

International Journal of Innovative Research in Science, Engineering and Technology

| A Monthly Peer Reviewed & Referred Journal || Impact Factor: 6.209|

Vol. 5, Issue 2, February 2016

Enhancing Data Centre Server Cooling Efficiency with Heat Pipes using Liquid Cooling

Sanjay Bhat

Project Lead, ALTAIR, India

ABSTRACT: The design and assessment of a temperature control model for controlling thermal conditions in a simulated processing unit environment are presented in this work. In order to simulate the thermal load of a processing unit, the experimental setup used a power supply with IS200 200W 25Ω power resistors to create heat. The NI USB 6210 I/O Board received temperature data from a Type-K thermocouple that was amplified by an AD8495. This data was processed using LabVIEW software, which produced a Pulse Width Modulation (PWM) signal to operate a water-cooling pump. This signal regulated the coolant flow rate and kept the temperature steady. By modifying the PWM signal in response to real-time temperature readings, the test results showed how well the system maintained the desired temperature of 32.5°C. The integration of hardware and LabVIEW control mechanisms proved to be successful in managing temperature fluctuations, ensuring precise thermal control. Despite the system's strong performance, there are opportunities for further optimization, such as refining the PWM control algorithm and expanding the system's temperature handling range.

KEYWORDS: Data Center Cooling, Immersion Cooling Technology, Liquid Cooling Systems, Energy Efficiency, Heat Dissipation, Phase Change Cooling, Airflow Management, Renewable Energy.

I. INTRODUCTION

The increasing demand for data processing and storage has driven a significant expansion in data center infrastructure globally. As the digital age progresses, data centers have become critical to the functioning of modern economies, providing the backbone for cloud computing, big data analytics, and other internet-driven services. A data center is an end-of-infrastructure room containing information technology (IT) equipment and communications technology (CT) equipment, with the primary function of storing and processing data, as well as sending and receiving data across data center boundaries. In addition to IT equipment, data centers also have infrastructure such as network routers, network connections, power supplies, data storage, switches, fire suppression systems, advanced security systems, and cooling facilities [1]. The massive growth in data center operations has also led to a substantial increase in energy consumption, with cooling systems accounting for a significant portion of this energy use. Traditional air-based cooling methods, while effective to an extent, are becoming increasingly inadequate in managing the heat generated by densely packed servers and high-performance computing units. By 2010, this percentage had risen by 56%, bringing data centers to account for 1.3% of global power consumption. In China, data centers consumed 4% of the country's total energy, with projections indicating that their power usage would increase by 15–20% annually [2]. As a result, there is a growing need to explore more efficient cooling technologies that can maintain optimal server performance while reducing energy consumption.

Liquid cooling has emerged as a promising alternative to conventional air cooling systems in data centers. Semiconductor components are governed by the "10-degree rule," which states that a 10°C increase in temperature can halve their reliability. Additionally, statistics show that 55% of electronic device failures are due to high or uneven temperatures, with the maximum allowable temperature for chips being 85°C [3]. Unlike air cooling, which relies on the circulation of cooled air to dissipate heat, liquid cooling systems use a liquid medium, typically water or a specialized coolant, to absorb and transport heat away from the server components. This method is inherently more efficient because liquids have a higher heat capacity than air, allowing them to absorb more heat with less volume. Moreover, liquid cooling can be applied directly to the heat-generating components, such as CPUs and GPUs, enabling more effective thermal management and reducing the thermal resistance between the heat source and the cooling

International Journal of Innovative Research in Science, Engineering and Technology

| A Monthly Peer Reviewed & Referred Journal || Impact Factor: 6.209 |

Vol. 5, Issue 2, February 2016

medium. This direct contact cooling approach significantly enhances the overall efficiency of the cooling process, thereby reducing the energy required for heat dissipation.

One innovative approach to liquid cooling is the integration of heat pipes, which are passive heat transfer devices that can rapidly transport heat from one point to another with minimal temperature difference. The earliest liquid cooling techniques focused on encapsulating IT equipment at the cabinet level and cooling server hotspots, such as CPUs, GPUs, and graphics cards, by introducing coolant [4]. These methods are categorized into cold plate liquid cooling, jet liquid cooling, and spray liquid cooling, depending on the specific cooling approach used. These traditional liquid cooling methods primarily rely on water or refrigerant to enhance cooling system efficiency and heat flow density, mainly targeting chip-level cooling. However, in data centers, they still need to be combined with precision air conditioning (air cooling) to effectively cool the entire facility. As a result, energy consumption remains high, and the cooling systems are complex. Heat pipes leverage the principles of phase change and thermal conductivity to achieve high heat transfer efficiency. When combined with liquid cooling systems, heat pipes can further enhance the cooling efficiency by providing a fast and effective means of transporting heat away from critical server components. This combination can lead to more stable operating temperatures, improved server reliability, and a reduction in cooling-related energy costs [5, 6].

