Status of Liquid Cooling of Data Centres: Some Answers and Some Questions

A.J. Robinson, PhD

Nomenclature

CDU	Coolant Distribution Unit
CHF	Critical Heat Flux
DLC	Direct Liquid Cooling
ERE	Energy Reuse Efficiency
LCA	Life Cycle Analysis
PUE	Power Usage Effectiveness
SPLIC	Single Phase Liquid Immersion Cooling
TPIC	Two Phase Immersion Cooling

1. Introduction

It is common knowledge that almost all of today's data centres use air as the coolant for managing component temperatures in servers. Air is dielectric, non-toxic, free, abundant, and as a chemical has no intrinsic global warming potential. However, owing to its poor thermophysical properties, keeping components below their safe temperature limits becomes progressively more challenging as power and heat flux levels escalate, and this is particularly true for the server processors (CPUs, GPUs etc.). Pushing the limits of air cooling is complicated by the fact that there is little to no headroom for improvement. This is because turbulent convective heat transfer intensity plateaus at high velocities, so forcing more air is a situation of diminishing gains. At the same time, acoustic noise and fan power dramatically escalate. If considered purely from a cooling perspective, the only option for air cooling is to increase the height of the server, say from 1U to 2U, in order to accommodate taller heat sinks with more heat transfer surface area. This is in fact the technical solution often used, though comes with its own set of problems. In particular, increasing the server dimensions decreases the rack compute density, which is surely a step backwards. Unfortunately, this is the option that is available, and is well known, since the limits of air cooling are well established and have been for some time [1]. Simply put, some CPU heat fluxes are at or near the air-cooling threshold, so alternative means of cooling must be sought [2-4].

Thermal management of server hardware is not solely a heat transfer problem, it is an energy problem [5]. In recent years, the spotlight has been shone on data centres over their electrical energy consumption and associated environmental impact [6], and there is growing public and governmental pressure to address this problem. Regardless, data centres are a business, so operating costs, and in current times energy security and looming legislative measures, are incentives to reduce the electrical energy consumption. With global average PUEs currently levelling off at around 1.57 [7], and air cooling being largely responsible for this, it must necessarily be targeted if significant improvements in data centre efficiencies are to be realized.

Considering the above, liquid cooling of data centres seems inevitable. Owing to their considerably better thermophysical properties, using liquids as a coolant not only has the capacity to solve the heat transfer and energy problems simultaneously, but gives new headroom for future development whilst facilitating heat recovery, transport and reuse.

In this report, three different liquid cooling technologies are reviewed, Single-Phase Immersion Cooling, Two-Phase Immersion Cooling and Direct Liquid Cooling. The overarching aim of the report is to provide useful technical information and hopefully give some clarity with regard to the state-of-the-art of liquid cooling technology as it pertains to data centre cooling and energy management. To achieve this, only data disclosed in primary references is reported.

2. Immersion Cooling

2.1 Single-Phase (Liquid) Immersion - SPLIC

2.1.1. How it works

As outlined by Chen and Li [8], *Single-Phase Liquid Immersion Cooling* (SPLIC) refers to the technology whereby electronic components are cooled by direct contact with a dielectric liquid coolant. Servers are placed vertically, side-by-side, in a tank. The tank is filled with a dielectric (non-electrically conductive) liquid to a level just above the servers. The liquid is forced to and from the tank by a pump in a closed loop between the tank and the Coolant Distribution Unit (CDU). The CDU contains a liquid-liquid heat exchanger which uses pumped facility water to cool the dielectric fluid. The now heated facility water is circulated by a secondary pump which routes it to the rooftop air cooling units (e.g. dry cooler). The cooled dielectric coolant then flows vertically through the servers where the heat generated by the electronics is transferred by convection to the liquid. As the liquid flows from entrance to exit of the servers, its temperature rises due to sensible heating of the liquid. The hot liquid is then routed back to the CDU to continue the heating-cooling cycle.

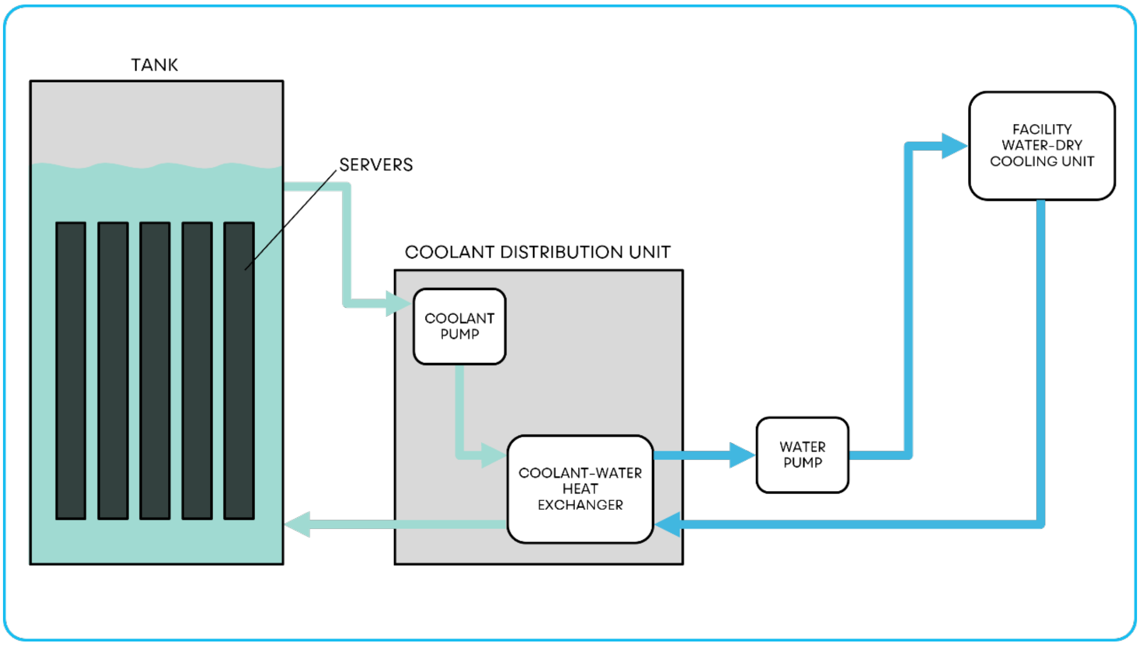

Figure 1: Schematic representation of single-phase immersion server cooling technology

2.1.2 Dielectric Liquids

A *dielectric liquid* is one which is electrically insulative (does not conduct electricity) so as to not interfere with the operation of the electronic device/system, which itself is conducting electricity to perform its required function. Thus, dielectric liquids are in direct physical contact with current-carrying components without electrically interacting with them. Air is an electrical insulator up to a certain voltage intensity, and this is one key reason why it is used as a coolant for computing systems, which are typically low voltage.

Uncontaminated dielectric liquids can achieve insulative properties about 10^7 times better than air [9]. Additionally, being a liquid, they possess better thermophysical properties for heat transfer applications compared to air. In particular, their thermal conductivity, specific heat capacity and density combine in such a way as to provide better intensity of convective heat transfer at much lower volumetric flow rates. Table 1 shows a comparison of key thermophysical properties of typical dielectric liquids (hydrocarbon mineral oil, synthetic oil, hydrofluoroether), with air and water, and can be used as a preliminary guide to assess the efficacy of fluids for heat transfer applications.

