THE ART OF TWO-PHASE DATA CENTER COOLING

Precision Science Driving Data Center Thermal Efficiency

Dr. Issam Mudawar

Betty Ruth and Milton B. Hollander Family Professor of Mechanical Engineering
Director, Purdue University Boiling & Two-Phase Flow Laboratory (PU-BTPFL)
Director, Purdue University International Electronics Cooling Alliance (PU-IECA)
Principal Investigator, NASA's Flow Boiling & Two-Phase Condensation Experiment (FBCE) for the International Space Station (ISS)

Introductions by Dr. Richard W. Bonner III, CTO of Accelsius

ACCELSIUS™
contact@accelsius.com
www.accelsius.com

INTRODUCTION

The discussion surrounding data center liquid cooling revolves around two key technology comparisons. The first is between immersion and direct-to-chip cooling, which is relatively straightforward to comprehend. The second involves comparing single-phase to two-phase heat removal solutions, which can be more challenging to grasp fully.

Most understand the fundamental fact: single-phase removes heat while maintaining its liquid state. In contrast, two-phase removes heat by changing from a dielectric liquid to a vapor state, and back to a liquid state in a closed cooling loop which we know enables a significantly greater amount of heat removal. But why is this true? And how does it work?

For those with a penchant for scientific exploration, we have invited Dr. Issam Mudawar, an expert in thermodynamics, to educate on "The Art of Two-Phase," a term he used to contrast Accelsius' NeuCool direct-to-chip platform to single-phase cooling. The NeuCool platform is designed to deliver true precision cooling with every component meticulously engineered to provide the most efficient and reliable method for cooling servers. Unlike single-phase cooling, two-phase does not rely on the brute force of increased flow rates to remove higher heat loads.

In addition, single-phase cooling often uses water (or Glycol mix) in it's cooling loop which has the potential to leak causing material damage to the server and rack-level IT gear costing customers millions of dollars. Two-phase, direct-to-chip cooling (like NeuCool) uses an A1 rated, safe, non-corrosive dielectric optimized for phase change cooling and with no potential for "leak" damage.

Dr. Mudawar is a renowned thermal scientist and founder of Purdue's Boiling and Two-Phase Flow Laboratory. He has played a pioneering role on both national and international stages in the development of thermal solutions for supercomputers, servers, laptops, chip testing, hybrid vehicle power electronics, and x-ray medical devices and systems.

CONTENTS

THE MERITS OF TWO-PHASE COOLING

In this first part of the series, Dr. Mudawar provides performance summary of single-phase versus two-phase cooling. The compelling two-phase cooling performance leads to an extended discussion on two of the key phase change processes involved in two-phase cooling: boiling and condensation. The description is written in a way that is relatively easy to understand and yet fundamentally rigorous.

Breakthroughs in many of today's cutting-edge technologies are becoming increasingly dependent upon the ability to safely dissipate enormous amounts of heat from very small areas. Examples include computer electronics, data centers, nuclear fission reactors, nuclear fusion blankets, x-ray medical devices, hybrid and electric vehicle power electronics, avionics, laser and microwave defense electronics, and advanced radar.

Cooling effectiveness is dictated by both cooling configuration and type of coolant used. As shown in Fig. 1, force-convection involving supply of coolant via a pump, compressor, or fan is superior to natural convection, which driven entirely by the buoyancy resulting from density gradients in the coolant, but the most demanding cooling configurations are managed with boiling.

Boiling, or *liquid-to-vapor phase change*, occurs when a liquid coolant used to cool a heat dissipating surface exceeds the liquid's boiling point. Without the phase change, pure liquid cooling relies entirely on sensible heat capacity of the liquid to extract the heat, which is manifest by an increase in the temperature difference between the surface and liquid that is proportional to the surface heat flux (heat transfer rate per unit surface area). For a surface that is dissipating a large amount of heat, this would culminate in a commensurate large increase in the surface temperature.