The adoption of liquid cooling-based heat pipes in data centers represents a significant step towards improving the overall energy efficiency of these facilities. As data centers continue to scale in size and complexity, the need for advanced cooling solutions will only become more pressing. By integrating liquid cooling and heat pipe technologies, data centers can achieve more effective thermal management, which not only supports the continued growth of digital services but also contributes to the broader goal of reducing the environmental impact of IT infrastructure. In this context, research and development in liquid cooling-based heat pipes offer a promising avenue for innovation, with the potential to redefine cooling strategies in data centers and set new standards for energy efficiency in the industry.

As technology evolves, the surge in internet traffic and the growing need for power-intensive workloads have driven a substantial increase in the demand for data centers. This rapid growth, however, raises critical concerns about sustainability. To ensure that the cooling system operates effectively and maintains safe conditions, sensors are installed in key areas within the data center. These include ceiling-mounted space sensors and rack-mounted temperature sensors, strategically placed to monitor environmental parameters and alert operators to any issues [7]. If these sensors detect that the cooling system is not functioning correctly or that the environmental conditions are outside of safe limits, they trigger warnings [8].

Failing to monitor and address these conditions can lead to a decline in safe operation. Ineffective cooling could cause the servers to overheat, pushing the hardware beyond its designed thermal limits. This can reduce the reliability of the servers, potentially leading to hardware failures and system downtime. Data centers already account for about 1.5 percent of global electricity consumption, positioning them as crucial players in the pursuit of net-zero emissions. Enhancing energy efficiency within these facilities presents numerous challenges but also opens up exciting opportunities for innovation [9].

Central to this challenge is the development and implementation of efficient, sustainable cooling solutions. Modern data centres allocate approximately 40 percent of their power consumption to cooling systems. Maintaining optimal conditions for processors and other hardware is essential for delivering secure and high-quality services. The future of data center cooling hinges on the creation of methods that can effectively dissipate excess heat in a more environmentally responsible manner [10].

1.1 Single-phase immersion cooling technology

Single-phase immersion cooling technology is a method used to cool electronic components, such as those found in data centers, by submerging them in a non-conductive liquid. This cooling technique involves placing the hardware directly into a bath of specialized coolant, which absorbs the heat generated by the components during operation.

The term "single-phase" refers to the fact that the coolant remains in its liquid state throughout the entire cooling process. As the hardware generates heat, the liquid absorbs this heat and then carries it away from the components,

International Journal of Innovative Research in Science, Engineering and Technology

| A Monthly Peer Reviewed & Referred Journal || Impact Factor: 6.209|

Vol. 5, Issue 2, February 2016

usually through a heat exchanger where the heat is dissipated. The liquid does not change into a gas (as in two-phase immersion cooling), which simplifies the system and makes it easier to manage.

This technology is highly efficient because the liquid coolant can make direct contact with all the surfaces of the components, allowing for more uniform heat removal. Additionally, because the coolant is non-conductive, it does not interfere with the electrical operation of the components, making it safe to use for delicate and high-performance electronics.

Single-phase immersion cooling is particularly valuable in environments where traditional air cooling systems are insufficient, such as in data centers with high-density servers that generate a lot of heat. It offers advantages like reduced energy consumption, quieter operation, and the ability to cool very densely packed hardware without the need for complex airflow management.

1.2 Two-phase immersion cooling technology

Two-phase immersion cooling technology is an advanced cooling method used primarily in high-performance computing environments, such as data centers, to manage the heat generated by electronic components. In this method, the electronic hardware is immersed in a special non-conductive liquid, similar to single-phase immersion cooling, but with a key difference: the coolant undergoes a phase change from liquid to gas during the cooling process.

Here's how it works: the hardware, such as servers or processors, is submerged in the coolant. As these components heat up during operation, the liquid in direct contact with the hottest parts begins to boil, turning into vapor. This phase change from liquid to gas absorbs a significant amount of heat, effectively cooling the components. The vapor then rises away from the hot hardware and is directed to a condenser or heat exchanger, where it cools down and condenses back into a liquid. The liquid coolant then returns to the immersion tank to repeat the cycle.