Table 1: Key thermophysical properties for heat transfer fluids

Fluid	Description	Thermal Conductivity k, (W/mK)	Specific Heat Capacity C_p, (kJ/kgK)	Density ρ, (kg/m^3)	Viscosity μ, (Pas)	Volumetric Heat Capacity ρC_p, (kJ/m^{3}K)
Air	N/A	0.026	1.0	1.2	0.000018	1.2
HFE 7100	Hydrofluoroether (HFE)	0.069	1.183	1510	0.00057	1798
STE Crystal Plus*	Hydrocarbon mineral oil	0.13	1.67	849.3	0.016	1,418
EC-100*	Synthetic oil	0.14	2.17	803.8	0.012	1,744
Water	N/A	0.6	4.18	997	0.0089	4,167

* Shah, Eiland et al. (2019) [21]; **Shah, Bhatt et al. (2019) [14]

From Table 1, a couple things immediately stand out. Dielectric liquids are notably more thermally conductive, dense and store/transport more energy per unit mass/volume than air. It is perhaps no surprize that submersing electronic devices and components in dielectric liquids is a very old technology, stemming back to the late 19th century [10], possibly earlier [11]. It was not long after the invention of computers that the identical technology was translated into to the compute industry, with the first know application by IBM in 1966 [12]. It is also clear that dielectric liquid coolants are still far from ideal when compared with water, though water cannot be used for cooling in direct-touch applications.

Key Take-Aways

- Single-phase liquid immersion as a cooling technology for electronics, including computers, is a very old technology.
- In terms of thermophysical properties that promote convective heat transfer, dielectric liquids are good, compared to air, but not great, compared to water.

2.1.3 Performance

As mentioned by Kanbur et al. [13], there still exists an extensive knowledge gap with regard to detailed system-scale performance immersed liquid cooling technology. This includes its dynamic performance under varying load and environmental conditions [14]. Despite the fact that it has been around for over a decade, most of what we know is currently from non-primary reference sources (e.g. media references), and can be considered anecdotal until rigorously verified. Regardless, some research exists which can be relied on to give a general sense of the performance attributes and limitations of this technology.

Eiland et al. [14] performed a clever experiment whereby a single Intel-based Open Compute server was considered in order to obtain unambiguous heat transfer and fluid dynamic performance data over a range of feasible coolant flow rates and inlet temperatures. The stock 2OU server was kept relatively intact, apart from a few modifications (removing fans etc.), the most notable being the hard disk drive having to be situated outside the tank, since it cannot operate submersed. Oil flow rates between 0.5-2.5 L/min were tested at inlet temperatures between 30°C-50°C. For the tests, the server was run at 75% CPU utilization with 20% memory allocation, with a total server power typically just below 225 W for both stock server air tests (at 25°C) and the white mineral oil-cooled scenario. Unsurprisingly, the liquid cooled configuration could beat or match cooling performance of the stock air-cooled server, and could do so up to about 40°C inlet temperature. Defining the Partial Power Usage Effectiveness (pPUE) as the ratio of the sum of the cooling power and IT load to the IT load (i.e. PUE not including facility infrastructure), they found that it ranged between about 1.03 and 1.17, with the lower values favouring the lowest flow rates and vice versa. Interestingly, increasing the oil temperature tended to improve the pPUE as the pumping power reduced, owing to reduced viscosity, which more than offset increases in IT power due to leakage current effects, showing an interesting interrelation between the IT and thermal-hydraulic systems in the context of thermal and energy performance.

From a thermal fluid perspective, what is noteworthy is the low flow velocities of the oil in SPLIC systems such as this, resulting in very low Reynolds numbers and thus highly laminar flow. This is detailed further in Shah, Bhatt et al. [15], who numerically investigated the flow and heat transfer in a submersed Third Generation Open Compute Server. Testing both white mineral and synthetic EC-100 oils, they investigated the influence of altering the case thickness from 1OU to 2OU for the same flow rates used by Eiland et al. [14]. Flow velocities ranged between approximately 0.0005 – 0.003 m/s, which is so low that the Reynolds numbers were between about Re~10-20. This can be considered extremely low for convective heat transfer applications. The flow rate and inlet oil temperature trends on the cooling effectiveness of this simulation study were broadly similar to those of Eiland et al. [14]. A notable contribution of this work is the estimation of the thermal resistance associated with the finned heat sink in oil cross-flow, which was within a relatively tight range of ~0.2 – 0.3 K/W, despite the different case dimensions. They key reason for this relative insensitivity is likely because laminar flow generally exhibits low sensitivity to the Reynolds number, to the extent that for fully developed laminar flow in ducts, the Nusselt number is in fact constant [16]. They also showed some minor improvements with the synthetic mineral oil EC-100 and hypothesize that Fluorochemical liquids may offer further advantage, though this may be debatable since they tend to have much lower thermal conductivity (e.g. HFE 7100, Table 1). In the context of immersed single-phase cooling, very little is available in the open literature in terms of assessing their heat transfer potential in this cooling configuration, though Cheng et al. [17] investigated HFE 7100 in an immersed PC compute system. Recently, Chhetri et al. [9] performed a very similar simulation study to Shah, Bhatt et al. [15], using the Third Generation Open Compute Server and the synthetic oil EC-110. For identical flow rates and inlet temperatures up to 40°C, the study focussed on thermal shadowing, which is where an upstream component (CPU1 in this case) heats the coolant to the extent that the downstream component (CPU2 in this case) runs notably hotter. This work showed a relatively severe thermal shadowing effect, and that for high heat fluxes could only be managed with either or both increasing the flow rate or reducing the inlet temperature, both of which having implications in terms of the energy demand of the cooling infrastructure, though this was not considered in this investigation. Regardless, this study raises a fairly critical point, which is that the design layout and orientation of the server must be considered very carefully hen implemented in SPLIC systems.

The above studies give some insight into the thermal-fluid-energy behaviour at the single-server level, which is relevant, though cannot be completely extrapolated, to the system-level, in particular with regard to the energy dynamics. One of the earlier studies that considered system-level performance was by Chi et al. [18], who performed what can be considered a like-for-like rack-level comparison between a hybrid air-cooled system with rear door water cooling, and a fully immersed rack. What was different here was that each server contained a closed volume of fluoro-organic liquid. The dielectric liquid transferred heat by natural convection to a water-cooled coldplate also embedded in the server; ostensibly a server-integrated liquid-liquid heat exchanger. In this way, the interaction with the water-cooled plate and the heated dielectric within the server cavity created buoyancy-driven flow to transfer the heat from the components, into the dielectric, and then out of the dielectric to the water. In their comparison, the hybrid-air system used facility water from the large-scale chiller whereas the liquid-cooled rack used a dry-cooler to transfer the heat directly to the outdoor environment. The key finding was that the (partial) PUE of the liquid cooled system achieved 1.14, compared to 1.48 for the hybrid system, a 34% reduction.