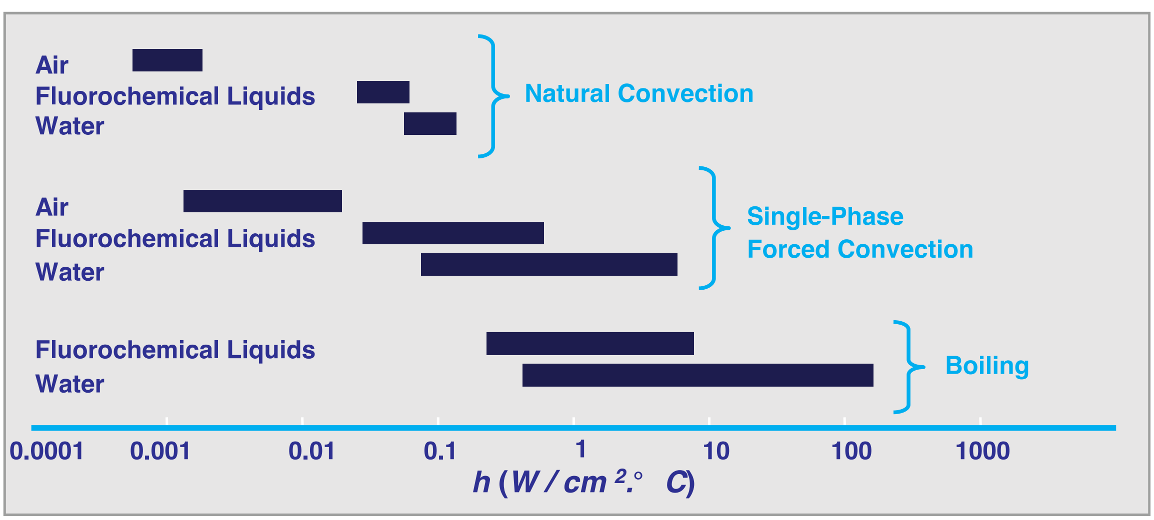

FIG. 1 Comparison of cooling effectiveness based on cooling configuration and coolant type

The primary advantage of boiling is reliance on the coolant's both sensible and latent heat, which allows for dissipation of a broad range of heat rates corresponding to only modest increases in surface temperature. Fig. 2 illustrates this advantage for an electronic device that dissipates a heat flux between 0 and 50 W/cm2 but whose surface temperature, because of design and safety reasons, cannot exceed 85°C. Acceptable cooling performance is identified by the region in the figure highlighted in orange color. The figure shows the response of surface temperature to different cooling options. With air cooling, the device temperature would exceed the 85°C limit at a miniscule heat flux, while with pure liquid cooling, the heat dissipation is greatly ameliorated, albeit failing to maintain acceptable surface temperature at the maximum heat flux of 50 W/cm2.

On the other hand, by using a coolant having a boiling point below the maximum temperature, the coolant will begin to boil, which greatly enhances the cooling performance relative to pure liquid cooling in two major ways: (1) causing the surface to incur only modest changes in the surface temperature corresponding to fluctuations in the wall heat flux (which is important for temperature sensitive devices as well as reduces the likelihood of delamination and debonding resulting from the temperature cycling), and (2) maintaining surface temperatures significantly below those with pure liquid cooling, 70°C versus 525°C, at the maximum heat flux. These constitute the primary heat transfer merits of boiling.

