference in temperatures for the extreme cases presented in both figures is 5°C in the case of PCH and almost 8°C for DIMMs. The values of these maximum temperatures were well below the specified limits of safe and reliable operation. These results also bolster the fact that isolating the server from the ambient environment will not have a detrimental effect on the cooling efficiency. It is also evident from Figs. 7 and 8 that the trends followed by DIMMs and PCH temperatures are similar to those of the cooling percentage and maximum core temperatures, respectively.

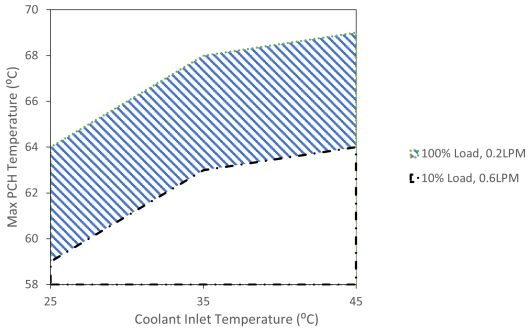

FIG. 7: Maximum PCH temperatures at various inlet temperatures and CPU utilization

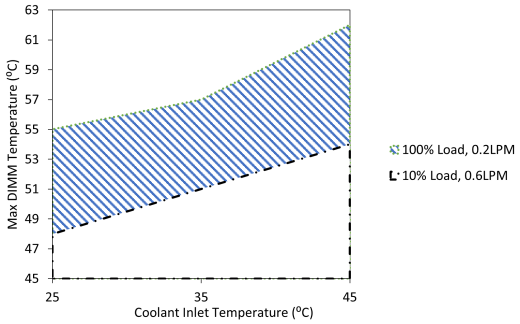

FIG. 8: Maximum DIMM temperatures at different inlet temperatures

Figure 9 describes the variation of cooling percentage with the percentage change in percentage IT power at 45°C coolant inlet temperature. For 0.6 lpm the cooling percentage drops by almost 17% and approximately 12% for 0.2 lpm as the CPU utilization changes from idle to 100% workload. The difference in cooling percentage at minimum CPU utilization is found to be approximately 7%, which reduces to 2% at maximum CPU utilization. The trend obtained shows that as the CPU utilization and inlet temperature increase, the cooling power savings reduces to almost 5%–7% for both flow rates. It can thus be concluded that there exists a trade-off value between the cooling power savings and equipment reliability due to higher component temperatures.

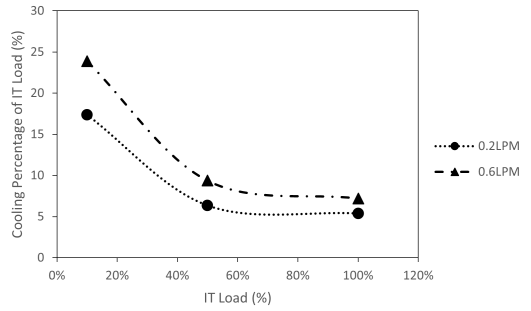

FIG. 9: Cooling percentage vs. total IT power percentage for different flow rates at a coolant inlet temperature of 45°C

The results obtained from the above experiments, as mentioned earlier, are the first step toward a novel concept of dynamic cooling or target flow delivery at the server level. This will be accomplished by using an actively controlled flow control device (FCD) integrated at the outlet of each server in a rack as shown in Fig. 10 (Shahi et al., 2022). This study details the design and development process of the FCD and characterizes the flow characteristics of the FCD. This device works on a temperature- and pressure-based control strategy. The device itself has an integrated v-cut ball valve attached to a shaft and servo motor assembly. The servo motor rotates the ball valve to desired angles based on the variation in outlet temperature of the warm coolant at the server. A server working at higher utilization will have a higher outlet temperature, which will trigger the servo to rotate the ball valve, thus allowing increased coolant flow rate to the server. Similarly, the coolant flow rate to servers working at a lower utilization will be reduced. The pressure drop change across the rack due to increased flow resistance caused by valve rotation gets compensated by varying the pump rpm. This overall optimization of pump power gives yields a lower pump power consumption per rack. For a data center with hundreds of such racks, the cost savings due to this energy optimization strategy will be significant. The experimental setup of a pseudo-rack arrangement integrated with 3D printed FCDs is being developed by the authors to report pumping power savings for a single rack. The data obtained from these experiments will be used to derive an approximate value of savings that can be achieved at the data center level. A pumping power savings of 64% has also been reported using the abovementioned control strategy computationally (Kasukurthy, 2019; Kasukurthy et al., 2021).