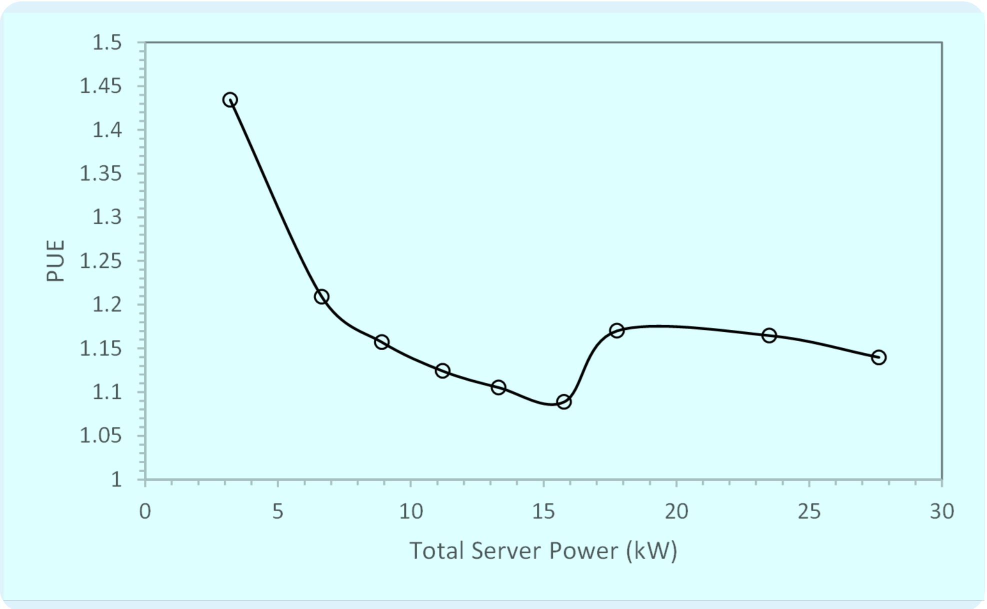

Figure 2: PUE versus total server power for commercial Single-Phase Liquid Immersion Cooling rated at 30 kW

More recently, this was considered in more detail by Kanbur et al. [13] who have provided some of the first reliable measurements of system-scale energy performance of an immersed bath SPLIC system. In this work, a commercial 42U SPLIC system with a maximum capacity of 30 kW was investigated, in a configuration resembling that of Figure 1 i.e. including a CDU and dry cooler. Supermicro GPU Superserver 1028 GQ servers, rated to 2 kW, were tested over a total power range between 3.2-27.6 kW. Data for the total server power, oil pump power, water pump power and dry cooler fan were tabulated and are reproduced in Figure 2 in terms of the PUE=(cooling power + IT power)/IT power. It is first noted that with increase in IT load between 3.8 kW to 15.8 kW, the PUE drops from 1.43 to 1.09. This a result of the dual-speed pumps operating in fixed yet low frequency mode, such that the cooling load remains about constant as the IT load increases. After an IT load of 15.8 kW, the oil pump triggers to high frequency mode as an oil temperature approaches just under 40°C., and the additional power draw causes the PUE to increase to 1.17. Further increase in IT power reduces at a marginally lower rate, and this is partially due to the water pump also switching to the high frequency mode of operation to further moderate the inlet temperature.

What is somewhat striking here is the wide variation and relatively high PUEs, in particular at the lower and higher IT power loads, considering that values typically quoted to be much lower [19]. However, this may be a result of quoting best achieved PUEs, highlighting the current need for more detailed system-level and dynamic data in the peer reviewed literature. As Kanbur et al. [13] mention, this may of course be in part due to this particular system having a secondary water cooling loop. However, the data suggests that this will only have a minor influence, in particular since oils are generally more viscous. Regardless, it is significantly better than the current global average PUE which is closer to 1.57 [7], dominantly due to the ineffectiveness of air cooling. However, this must be considered in the context that some air-cooled data centers have reporting PUEs better than 1.2 [20,21].

Key Take-Aways

- Very little reliable data exists in the open literature to critically assess server-scale and system-scale thermal-hydraulic-energy performance of SPLICs.
- Forcing liquids through the large open area of a server results in very low flow velocities, even for reasonably high flow rates (circa 1-2 L/min per server).
- 40°C appears to be an approximate upper threshold for inlet coolant temperature.
- Layout of CPUs and server orientation can influence cooling performance.
- Dynamic IT power loading may cause wide variation in PUE, which is realistically between 1.09-1.4; better than global average PUE though close on average to some of those reported for air-cooled data centres.
- High IT power loads require higher flow rates and are detrimental to PUE.

2.1.4 Some Key Open Questions

Here, some additional considerations are outlined in terms of the implementation of SPLIC systems, with particular focus on oils, which currently are the most studied and deployed fluid for this application. A similar section can be found in Section 3.4 for fluids like Hydrofluoroethers, as they are more common for two phase immersed cooling applications.

Material Compatibility: Probably the most important unanswered question regarding SPLIC systems is the material compatibility of the oils with the numerous materials which they contact. As pointed out by Shah et al. [22], any remarks regarding a lack of detrimental effects on component immersed in mineral oil are anecdotal. Simply put, there is not enough data to draw this type of conclusion, although some progress is being made [22,23]. Problematically, since full immersion results in direct contact of the oils with all exposed material in a server, and can be absorbed to sub-surface material, there is a myriad of materials to test and the tests must necessarily be taken over an extended time frame and require that multiple thermal-mechanical-electrical properties be assessed. This includes the different potential oils which themselves can degrade. It is a mammoth task, yet necessary for server equipment suppliers to not void warranties and/or regulatory compliance subsequent to immersion in oils.

Life Cycle Analysis (LCA): Hydrocarbon-based mineral oils are distilled from petroleum. There is an embedded energy cost associated with their manufacture, transport, and recycling. Any energy and CO2 savings associated with SPLIC operation must be considered in the context of an LCA. Synthetic oils, albeit a different chemical make-up and synthesis process, have been shown to have lower, yet comparable embedded energy intensity per litre as mineral oils [24]. Considering the large volumes of oil in these systems, along with inevitable recharging, there is an open question around the impact of an LCA on energy performance of SLPIC.

Flammability: Oils are combustible and there is a large volume of potential fuel in SPLICs. Although the flash points are generally much higher than the temperatures in properly cooled servers, the fire hazard raises concerns around safety and standards. It is also of note that increasing the flash point of oils typically increases the viscosity, so there is additional consideration required in this regard [25]

Other Considerations [22,25]:

- Occupational safety e.g. toxicity, handling, service life of oils
- Signal compatibility e.g. interference with I/O signal transmission

- Further environmental impacts e.g. recycling; soil and groundwater contamination
- The unexpected e.g. lifting of barcode stickers, dissolving ink labelling, debris accumulation etc. [26]

Key Take-Aways

- Material compatibility of dielectric oils and their impact on the chemical and mechanical properties of server components are justified concerns owing to lack hard data.
- It is unknown how a Life Cycle Analysis will influence the energy/CO2 dynamics of SPLICs.

2.2 Two-Phase (Boiling-Condensation) Immersion Cooling - TPIC

2.2.1. How it works

Like SPLICs discussed above, *Two-Phase Immersion Cooling* (TPIC) refers to the technology whereby electronic components are cooled by direct contact with a dielectric coolant. The key differentiators are that in TPICs, a specialized dielectric fluid that has a relatively low (typically <65°C) boiling point is used. Once again, servers are placed vertically, side-by-side, in a tank. The tank is filled with the low boiling point dielectric (non-electrically conductive) liquid to a level just above the servers and the tank is sealed. The higher heat flux components will heat to above the saturation temperature of the fluid (termed *superheat*) and this will cause vapor bubbles to form, grow, and depart from the surface; this is called *nucleate pool boiling* [27] In this way, the energy from the heated components is now converted to latent heat, which is the energy stored in the fluid during the process of converting liquid to vapor, opposed to the sensible heat transfer process of the SPLIC technology which raised the temperature of the liquid during the energy transfer-storage process. Another main difference between the two immersion technologies is the method by which the stored energy is transported from the servers. In the SPLIC system, pumped liquid continually moves the heat energy from the heated parts out of the immersion tank. In the TPIC system, when the bubbles containing the latent heat depart, they rise due to gravity-induced buoyancy until they breach the surface of the pool, releasing the vapor and absorbed energy, into the upper region of the chamber. In the upper chamber is a heat exchanger (usually bare copper pipe or similar) that contains flowing facility water, which is cold enough to condense the vapor, causing a release of the latent heat into the flowing water stream. Mechanically, gravity is again leveraged to cause the heavy condensate to flow back into the lower pool, thus closing both the fluidic and thermodynamic cycle.