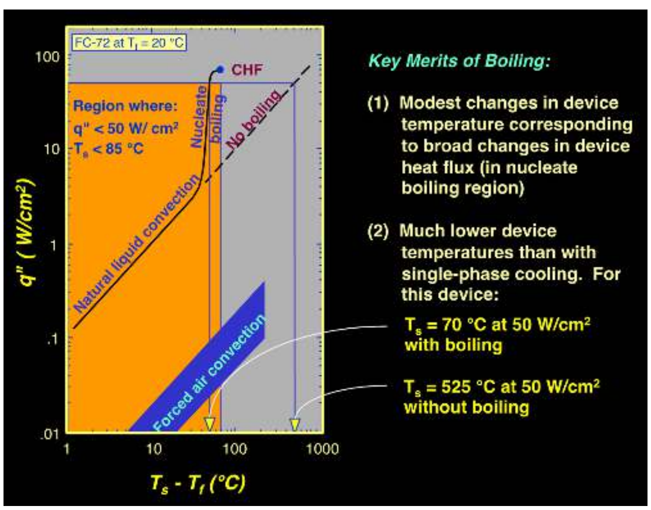

FIG. 2 Heat transfer merits of boiling compared to both pure liquid cooling and air cooling.

PHYSICS OF BOILING

On a fundamental basis, these merits of boiling are rooted in the nucleation, growth, and departure of vapor bubbles from the surface. These interrelated processes draw bulk liquid toward the surface to continuously replace liquid that has been converted into vapor. These take place at very high frequency, accounting for the enormous heat transfer benefits realized in what is called the nucleate boiling regime. Fig. 3 illustrates these processes during pool boiling on a surface of an electrically heated wire.

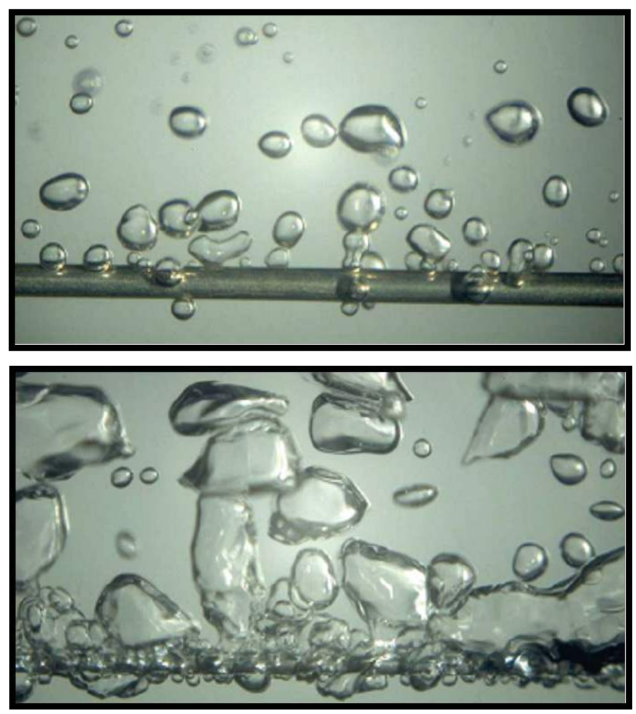

FIG. 3 Images depicting nucleate boiling regime (left) and critical heat flux (CHF) limit (right) for boiling along the surface of an electrically heated wire.

The vapor-liquid exchange process that is responsible for much of the cooling effectiveness within this regime requires uninterrupted liquid access to the device's surface. Higher heat flux fluxes are dissipated by production of more vapor bubbles per unit surface area. Increased bubble coverage of the surface will eventually lead to significant vapor coalescence and begins restricting liquid access to the surface. Once the vapor-liquid exchange process is interrupted, the power dissipated in the device itself will no longer be rejected and the surface temperature would begin to escalate uncontrollably. This condition, which is also depicted in Fig. 3, is the upper-most heat flux limit for the nucleate boiling regime and is termed critical heat flux (CHF). Therefore, effective design of the cooling system consists or maintain operation within the nucleate boiling region, but safely below the CHF limit.

"Higher heat flux fluxes are dissipated by production of more vapor bubbles per unit surface area."