4. CONCLUSION

The main objectives of the present investigation were to test the concept of dynamic cooling by varying the flow rates of a server that is isolated from the ambient environment. CPU utilization, coolant inlet temperatures, and flow rates were varied, and core temperatures were constantly

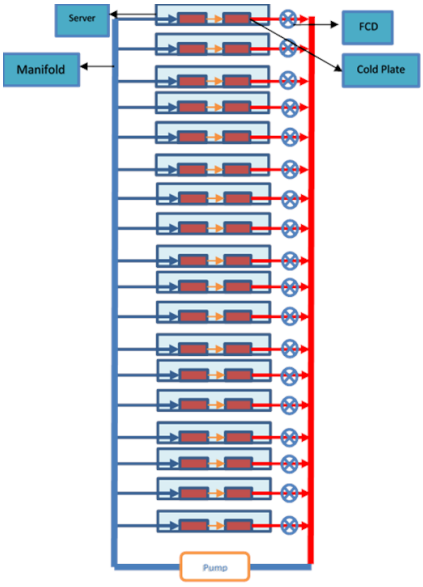

FIG. 10: Operational strategy for rack level dynamic cooling using flow control device

examined. The effect of an increase in inlet coolant temperature and recirculating inside air of the server was studied by monitoring the PCH and DIMM temperatures. The cooling efficiency by varying the pumping power at the component level was studied to substantiate the fact that controlling the flow at the server level based on the server workload utilization can yield energy savings. Both PCH and DIMM temperatures were successfully maintained well below the critical temperature limits. The maximum temperature obtained at PCH in the extreme case of bare minimum flow rate and 100% CPU utilization was 69°C. The hot circulating air inside the server even at high inlet temperatures kept the DIMMs temperature to a maximum of 61°C. No change in cooling power consumption was observed until 35°C, after which the cooling power increases. Higher-resolution experimental readings between 35°C and 45°C will give a better insight into what exact value of temperature this change happens. The cooling power savings reduce with an increase in inlet temperature, as expected, and become constant for higher IT loads.

The proposed concept of dynamic cooling of servers that are isolated from the surrounding environment can offer significant advantages in terms of cooling and improved reliability. The data center room will still require marginal cooling due to heat radiated from the server chassis, which can be cooled using outside air throughout the day with standard filtration techniques. The biggest issue of the impact of gaseous and particulate contamination related to free air-cooled data centers (Saini et al., 2022) on IT equipment can be minimized using the proposed cooling technique. This is the first time the concept of dynamic cooling at the rack level is being introduced to further improve energy savings in liquid-cooled data centers.

REFERENCES

Almoli, A., Thompson, A., Kapur, N., Summers, J., Thompson, H., and Hannah, G., Computational Fluid Dynamic Investigation of Liquid Rack Cooling in Data Centres, *Appl. Energy*, vol. **89**, no. 1, pp. 150–155, 2012.

ASHRAE T.C. 9.9, Liquid Cooling Guidelines for Datacom Equipment Centers, 2nd ed., ASHRAE, Atlanta, GA, 2014.

ASHRAE T.C. 9.9, Thermal Guidelines for Liquid Cooled Data Processing Environments, American Society of Heating, Refrigerating and Air-Conditioning Engineers, Atlanta, GA, 2011.