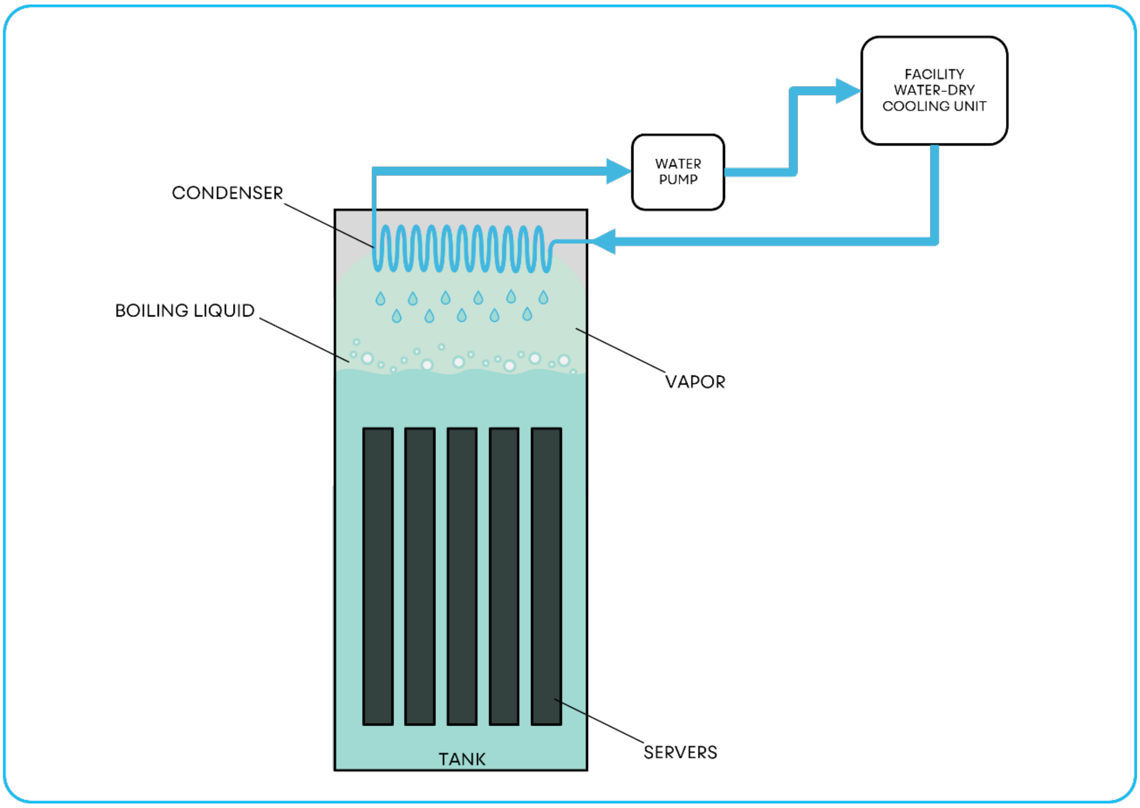

Figure 3: Schematic representation of two-phase immersion server cooling technology

2.2.2 Dielectric Liquids and Boiling

The use of boiling dielectric fluids to cool heated server electronics of course entails a different set of physical mechanisms compared with the case of single-phase convection in the SPLIC system. As such, a different set of thermophysical properties dictate the intensity of heat removal. Table 2 lists the properties which are well known to influence boiling heat transfer intensity for some of the more common dielectrics used in immersed boiling (FC 72, Novec 649, HFE 7100). Also tabulated is a Figure of Merit which can be used to roughly compare different fluids [28]

$$FOM = \frac{10^3}{\left(\frac{h_{fg}}{C_p}\right)\left(\frac{C_p\mu}{k}\right)^n\left(\frac{1}{\mu h_{fg}}\sqrt{\frac{\sigma}{g(\rho_L-\rho_V)}}\right)^{0.n}}$$

where n~1 for water on metals and n~2 for refrigerants on metal [29].

Table 2: Key thermophysical properties for two-phase heat transfer fluids

Fluid	Description	BP	k	σ	h_{fg}	ρ_L	μ	C_p	FOM	GWP
FC 72	Fluorocarbon	56	0.057	0.01	88	1680	0.00064	1.1	3.3	9,300
Novec 649	Fluorinated ketone	49	0.059	0.0108	88	1600	0.00064	1.103	3.4	1
HFE 7100	Hydrofluoroether	61	0.069	0.0136	112	1510	0.00057	1.183	4.3	297
Water	N/A	100	0.61	0.72	4.18	997	0.00089	4.18	18.3	0

BP=Boiling Point (°C), k=Thermal Conductivity (W/mK), σ= Surface Tension (N/m), h_{fg}=Latent Heat of Vaporization (kJ/kg), ρ= Density (kg/m^3), μ =Viscosity (Pas), Cp=Specific Heat Capacity (kJ/kgK), FOM=Figure of Merit, GWP= Global Warming Potential

From Table 2, it is evident that the low boiling point dielectrics are broadly similar, with HFE 7100 showing better potential as a coolant. Furthermore, these fluids would not be expected to produce nearly as intense boiling heat transfer compared with water, and this is well known. Though boiling water on server components is impractical due to its high boiling point and electrically conducting properties, this does illustrate that there are limits imposed by the nature of the fluids with regard to how intense a cooling effect they can generate, and low boiling point dielectrics are relatively poor boiling fluids in the context of the spectrum of engineering fluids used in two-phase equipment. Beyond this, boiling heat transfer has associated with it a severe failure mode, called the Critical Heat Flux (CHF), which occurs when the vapor generation rate is so high that it prevents new liquid from rewetting the surface. This creates an insulating vapor layer causing a thermal runaway condition, typically leading to burnout of electrically heated devices. Again, the fluid properties, in particular the latent heat of vaporization, play a key role in determining this failure limit. For perspective, for water at atmospheric pressure CHF~100 W/cm^2 [27], whereas for the dielectrics listed in Table 2 CHF~20 W/cm^2 [30,31], which is dangerously close to the generated local heat fluxes of some modern computer processors. A final consideration is the phenomenon of wall temperature excursions at boiling incipience. This is particularly relevant to highly wetting and low surface tension fluids, like low boiling point dielectrics, since they tend to penetrate the small pits and cavities on the surface which act as nucleation sites for boiling. As the heat flux is increased, the onset of boiling is delayed until such time as the surface temperature is high enough to instigate boiling. Often, and in particular with highly wetting fluids, the temperature even drops subsequent to the onset of nucleate boiling, as the mode of heat transfer switches from convection to liquid to more intense boiling, and this is part of a process known as *boiling hysteresis* [32]. Depending on the smoothness of the surface and heater orientation, the surface temperature required to initiate boiling of highly wetting fluids can be substantial, and the overshoot temperature quite severe (>10 K) [30,33].

The relatively poor boiling heat transfer coefficients and low CHF of low boiling point dielectrics has been recognized as a problem for enabling TPIC technology for some time [34]. In order for boiling to be viable for this class of fluids; (i) the maximum heat flux must be managed or the CHF increased, (ii) the boiling heat transfer coefficient must be enhanced, (iii) the boiling incipience temperature lowered, and (iv) boiling hysteresis eliminated. These must all be achieved with stable, reliable, and cost effective design interventions.