PHYSICS OF CONDENSATION

Finally, it is important to emphasize that any phase change cooling scheme wherein a liquid coolant is converted to vapor will require a condenser to return the vapor produced back to liquid state as part of a closed system. Several types of condensers are popular for electronic cooling applications, including crossflow air-cooled, plate type, and shell-and-tube. Aside from internal design differences, these condenser types also differ in the type of external coolant used to extract the heat from the system to condense the vapor back to liquid state. Within the inner flow passages of the condenser, as shown in Fig. 8, flow structure typically consists of a thin film that sheathes the flow passage inner walls, shear-driven by a central, fast-moving vapor core. Condensation along the film's interface gradually decreases flow rate of the vapor while increasing the flow rate of the liquid. In a practical closed loop cooling system, it is often desired to achieve full condensation of the vapor before returning the liquid to cool the heat dissipating devices.

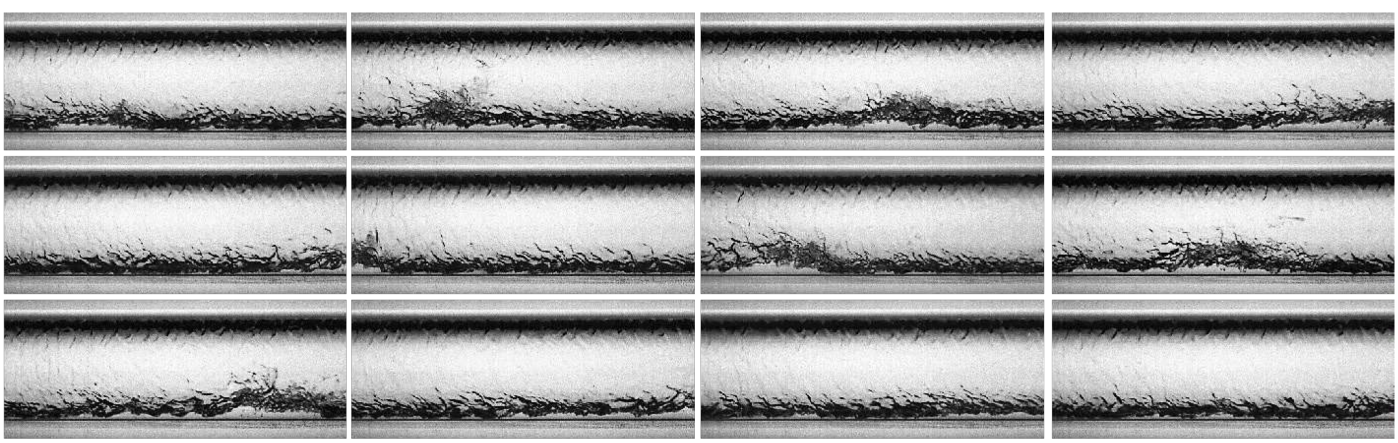

FIG. 8 Returing vapor to liquid state along horizontal flow passages of a condenser

TWO-PHASE COOLING STRUCTURES

The phase change process in two-phase cooling lends itself to some unique opportunities for cooling enhancement. Here Dr. Mudawar discusses some of these structures that are often exploited by thermal engineers. The first two technologies discussed, pool boiling thermosyphons and pumpless cooling loops, use the large density differences between liquid and vapor to passively circulate flow using nothing but the applied thermal energy. Pumped flow loops capable of removing the highest heat fluxes are then covered. Finally, a summary of the different channel flow boiling modules used in practice to enhance the surface area for boiling is provided.

The past four decades have witnessed unprecedented interest in application of boiling to dissipate high fluxes while maintaining safe device temperatures using a variety of cooling schemes or configurations. Among those are *passive schemes*, relaying on buoyancy to achieve coolant circulation, and *active schemes* relaying on a pump to achieve the circulation.