ASHRAE T.C. 9.9, Equipment Thermal Guidelines for Data Processing Environments, American Society of Heating, Refrigerating and Air-Conditioning Engineers, 5th ed., Atlanta, GA, 2021.

Bansode, P.V., Shah, J.M., Gupta, G., Agonafer, D., Patel, H., Roe, D., and Tufty, R., Measurement of the Thermal Performance of a Custom-Build Single-Phase Immersion Cooled Server at Various High and Low Temperatures for Prolonged Time, *ASME J. Electron. Packag*., vol. **142**, no. 1, Article ID 011010, 2019.

Chu, R.C., Hwang, U.P., and Simons, R.R., Conduction Cooling for an LSI Package: A One-Dimensional Approach, *IBM J. Res. Dev.*, vol. **26**, no. 1, pp. 45–54, 1982.

Coles, H. and Greenberg, S., Direct Liquid Cooling for Electronic Equipment, Earnest Orlando Lawrence Berkely National Laboratory, Berkley, CA, LBNL Rep. LBNL 6641 E, March 2014.

Copeland, D., Review of Low Profile Cold Plate Technology for High Density Servers, accessed October 2021, from https://www.electronics-cooling.com/2005/05/review-of-low-profile-cold-plate-technology-for-high-density-servers/, 2005.

David, M.P., et al., Cooling the Cloud: Energy-Efficient Warm-Water Cooling of Servers, *Electronics Cooling*, accessed October 2021, from https://www.electronics-cooling.com/2013/12/cooling-cloud-energy-efficient-warm-water-cooling-servers/, 2013.

Ellsworth, M.J., Jr. and Iyengar, M.K., Energy Efficiency Analyses and Comparison of Air- and Water-Cooled High-Performance Servers, *Proc. of the ASME 2009 InterPACK Conf. Collocated with the ASME 2009 Summer Heat Transfer Conference and the ASME 2009 3rd International Conference on Energy Sustainability*, San Francisco, CA, vol. **2**, pp. 907–914, 2009.

Fernandes, J.E., Minimizing Power Consumption at Module, Server and Rack-Levels within a Data Center through Design and Energy-Efficient Operation of Dynamic Cooling Solutions, PhD, The University of Texas at Arlington, 2015.

Jennings, C., Data Center Energy Consumption, accessed November 15, 2021, from http://www.hainc.com/blog/entry/data-center-energy-consumption/, 2015.

Kasukurthy, R., Design and Optimization of Energy Conserving Solutions in Data Center Application, PhD, The University of Texas at Arlington, 2019.

Kasukurthy, R., Rachakonda, A., and Agonafer, D., Design and Optimization of Control Strategy to Reduce Pumping Power in Dynamic Liquid Cooling, *ASME J. Electron. Packag.*, vol. **143**, no. 3, Article ID 031001, 2021.

Khalili, S., Mohsenian, G., Desu, A., Ghose, K., and Sammakia, B., Airflow Management Using Active Air Dampers in Presence of a Dynamic Workload in Data Centers, *Proc. of 35th Semiconductor Thermal Measurement, Modeling and Management Symposium (SEMI-THERM)*, San Jose, CA, pp. 101–110, 2019.

Niazmand, A., Murthy, P., Saini, S., Shahi, P., Bansode, P., and Agonafer, D., Numerical Analysis of Oil Immersion Cooling of a Server Using Mineral Oil and Al_2O_3 Nanofluid, *Proc. of the ASME 2020 International Technical Conference and Exhibition on Packaging and Integration of Electronic and Photonic Microsystems*, Virtual, Online, October 27–29, 2020.

Open Compute Wants to Make Biodegradable Servers, Online, *Data Center Knowledge*, accessed on November 2021, from https://www.datacenterknowledge.com/archives/2012/11/08/open-compute-eyes-biodegradable-servers, 2012.

Patterson, M.K., Liquid Cooling Guidelines, accessed October 2021, from http://delivery.acm.org/10.1145/2160000, 2011.