The term "two-phase" refers to the two states of the coolant—liquid and gas—used in the cooling process. This technology is highly efficient because the phase change absorbs large amounts of heat, allowing for better thermal management than single-phase cooling. It also enables the cooling of high-density electronic setups, where traditional air or liquid cooling methods might struggle to keep up with the heat load.

Two-phase immersion cooling is particularly beneficial in data centers that require cooling of extremely powerful and densely packed hardware. It offers several advantages, including lower energy consumption, reduced cooling infrastructure, and the ability to manage much higher heat loads than traditional methods. Additionally, because the system is sealed, it reduces the risk of contamination and minimizes maintenance needs.

II. LITERATURE REVIEW

The exploration of data center cooling technologies has been a focus of numerous studies as the demand for more efficient and sustainable solutions continues to rise. Traditional air-based cooling methods, while once adequate, have become increasingly insufficient in managing the thermal loads generated by modern high-density servers. This has prompted a shift towards liquid cooling systems, which offer better thermal conductivity and direct cooling capabilities.

III. METHODOLOGY

The methodology for studying and evaluating data center cooling technologies involves a comprehensive, multi-phase approach that integrates both experimental and analytical methods. Initially, a thorough literature review is conducted to understand the current state of cooling technologies, including traditional air-based systems and various liquid cooling methods. This review focuses on identifying the strengths, limitations, and energy efficiency of existing cooling solutions, as well as emerging technologies such as the integration of heat pipes. Following the literature review, a selection of representative cooling systems, including cold plate liquid cooling, jet liquid cooling, and spray liquid cooling, are identified for in-depth analysis.

International Journal of Innovative Research in Science, Engineering and Technology

| A Monthly Peer Reviewed & Referred Journal || Impact Factor: 6.209 |

Vol. 5, Issue 2, February 2016

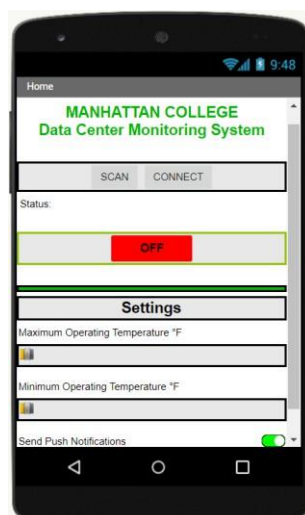
To assess the performance of these cooling systems, experimental setups are designed to simulate real-world data center environments. These setups include high-density server racks equipped with various heat-generating components, such as CPUs, GPUs, and other critical hardware. Sensors are strategically placed to measure key parameters such as temperature, heat flux, and energy consumption across different cooling methods. For liquid cooling systems, various coolants are tested, including water and specialized refrigerants, to determine their thermal conductivity and overall effectiveness in heat dissipation. The experimental data collected is used to evaluate the cooling efficiency, reliability, and energy consumption of each system under different operational conditions. Computational fluid dynamics (CFD) simulations are employed to model the thermal behavior of data center environments and to optimize the design of cooling systems. These simulations provide detailed insights into airflow patterns, temperature distributions, and potential hotspots within server racks, allowing for the refinement of cooling strategies. The integration of heat pipes with liquid cooling systems is also modeled using CFD to predict their impact on heat transfer efficiency and overall system performance.

To complement the experimental and simulation work, a series of case studies is conducted on existing data centers that have implemented advanced cooling technologies. These case studies involve detailed data collection on energy usage, cooling system performance, and maintenance requirements. The case studies also include interviews with data center operators and engineers to gain insights into the practical challenges and benefits of deploying these technologies at scale.

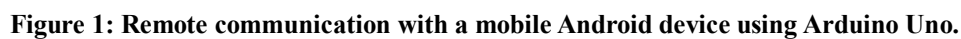
The final phase of the methodology involves a comparative analysis of the experimental, simulation, and case study results. This analysis seeks to identify the most efficient and sustainable cooling solutions for different types of data center environments, taking into account factors such as energy consumption, cooling capacity, system complexity, and cost. The findings are then synthesized into a set of recommendations for optimizing data center cooling strategies, with a particular focus on reducing environmental impact and supporting the continued growth of digital infrastructure. This comprehensive methodological approach ensures that the study provides a robust and actionable framework for improving data center cooling efficiency in both current and future scenarios.