There is a wealth of publications in the open literature on methods and means to enhance boiling heat transfer [33,35]. In recent times, most of these are motivated by the escalating heat fluxes of electronic components and the potential use of two-phase cooling to manage their temperature within strict limits. Enhancement strategies range from micro/nano features to porous coatings, with the latter likely showing the most promise for the case at hand. A good example is illustrated in Figure 3, which shows a schematic representation of a surface-enhanced Integrated Heat Spreader (IHS) beneath which is attached a high-powered electronic component (e.g. processor die). This can be viewed potentially as a porous coating applied directly to a commercial CPU package. This type of technology can (i) raise the CHF via the enhanced porous coating whilst lowering the surface heat flux via heat spreading in the IHS, (ii) enhance the boiling heat transfer coefficient via, among other things, increased nucleation site density, (iii) reduce the boiling incipience to a few degrees of superheat, (iv) eliminate boiling hysteresis, and (v) eliminate TIM2 (Thermal Interface Material between IHS and attached heat sink) if the coating is directly applied to the IHS of the commercial processor package. However, one must be very cautious, as these are typically highly engineered surfaces and, to the best of knowledge, are still very much in the development phase as the research community continues to develop an understanding of why and how they work, or do not in some cases. This was borne out by Wu et al. [36] who investigated two phase immersion cooling of a server with both an attached plate with an engineered porous copper coating as well as the stock aluminium heat sink. Surprisingly, the stock heat sink performed better, which is somewhat inconsistent with the existing body of knowledge on the subject, and highlights potential pitfalls associated with deploying underdeveloped technology.

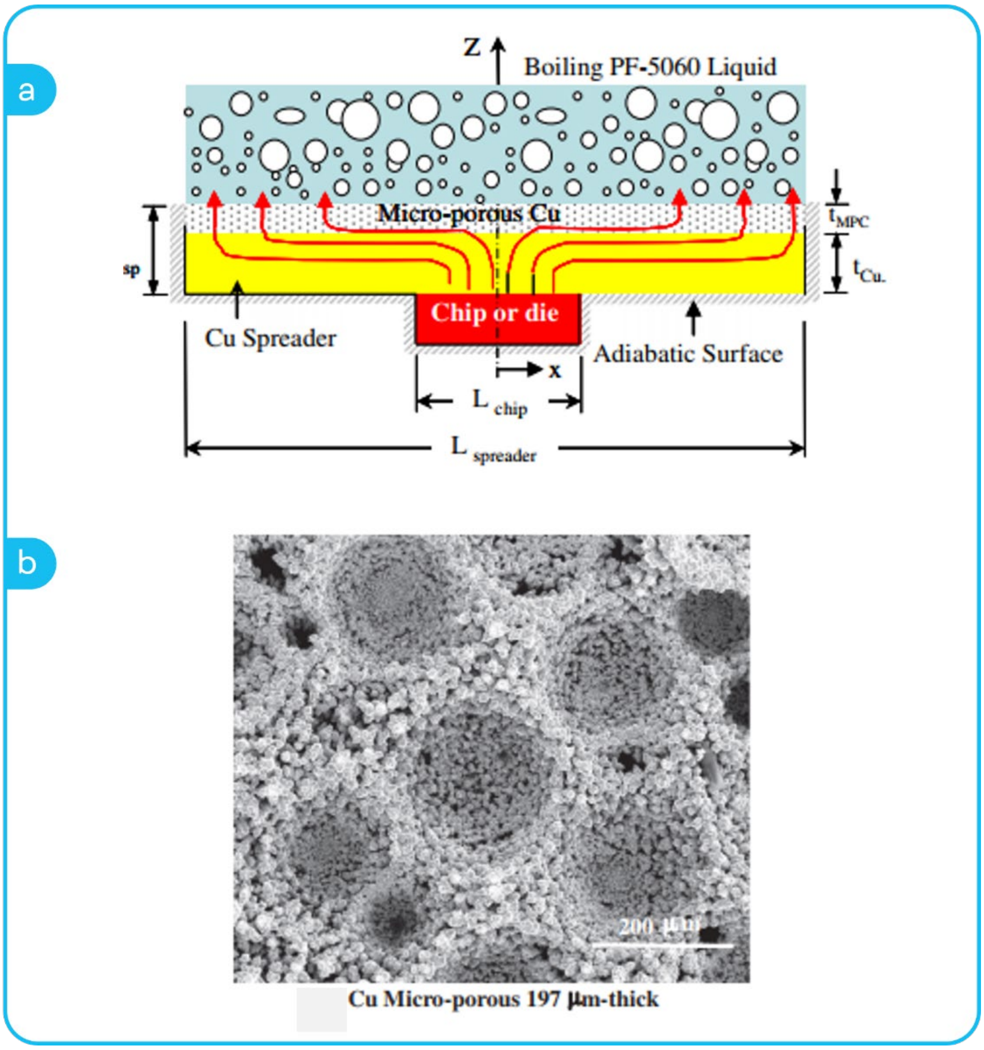

Figure 3: (a) Diagram of cooling of a heated die equipped with an IHS coated with porous material to enhance nucleate pool boiling heat transfer, (b) Enhanced copper micro-porous surface coating [37]

What has not been mentioned yet is the boiling point. The typical range of feasible boiling points is in a relatively narrow range of about 50°C- 60°C. The upper end is in part defined by the boiling heat transfer, in the sense that the temperature of the die will be $T_{die}=T_{sat}+Q_{die}R_{TH}$ where Q_{die} is the power of the device and R_{TH} is the source-to-sink thermal resistance, which will include the die, TIM, IHS and boiling etc. Since the $Q_{die}R_{TH}$ temperature rise term can be of the order of 10 °C - 30°C, depending on many factors, the saturation temperature of the fluid, T_{sat}, must not be so high as to push the die temperature above its operational limit (~80 °C –

90°C). One would then expect that lowering the saturation temperature would be beneficial, which can be true, though to a point. The condenser operates by providing water at a colder temperature than the vapor, which is at T_{sat}. As the saturation temperature gets closer to the water coolant temperature, the condensation heat transfer intensity diminishes i.e. in the limit if they were equal, there would be no condensation, so the system would pressurize to raise the saturation temperature (and thus the die temperature) to establish the differential required for condensation. Therefore, for TPIC operation at or near atmospheric pressure, maintaining a sufficient temperature difference between the water coolant and the dielectric fluid limits the minimum feasible boiling point of the coolant.

Key Take-Aways

- Two-phase liquid immersion as a cooling technology utilizes different, and more complex, mechanisms and systems for immersed cooling compared with its single-phase counterpart.
- In terms of thermophysical properties that promote intense boiling heat transfer, these fluids are not ideal and may require some relatively sophisticated enhancement technologies to ensure safe and reliable operation.