Pool boiling thermosyphon is the simplest of the passive schemes wherein the heat dissipating electronic devices are submerged in a pool of dielectric liquid in a closed vessel. As shown in Fig. 4, the heat is dissipated by boiling the liquid directly upon the surface of the device. The generated vapor rises by buoyancy to the top section of the vessel, where it is returned to liquid state by condensing on the surfaces of an air-cooled or water-cooled condenser; the condensed liquid drips back to the liquid pool. Fig. 5 shows increased heat dissipation on the device surface increases the surface area incurring boiling.

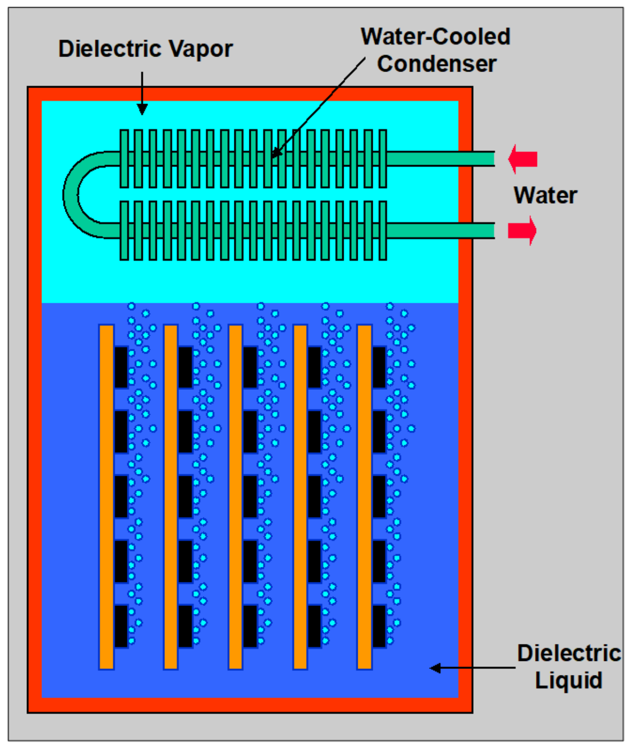

FIG. 4 Schematic of a pool boiling thermosyphon used to cool a multitude of circuit boards

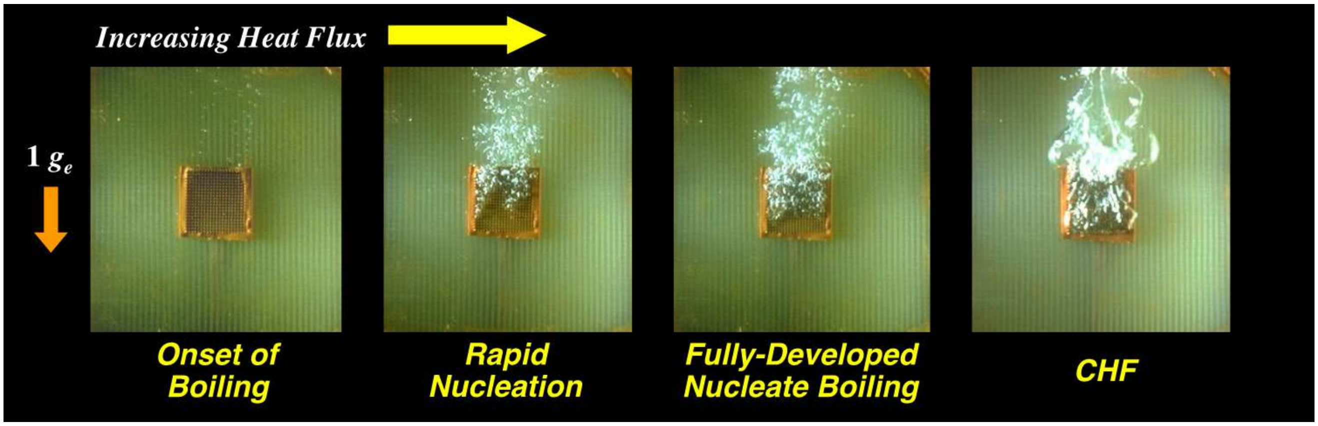

FIG. 5 Evolution of the pool boiling process with increasing surface heat dissipation.