Rubenstein, B.A., Zeighami, R., Lankston, R., and Peterson, E., Hybrid Cooled Data Center Using above Ambient Liquid Cooling, *Proc. of 12th IEEE Intersociety Conf. on Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm)*, Las Vegas, NV, pp. 1–10, 2010.

Sahini, M., Experimental and Computational Study of Multi-Level Cooling Systems at Elevated Coolant Temperatures in Data Centers, PhD, The University of Texas at Arlington, 2017.

Sahini, M., et al., Rack-Level Study of Hybrid Cooled Servers Using Warm Water Cooling for Distributed vs. Centralized Pumping Systems, *Proc. of 33rd Thermal Measurement, Modeling & Management Symposium (SEMI-THERM)*, San Jose, CA, pp. 155–162, 2017.

Saini, S., Adsul, K.K., Shahi, P., Niazmand, A., Bansode, P., and Agonafer, D., CFD Modeling of the Distribution of Airborne Particulate Contaminants inside Data Center Hardware, *Proc. of the ASME 2020 International Technical Conf. and Exhibition on Packaging and Integration of Electronic and Photonic Microsystems*, Virtual, Online, October 27–29, 2020.

Saini, S., Shah, J.M., Shahi, P., Bansode, P., Agonafer, D., Singh, P., Schmidt, R., and Kaler, M., Effects of Gaseous and Particulate Contaminants on Information Technology Equipment Reliability—A Review, *ASME. J. Electron. Packag*., vol. **144**, no. 3, p. 030801, 2022.

Saini, S., Shahi, P., Bansode, P., Siddarth, A., and Agonafer, D., CFD Investigation of Dispersion of Airborne Particulate Contaminants in a Raised Floor Data Center, *Proc. of 36th Semiconductor Thermal Measurement, Modeling & Management Symposium (SEMI-THERM)*, San Jose, CA, pp. 39–47, 2020.

Shah, J.M., Anand, R., Saini, S., Cyriac, R., Agonafer, D., Singh, P., and Kaler, M., Development of a Technique to Measure Deliquescent Relative Humidity of Particulate Contaminants and Determination

of the Operating Relative Humidity of a Data Center, *Proc. of the ASME 2019 Int. Technical Conf. and Exhibition on Packaging and Integration of Electronic and Photonic Microsystems*, Anaheim, CA, 2019.

Shah, J.M., Bhatt, C., Rachamreddy, P., Dandamudi, R., Saini, S., and Agonafer, D., Computational Form Factor Study of a 3rd Generation Open Compute Server for Single-Phase Immersion Cooling, *Proc. of the ASME 2019 Int. Technical Conf. and Exhibition on Packaging and Integration of Electronic and Photonic Microsystems*, Anaheim, CA, 2019.

Shahi, P., Agarwal, S., Saini, S., Niazmand, A., Bansode, P., and Agonafer, D., CFD Analysis on Liquid Cooled Cold Plate Using Copper Nanoparticles, *Proc. of the ASME 2020 Int. Technical Conf. and Exhibition on Packaging and Integration of Electronic and Photonic Microsystems*, Virtual, Online, October 27–29, 2020.

Shahi, P., Deshmukh, A.P., Hurnekar, H.Y., Saini, S., Bansode, P., Kasukurthy, R., and Agonafer, D., Design, Development, and Characterization of a Flow Control Device for Dynamic Cooling of Liquid-Cooled Servers, *ASME. J. Electron. Packag.*, vol. **144**, no. 4, p. 041008, 2022.

Xu, H., Feng, C., and Li, B., Temperature Aware Workload Management in Geo Distributed Data Centers, *IEEE Trans. Parallel Distributed Syst.*, vol. **26**, no. 6, pp. 1743–1753, 2014.

Zimmermann, S., Meijer, I., Tiwari, M.K., Parades, S., Michel, B., and Poulikakos, D., Aquasar: A Hot Water-Cooled Data Center with Direct Energy Reuse, *Energy*, vol. **43**, no. 1, pp. 237–245, 2012.