2.2.3 Performance

Although many publications exist on the topic of immersed boiling on plane and enhanced surfaces that emulate electronics (e.g. Gess et al. [38]), there are very few on practical electronic packages, like CPU processors. Also, like SPLICs, reliable data from primary references is very limited for TPICs, both at the single server level and at the system level. Again, much of what is reported in terms of thermal and energy performance is anecdotal. However, publications in the open literature are beginning to emerge which are providing some deeper insight into more system-level immersed two-phase cooling performance of compute devices. For example, Liu and Yu [39] performed experiments on ASIC crypto-miners (3 x T2T-25T boards with 140 T2T CPUs each) immersed in HFE 7100. The CPUs are relatively low power (~5 W), as is their heat flux (~8 W/cm^2), though combined could achieve up to around 2 kW. Tests were performed in a bespoke sealed case that included a condenser and associated water loop that cooled the water in an air-cooled CDU-type unit rated at 2.4 kW. The results showed that the system was capable of cooling the CPUs adequately at room temperature (~22°C), with core temperatures reaching just over 73°C at a total power of 1580

W, This is maybe not so surprizing considering the low heat flux of these CPUs. In terms of the PUE of the system, it ranged between 1.05-1.04, with the power penalty being associated with the water circulation pump and fan in the CDU. It is noteworthy that the condenser water temperature remained about 25°C- 35°C below the HFE 7100 saturation temperature of 61°C. In a subsequent study with a similar apparatus, Liu and Yu [40] considered a single immersed Asus Prime Z590M-Plus motherboard with an Intel Core i9-10900K CPU (TDP=125 W) and Gigabyte RTX 260 Super Gaming OC graphics card. Boiling HFE 7100 on the bare IHS of the CPU, this TPIC system was not able to achieve the TDP of the CPU, as it throttled once the core temperature reached 100°C. This illustrates the point made earlier, that for high power and heat flux CPUs, specialized surfaces are required to enhance the boiling heat transfer, preferably directly on the IHS surface to avoid requiring an additional layer of thermal interface material. The TIM issue was raised by Ramakrishnan et al. [41] who used an attached block with an exposed porous surface as boiling heat sinks, using a solder joint to avoid this potential chemical degradation. To the best of knowledge, the chemical interaction of low boiling point dielectrics and conventional TIMs, such as thermal grease, is an open question.

System-level performance data on TPICs is only recently being published, and from only one research group at Nanyang Technological University in Singapore [36,42,43]. This research group has constructed a TPIC that housed (i) 8 Supermicro GPU Superserver 1028 GQ servers with Intel Xeon E5-2683 CPUs (TDP=120 W) and total combined power of up to 2 kW, (ii) dummy boards with ceramic heaters to mimic additional servers, (iii) 545 L of low boiling point dielectric fluid (HFE 7100 / Novec 649), (iv) a serpentine coiled copper condenser, and (v) a pump and rooftop dry tower. In Wu et al. [36], HFE 7100 was used as the working fluid and discussed preliminary test results of a long-term test (24/7 for a month) with varying load from idle to full power. In Kanbur et al. [42], again using HFE 7100, provided more detailed information for a test campaign that spanned 6 days over 6 power load conditions [3.4 kW (Case 6)- 9.2 kW (Case 1)] and ambient temperatures that ranged from from 25°C to 29.2°C. Water in the condenser-dry tower loop was supplied at approximately 83 L/min. No information was provided with regard to the CPU utilization percentage or temperatures. Figure 4 shows the key finding in terms of energy performance, with the PUE being as high as 1.4 at about 35% full load which drops to 1.15 at 100% full load. Referring to Figure 2, it is noted that this performance is broadly similar to the SPLIC PUE levels over the same load interval, though this comparison is made cautiously considering the differences in the two system designs and operating conditions. This point is highlighted when considering the recent Kanbur et al. [43] results on the same facility, though with Novec 449 as the working fluid and a lower water coolant flowrate of 55 L/min. The tabulated data provided has been used here to calculate the PUE in this test campaign, which was system-level the same as the earlier Kanbur et al. [42] study, except the load range varied from 6.9 kW to 15.9 kW.

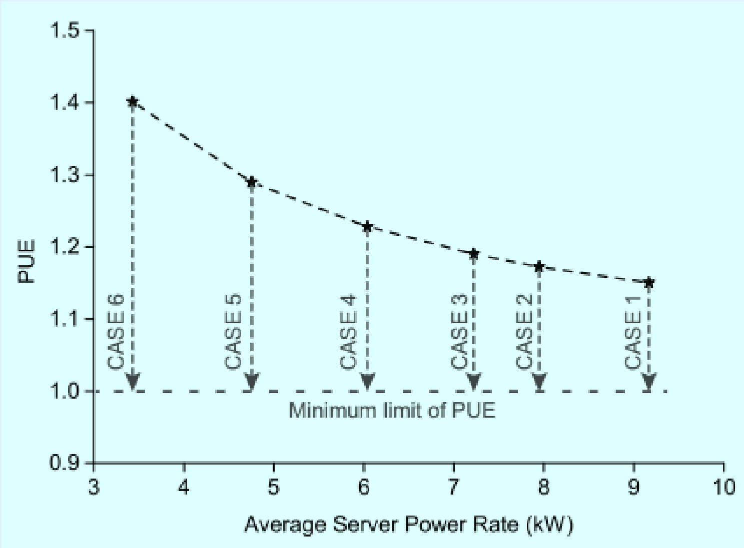

Figure 4: Average PUE of the TPIC system for varying load from 3.4 kW (Case 6) to 9.2 kW (Case 1) (plotted from Kanbur et al. [42] data)

Figure 4 shows the PUE, which can be considered partial since it does not include other facility power, though representative of the IT load and cooling infrastructure. Compared with their earlier study, the PUE levels are lower, and this must be due to the reduction in water flow rate from 83 L/min to 55 L/min. At an equivalent power loading of 9.2 kW, the earlier study determined a PUE=1.15, whereas here it is PUE=1.09. It is doubtful that the use of Novec 649 opposed to HFE 7100 plays a role since it has slightly worse properties in terms of those that influence pumping power (Table 2). For the maximum load that they could achieve for safe operation of the dummy boards with the electrical heaters (15.9 kW), the minimum PUE=1.05 was achieved. Compared with the SPLIC system tested in the same study, this is lower than the PUE=1.09 at roughly the same load power, though the SPLIC system was able to operate up to almost 30 kW, provided the pumps switched to higher speed operation to manage the liquid oil coolant temperature. It is also noteworthy that in this Kanbur et al. [43] study, the water coolant inlet/outlet temperatures were 34°C/38°C at the highest server load condition, showing that the system could operate at reasonably high operating temperatures, even with the lower boiling point (49°C) Novec fluid.

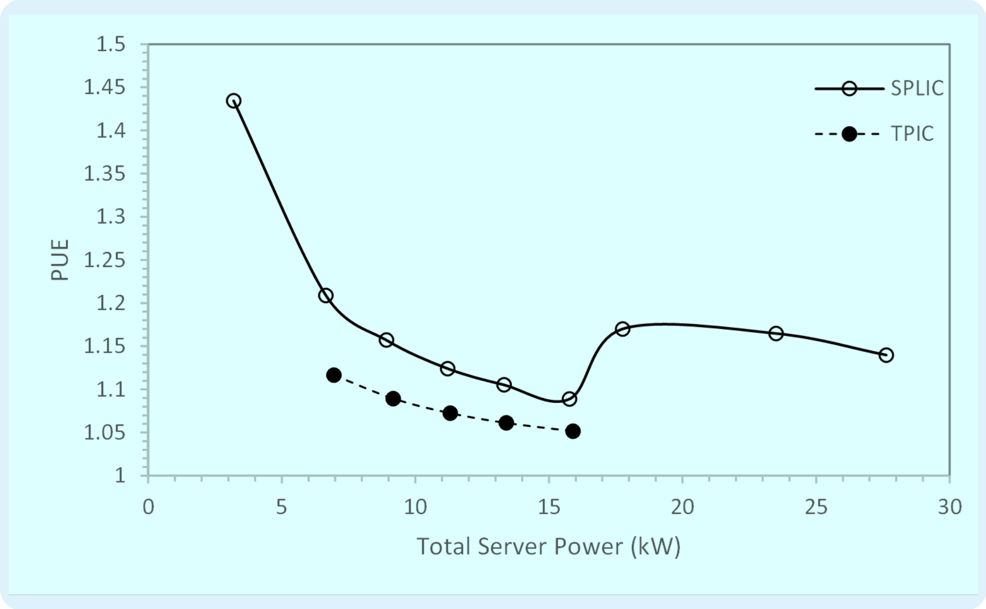

Figure 5: Average PUE of the TPIC system for varying load from 6.9 kW to 9.2 kW (plotted from Kanbur et al. [43] data)