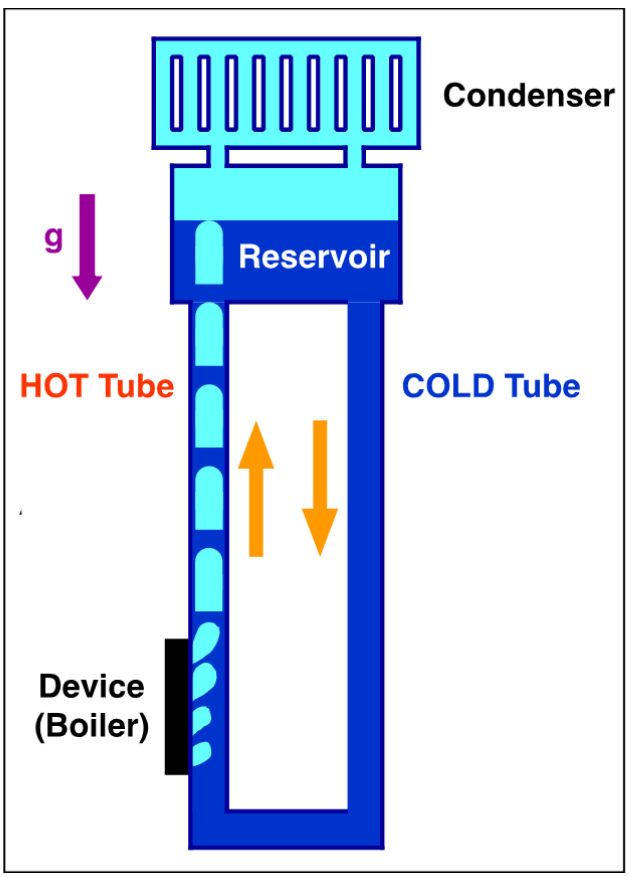

FIG. 6 Passive parallel tube cooling scheme

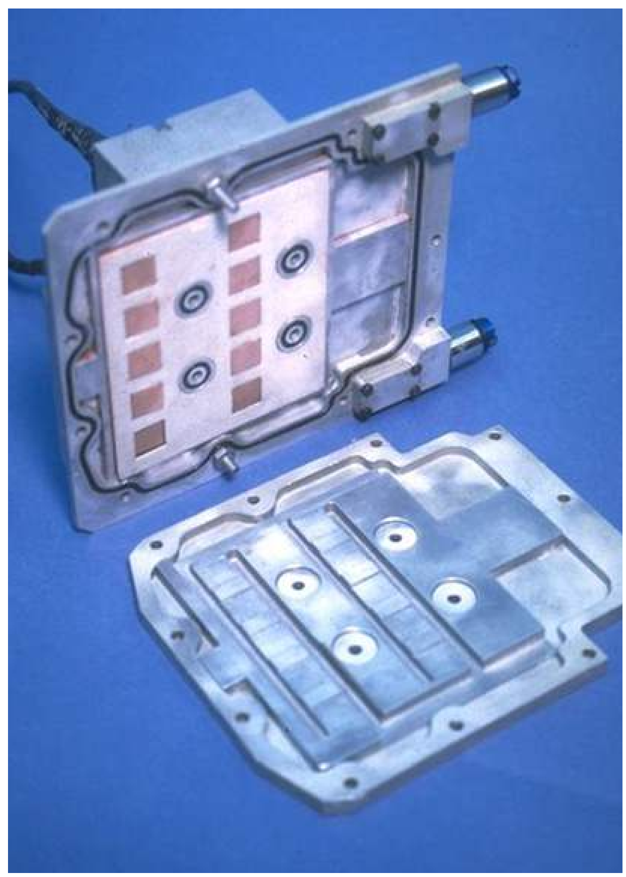

FIG. 7 Channel flow boiling module – cold plate – enclosing a multitude of heat dissipating devices. Coolant inlet and exit are managed using two side fluid couplers.