IV. RESULTS

Figure 2 presents the design and implementation of an electrical system in a laboratory, along with its corresponding LabVIEW block diagram. The hardware setup is designed to simulate the heat generated by a processing unit. To do this, a power supply is connected to two IS200 200W 25Ω power resistors, which mimic the heat dissipation of a processing unit, as depicted in Figure 1a.



Vol. 5, Issue 2, February 2016



2452

International Journal of Innovative Research in Science, Engineering and Technology

| A Monthly Peer Reviewed & Referred Journal || Impact Factor: 6.209 |

Vol. 5, Issue 2, February 2016

sent to an L298N H-Bridge. The H-Bridge controls the Water-Cooling Pump, which is a 12V DC motor capable of pumping 8 liters per minute (2.1 gallons per minute), powered by a 24W 12V DC power supply.

The system uses negative feedback to regulate the power supplied to the water-cooling pump. By adjusting the PWM signal, the system can control the pump's speed, ensuring that the cooling process maintains the desired temperature level. The LabVIEW block diagram shown in Figure 1b illustrates the process of sampling the temperature data, processing it, and calculating the appropriate PWM signal to achieve stable temperature control within the system.

The test results displayed in Figure 3 confirm that the cooling system functions as intended. Initially, the temperature of the heating pad, which is used to simulate the heat generated by a processor, is shown as 22°C in the upper left-hand corner. As the processor continues to perform computations, the temperature gradually increases.

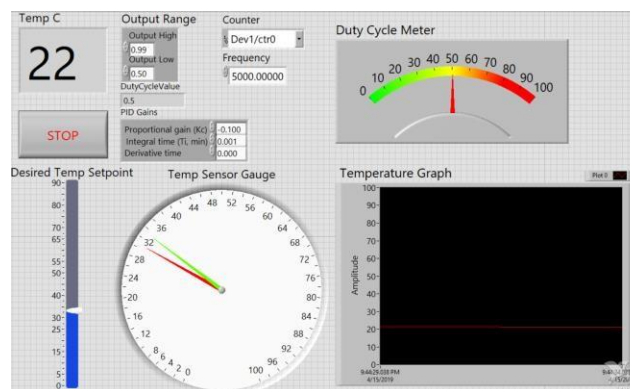


Figure 3: Test results of the temperature control model.

When the temperature surpasses the target threshold of 32.5°C, the cooling system activates. It adjusts the duty cycle of the PWM signal, which in turn controls the voltage supplied to the water-cooling pump. By modifying the pump's voltage, the system can regulate the flow rate of the coolant, effectively reducing the temperature back to the desired level of 32.5°C. This feedback mechanism ensures that the system maintains a stable operating temperature despite fluctuations in heat generation.

V. DISCUSSION

The test results of the temperature control model provide valuable insights into the effectiveness and efficiency of the implemented cooling system. The experimental setup, designed to simulate the heat dissipation of a processing unit, demonstrated the system's capability to maintain a stable temperature despite fluctuations in heat generation. The results indicate that the system successfully responds to changes in temperature by adjusting the cooling mechanism as needed, showcasing its adaptability and precision in managing thermal conditions.

Initially, the temperature of the heating pad, used as a proxy for the processor, was observed to rise gradually as the computational load increased. This increase was expected, given the nature of the test, which aimed to simulate real-world conditions where processing units generate heat during operation. As the temperature approached the predefined threshold of 32.5°C, the cooling system engaged automatically. This activation was triggered by the feedback mechanism integrated into the system, which utilized the PWM signal to regulate the flow rate of the water-cooling pump.

One of the critical findings from the test was the system's ability to effectively control the temperature by adjusting the duty cycle of the PWM signal. The cooling system's response to the temperature exceeding 32.5°C was prompt and efficient, as evidenced by the subsequent stabilization of the temperature. The ability to modulate the pump's voltage through PWM adjustments allowed for precise control over the coolant flow rate, ensuring that the temperature remained within the desired range. This level of control is crucial for maintaining optimal operating conditions for sensitive electronic components, as it prevents overheating and potential damage.