One key outcome of this series of research studies is that it has provided some of the first data that gives a reliable indication of TPIC energy performance operation over a reasonable range of operating conditions. Clearly, the TPIC technology shows potential in the context of its more rational use of energy, with (partial) PUEs as low as 1.05 at high IT loads and lower water flow rates, though can be up to 1.4 for low IT loads and higher flow rates. What is clear is that there is a large difference between PUE over the range of IT load, and this is true for the SPLIC system as well. IT loads fluctuate as do the PUEs (even if it is a partial PUE in these cases), and this is mainly due to the cooling equipment (pumps, fans) running at fixed speeds over large ranges of IT loads. For the SPLIC system, the pumps and potentially fans must be managed to ensure that the liquid temperatures do not escalate to the point that the chip temperatures exceed safe levels i.e. they are roughly proportionate to each other. For TPIC systems, the server heat is dissipated to the coolant at its saturation temperature, which is set by the pressure of the system. An under-sized condenser and/or too low flow rate through it will cause the system pressure to rise in order to create a higher temperature differential between the vapor and condenser to compensate. This will not only cause the chip temperatures to rise accordingly, but creates new problems around complexity, safety and vapor leakage at elevated pressure. At present, no design guidelines appear to exist for the condenser in TPIC systems, despite it being a key engineering component. It seems sensible that they should be designed for the worse case operating conditions (climate, IT load etc.).

Considering this, it is clear that there is still missing data in the literature with regard to real-life PUE in both SPLIC and TPIC systems. IT equipment in data centre servers operate at variable IT loading, generally, so some clarity and differentiation between ideal PUE and actual PUE in service is needed.

Key Take-Aways

- There is a severe lack of system-level data to draw definitive conclusions regarding the cooling effectiveness and energy performance of TPIC technology.
- Data that is only beginning to emerge is mixed, though can be considered promising from an energy perspective, since it can under certain circumstances achieve PUEs much lower than conventional air-cooled systems.
- There needs to be a differentiation between 'best' PUE and time-averaged PUE reporting, and preferably under real data centre conditions, where CPU utilization fluctuates.

2.2.4 Some Key Open Questions

Here, some additional considerations are outlined in terms of the implementation of TPIC systems.

Material Compatibility: In the literature, more information is available for oil compatibility with IT hardware than for low boiling point dielectrics. Manufactures, such as 3M, do however provide fairly extensive documentation with regard to chemical interaction with many materials. However, this still needs independent verification and a much wider material suit that is specific to the wetted components in servers. Reliable data on how these coolants interact chemically with TIMs is certainly missing, especially considering that these fluids are also marketed as solvent cleaners.

Life Cycle Analysis (LCA): Fluorinated fluids are synthesized and there is an embedded energy and environmental cost associated with their manufacture, transport, and recycling. Any energy and CO2 savings associated with TPIC operation must be considered in the context of an LCA.

Global Warming Potential: The GWP of low boiling point dielectric fluids are non-zero, and vary considerably based on their chemical composition. GWP is a multiplier that quantifies the CO2 equivalent of released fluid. As per Table 2, Novec 649 has a low GWP=1, meaning that a kilogram of released vapor has an equivalent global warming impact of the same mass of CO_2. Other fluids like HFE 7100 is moderate, though FC 72 is very high, likely too high to be considered viable in the current environmental landscape. The importance of considering the GWP is exacerbated when considering how highly volatile these fluids are and that there is potential for vapor release to the atmosphere, in particular if the system is pressurized to mitigate air ingress. This is concerning to the point that some EU countries are implementing measures to reduce consumption of high-GWP fluorinated fluid, and this could include GWP-weighted taxes etc. [44].

Other Considerations [25]:

- Occupational safety e.g. toxicity, handling, service life
- Sealing requirements and evaporations losses
- Signal compatibility e.g. interference with I/O signal transmission may be more problematic with highly wetting low surface tension fluids
- Further environmental impacts e.g. recycling; soil and groundwater contamination
- The unexpected e.g. lifting of barcode stickers, dissolving ink labelling, debris accumulation etc.
- Cost, since fluorinated liquids are disproportionately expensive compared with more conventional coolants, large charge volumes are required, and loss of coolant more likely due to high volatility.

Key Take-Aways

- Chemical compatibility across the myriad of materials in compute electronics still needs significant work, including thermal interface materials
- It is unknown how a Life Cycle Analysis will influence the energy/CO_2 dynamics of TPICs.
- The Global Warming Potential of fluorinated liquids, along with other environmental factors, like disposal, recycling, and soil/groundwater contamination, must be taken very seriously.

2.2.7 Heat Recovery Potential of Immersion Cooling

To the best of knowledge, heat recovery from immersed cooling has not been demonstrated in the open literature. Theoretical scoping studies have been performed for heat energy reuse in a polymer electrolyte membrane [45] and desalination system [46] for two-phase immersion with HFE 7100 as the working fluid, otherwise it is unproven.

In SPLIC systems, heat recovery is likely a difficult proposition since it seems to be limited to coolant temperatures in the region of 40°C. It may then be better suited in a scenario where heat pumps are used to elevate the temperature, though this adds investment cost and increases electricity consumption to such an extent that its economic viability is questionable [47]. TPIC systems may have the potential to achieve higher recovery temperatures. However, this is limited by its boiling point. HFE 7100 would be the best choice in this scenario, though this would have to be balanced against its GWP being orders of magnitude higher than Novec 649. Regardless, the recovery temperature will never achieve the boiling point as subcooling is required to activate condensation. From the data currently available, the level of subcooling is too high to be considered viable for heat recovery at feasible condenser water outlet temperatures [43].

3. Direct Liquid Cooling - DLC

3.1 Conventional Direct Liquid Cooling

3.1.1. How it works

Direct Liquid Cooling technology implements coldplate heat exchangers, typically microchannel-type. Liquid is used as the coolant, and can be water-based since it does not 'touch' any electrically conducting components. The coldplates are mounted directly on the high-powered and thus high heat flux generating processors (CPU/GPUs). Sometimes, other moderately powered components, such as the DIMMs, are also liquid-cooled. The coolant can be supplied directly from the facility (single loop) or from a CDU (dual loop), and is pumped to the cabinets and then distributed to the servers in a distribution manifold or header. Like conventional air-cooled data centres, cooled air is forced into server by internal fans to cool auxiliary low powered components (DIMMs, hard drives, power supplies, VRM etc.). The air is cooled/conditioned/distributed by the DC infrastructure (CRAC/CRAH systems) as in conventional air-cooled data centres.