The *pumpless cooling loop* is another passive cooling scheme which consists of two vertical parallel tubes that are connected at the bottom, and atop to a liquid reservoir fitted with a condenser. The heat-dissipating devices are connected to one tube (hot tube). As the vapor is generated by boiling, average fluid density and hydrostatic pressure in the hot tube are greatly reduced compared to those in the other tube (cold tube). This imbalance triggers passive fluid circulation in the loop, with liquid from the cold tube flowing downwards towards the hot tube, and liquid-vapor mixture in the hot tube rising upwards to the liquid reservoir, where the vapor is re-condensed to liquid state.

There are additional cooling schemes that are used to tackle very high heat dissipation scenarios. Most popular among which are those utilizing mini/micro-channels, jet impingement, and sprays, as well as hybrid combinations thereof.

When passive cooling systems fail to tackle high device heat fluxes, a pumped flow loop is required. The simplest of the pumped schemes is channel flow boiling wherein the heat dissipating devices are mounted in a variety of series and parallel arrangements inside a closed cooling module – cold plate – as shown in Fig. 7, with the coolant flowing parallel to the device surfaces. The primary benefit of this cooling scheme is the heat transfer enhancement realized with fluid motion. However, there is also great flexibility in controlling coolant temperature in the form of either saturated boiling, where the coolant is supplied at its boiling point, or subcooled boiling, where the coolant temperature is lower than the boiling point. Subcooling provides the advantages of greatly enhancing the CHF limit and ensuring that the vapor partially or fully recondenses in the bulk liquid; this would simplify the cooling loop from a two-phase liquid-vapor configuration to a single-phase liquid configuration.

There are additional cooling schemes that are used to tackle very high heat dissipation scenarios. Most popular among which are those utilizing mini/micro-channels, jet impingement, and sprays, as well as hybrid combinations thereof. There are also numerous methods for ameliorating cooling performance in all mentioned cooling schemes, including modifying the surface with micro texturing or use of fins, treating the surface with carbon nanotubes (CNTs), or using a nanofluid containing nano-sized conducting particles.

TWO-PHASE COOLING IN DATA CENTERS

In the final part of the Art of Two-Phase Cooling Series, Dr. Mudawar connects the dots between the previously covered fundamental concepts and the data center industry. In the data center world, two-phase cooling using dielectric coolants is the most logical approach to cool next generation servers while maintaining the reliability that the data center industry has come to expect.

Two-phase cooling with microchannel coolers is discussed in greater detail. This variation of cold plate technology is considered the "best-of-breed" regarding cooling performance. Dr. Bonner, Accelsius' CTO, closes the series by explaining how Accelsius NeuCool platform makes it easy for data center clients to take advantage of the state-of-the art technologies discussed by Dr. Mudawar in this special series.

Recent technological advances in high-performance computing and artificial intelligence are paralleled by rapid increases in the amount of heat dissipation. Data centers are comprised of a multitude of electronic cabinets, each housing a large number of cooling modules – cold plates - for cooling of circuit boards. The dense architecture of data centers, coupled with adoption of increasingly more advanced electronic devices, has led to unprecedented increases in the amount of heat that needs to be removed at the device, cold plate, cabinet/rack, and system levels.

While the heat removal in current data centers is tackled via air cooling or single-phase liquid cooling, the increased heat dissipation in very high performance data centers is beginning to exceed the capabilities of present cooling methods, which points to accelerating needs to adopt more effective two-phase cooling methods. With such scenarios, a highly effective cooling system hierarchy is one wherein cold plates, using a dielectric coolant and modified internally to take advantage of micro-channel cooling, as depicted in Fig.7, would be connected in a variety of series and parallel arrangements to a two-phase cooling loop managing an entire cabinet. The two-phase loop would be thermally interfaced to those tackling the heat removal in the other data center cabinets to enable ultimate heat rejection from the entire system to the ambient.