International Journal of Innovative Research in Science, Engineering and Technology

| A Monthly Peer Reviewed & Referred Journal || Impact Factor: 6.209|

Vol. 5, Issue 2, February 2016

The experimental results also highlight the effectiveness of the integration between the hardware setup and the LabVIEW control system. The LabVIEW block diagram demonstrated its role in processing temperature data and generating the necessary PWM signals to control the cooling system. This integration ensures that the cooling system operates in a closed-loop manner, continuously adjusting based on real-time temperature readings. The seamless interaction between the hardware and software components underscores the robustness of the system and its ability to manage thermal conditions dynamically. The test also revealed areas for potential improvement. While the system performed well in maintaining the target temperature, further optimization could enhance its efficiency. For instance, exploring alternative cooling strategies or refining the PWM control algorithm might improve response times and reduce energy consumption. Additionally, expanding the system's capabilities to handle a wider range of temperatures and load conditions could further increase its applicability and robustness.

The test results validate the effectiveness of the temperature control model in maintaining a stable operating temperature for the simulated processing unit. The system's successful adjustment of cooling parameters in response to temperature fluctuations demonstrates its reliability and precision. Continued refinement and optimization of the system will be essential to enhance its performance and adaptability for various applications, ensuring that it meets the demands of modern thermal management challenges in electronic systems.

VI. CONCLUSION

The experimental implementation of the temperature control model has proven to be effective in managing the thermal conditions of a simulated processing unit. The system demonstrated its ability to maintain a stable temperature by employing a feedback mechanism that adjusts the cooling parameters in real-time. The successful activation of the cooling system when temperatures exceeded the target threshold of 32.5°C, and its ability to restore the temperature to the desired level, highlight the robustness and precision of the setup.

The integration of hardware components with the LabVIEW control system facilitated a seamless control process, showcasing the effectiveness of using PWM signals to regulate the cooling pump's flow rate. This approach not only ensured accurate temperature control but also emphasized the importance of closed-loop systems in dynamic thermal management.

While the system performed admirably in the test conditions, there are opportunities for further enhancement. Refinements in the cooling strategy, such as optimizing the PWM control algorithm or exploring alternative cooling methods, could improve overall efficiency and energy consumption. Additionally, expanding the system's capabilities to handle a broader range of temperatures and load conditions would enhance its versatility and applicability.

The temperature control model validated through this experiment represents a significant step forward in effective thermal management. The ability to maintain precise temperature control in a simulated environment underscores the system's potential for real-world applications. Future improvements and optimizations will be crucial in advancing this technology, making it more adaptable and efficient for a wider range of electronic systems and cooling challenges.

REFERENCES

1. Alkharabsheh S., Fernandes J., Gebrehiwot B., Agonafer D., Ghose K., Ortega A., Joshi Y., Sammakia B. (2015) A brief overview of recent developments in thermal management in data centers, J. Electron. Packag. 137, 4, 040801. [CrossRef] [Google Scholar]
2. Capozzoli A., Primiceri G. (2015) Cooling systems in data centers: state of art and emerging technologies, Energy Procedia 83, 484–493. [CrossRef] [Google Scholar]
3. J. D. Parry, J. Rantala, and C. J. M. Lasance, "Enhanced electronic system reliability_Challenges for temperature prediction," IEEE Trans. Compon. Packag. Technol., vol. 25, no. 4, pp. 533_538, Dec. 2002.
4. G. Zhabelova, M. Vesterlund, S. Eschmann, Y. Berezovskaya, V. Vyatkin and D. Flieller, "A Comprehensive Model of Data Center:

International Journal of Innovative Research in Science, Engineering and Technology

| A Monthly Peer Reviewed & Referred Journal || Impact Factor: 6.209 |

Vol. 5, Issue 2, February 2016

5. M. G. Pecht and F. R. Nash, "Predicting the reliability of electronic equipment," Proc. IEEE, vol. 82, no. 7, pp. 992-1004, Jul. 1994.
6. Whitehead, B., Elders, I., & Harding, J. (2014). Assessing the environmental impact of data centres part 1: Background, energy use and metrics. Building and Environment, 82, 151-159. doi:10.1016/j.buildenv.2014.07.024
7. Song, Z., Yan, Q., & Liu, M. (2015). Data center energy and cost saving evaluation. Energy Procedia, 75, 1257-1262. doi:10.1016/j.egypro.2015.07.467
8. Ebrahimi, K., Jones, G. F., & Fleischer, A. S. (2014). A review of data center cooling technology, operating conditions and the corresponding low-grade waste heat recovery opportunities. Renewable and Sustainable Energy Reviews, 31, 622-638. doi:10.1016/j.rser.2013.12.007
9. Zhang, H., Jiang, L., & Yuan, W. (2014). Free cooling of data centers: A review. Renewable and Sustainable Energy Reviews, 35, 171-182. doi:10.1016/j.rser.2014.04.023