MICRO-CHANNEL COOLING

Of the different pumped cooling schemes, two-phase micro-channel cooling, which is adopted within the cold plate shown in Fig. 7, has gained unprecedented popularity among device and system manufacturers because of several attributes, including (i) simple construction, (ii) compactness, (iii) low coolant inventory, and (iv) very high heat transfer performance. The cooling effectiveness of micro-channels is largely the result of small coolant passage diameter, typically less than one millimeter. Optimum design of micro-channels for high heat flux application is achieved using a variety of two-phase numerical and computational models. Key outcomes of such models are (i) ability to dissipate the heat while maintaining acceptable device temperatures, (ii) maintaining operation safely below the critical heat flux (CHF), and (iii) achieving required cooling without escalating pressure drop and therefore pumping power.

DATA CENTER COOLING USING ACCELSIUS' NEUCOOL PLATFORM

By Dr. Richard Bonner

Accelsius conveniently offers the patented (pending) benefits of two-phase cooling through the various products offered under its NeuCool platform (shown in Figure 8). Accelsius' Vaporator (seen in Figure 9) attaches directly to the high-heat-flux CPUs and GPUs found in many of today's servers. The Vaporator uses micro-channel features and precision fluid ducting to push liquid towards the hottest parts of the chips, pull the heat directly from the chip, and shuttle the generated vapor towards the exit of the cold-plate.

The Vaporator is constructed robustly to minimize the threat of leaks. However, if a leak does occur, the entire NeuCool platform features low GWP, dielectric coolants that will not harm your server in any way or damage the environment. The NeuCool platform's iPCU also features intelligent flow control and general communication to the data center infrastructure management software (DCIM). Embedded in the iPCU is a heat exchanger that uses facility water to remove heat from the di-electric coolant used to cool the server: water never goes into the server. The iPCU is also located below the server, such that facility water always remains located below your server.

If your server cooling requirements are pushing you towards something that air cannot handle, please reach out to our team at Accelsius to discuss the NeuCool platform and its two-phase cooling capability in greater detail.

FIG. 8 The NeuCool Platform

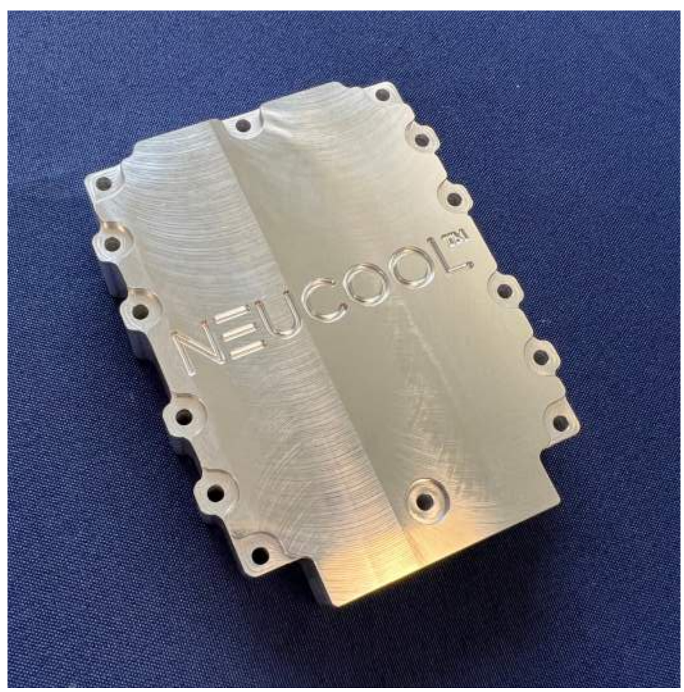

FIG. 9 The Accelsius Vaporator