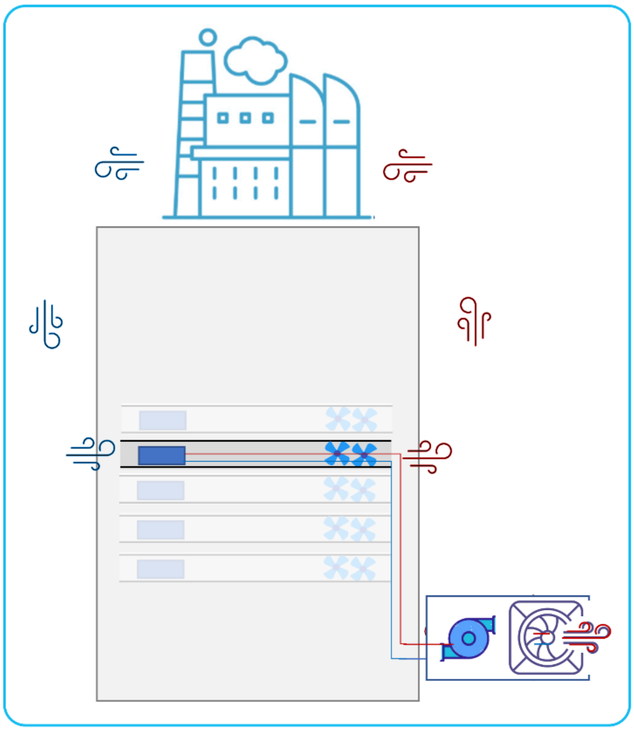

Figure 6: Schematic representation of direct-attach liquid cooling technology

3.1.2 Liquid Coldplate Heat Exchangers: Pros & Cons of the Microchannel Heat Sink

It has already been discussed why liquid is a better coolant than air owing to key thermophysical properties that influence the intensity of convective heat transfer (Table 1). Further drivers towards the adoption of DLC include escalating TDPs, reducing operating cost and PUE, chip reliability and heat recovery [48]. However, the heat exchanger design also plays a significant role in achieving these improvement goals. Specifically, what is needed in this particular instance is a compact heat exchanger that reduces, as far as feasible, the difference between the processor core temperatures and the coolant, and this must be achieved without too severe a pressure drop penalty. Reducing the chip-to-coolant temperature differential is in part achieved by reducing the thermal resistance of the coldplate. In the context of electronics cooling, the microchannel heat sink is the most researched and is the one which has seen the highest level of deployment in industry over the past decade or so, and this is particularly true for cooling of CPUs.

Liquid cooled microchannel heat exchanger technology dates back to the early 1980s and the seminal work of Tuckerman and Pease [49]. What was recognized here was that convective heat transfer intensification can occur when the scale of the heat exchanger was reduced to dimensions of the order of hundreds of micrometres. For example, for fully developed laminar flow in a channel, the convective heat transfer coefficient is proportional to the inverse of the channel dimension, $h \propto D^{-1}$. Compounding this is the notion of area intensification. Since the convective thermal resistance is the inverse of the product of the heat transfer coefficient and heat transfer surface area, $R_{TH} \propto (\boldsymbol{hA})^{-1}$, including microchannels onto an otherwise flat surface has the effect of substantially increasing both *h* and *A,* which together result in extremely low achievable overall thermal resistances. This is of course crucial since, from an engineering perspective, electronic components can be considered disproportionately powerful for their small size, and thus pose a significant challenge to cool effectively. However, intense convective cooling does not come without a penalty. For microchannels and other microfluidic heat exchangers, this is the hydraulic pressure drop and associated pumping power. For the fully developed flow in a channel case as an example, the pressure drop can be shown to increase as $\Delta P \propto D^{-4}$. Thus, the pressure drop increases at a rate that is disproportionally higher than the rate at which the heat transfer coefficient increases with decreasing channel dimension. This drastic pressure drop increase with reduced channel size is one key limitation in terms of how 'small' is feasible in terms of microchannel heat exchangers. Other challenges, such as manufacturing and fouling/clogging of channels are also important considerations at this scale.

Very briefly, microchannel heat exchangers have a long history of scientific and engineering development. The scientific development phase lasted from the early 1980s to the early 2000s, when the open questions were finally closed [e.g. 50-52]. Subsequent to this, and in parallel with the advances in CFD simulation potential and computing power, a surge in publications on the topic of microchannel science and technology has occurred over the past 15 years or so, with the preponderance being focussed on engineering of novel microchannel-type heat exchangers. Over this relatively short time span, in the region of 8000 papers have been published on the topic, representing about 90% of the total archived publications in the field since it commenced in the early 1980s. This has largely been fuelled by the problems facing the electronics packaging industry due to the escalating heat fluxes and the requirement of compact thermal management solutions.

A key challenge in the commercial translation of liquid-cooled microchannel coldplate technology is manufacturing. This is particularly complicated when these small features are fabricated in dense and ductile metals, such as copper, which is ideal because of its high thermal conductivity. Conventional subtractive manufacturing, like milling, of hundreds of high aspect ratio channels at the scale of 200-400 μm is challenging, and this influences production rate and cost. To the best of knowledge, the manufacturing method of commercial microchannel coldplates is predominantly skiving.

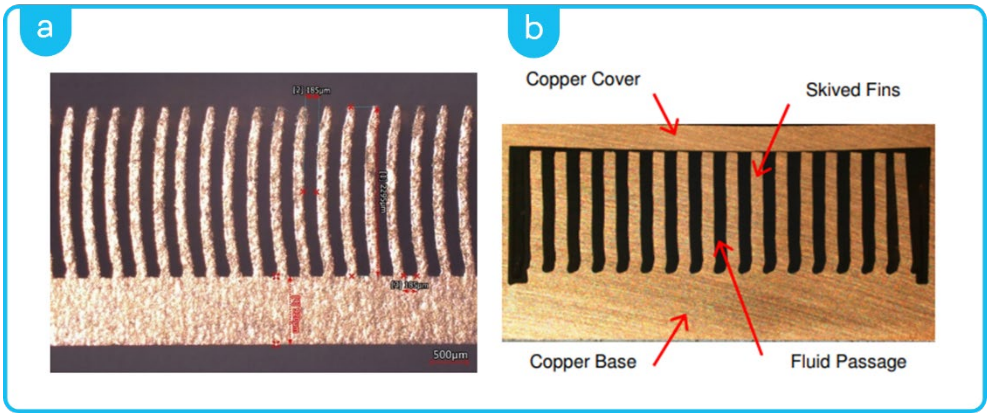

Figure 7: Skived copper microchannels. (a) 185 µm width [53], (b) 428 µm width [54]

Figure 7 shows examples of skived copper microchannel heat sinks, which illustrates some typical attributes. First is the width dimension, which is in the range of ~ 200 µm -500 µm, as anything smaller for manufacturing, structural integrity and pressure drop becomes problematic. Anything larger begins to diminish both the convective heat transfer and surface area intensities, thus is detrimental to the thermal resistance. It is also noticed that the aspect ratio of the channels is quite high, in the sense that the channel heights are significantly larger than the channel width, by a factor of about 10. This of course is motivated by increasing of the area intensity, though fins cannot be too high and thin since (i) fin efficiency will reduce resulting in diminishing gains, (ii) for a given total volumetric flow rate, the increasing channel cross-sectional area will decrease the flow per channel, resulting in higher sensible heating of the fluid (i.e. the fluid temperature along the channel will rise more diminishing local heat flux and overall heat transfer effectiveness). Fortunately, since liquid microchannels tend to operate in the laminar flow regime and is fully developed along most of its length, reducing the flow rate, and thus velocity, does not affect the heat transfer coefficient, since the Nusselt number is constant for this flow regime (i.e. independent of Reynolds number). However, changing the aspect ratio also has the influence of changing the channel hydraulic diameter, which does influence the Nusselt number. This entanglement of aspect ratio and hydraulic diameter was resolved in a clever paper by Sahar et al. [55]. Regardless, the simple geometry, laminar flow conditions and repeatable geometry with symmetry planes, allow for relatively straight forward and accurate CFD simulation models to be created, and these can and should include temperature dependent fluid properties, conjugate heat transfer and viscous heating influences.

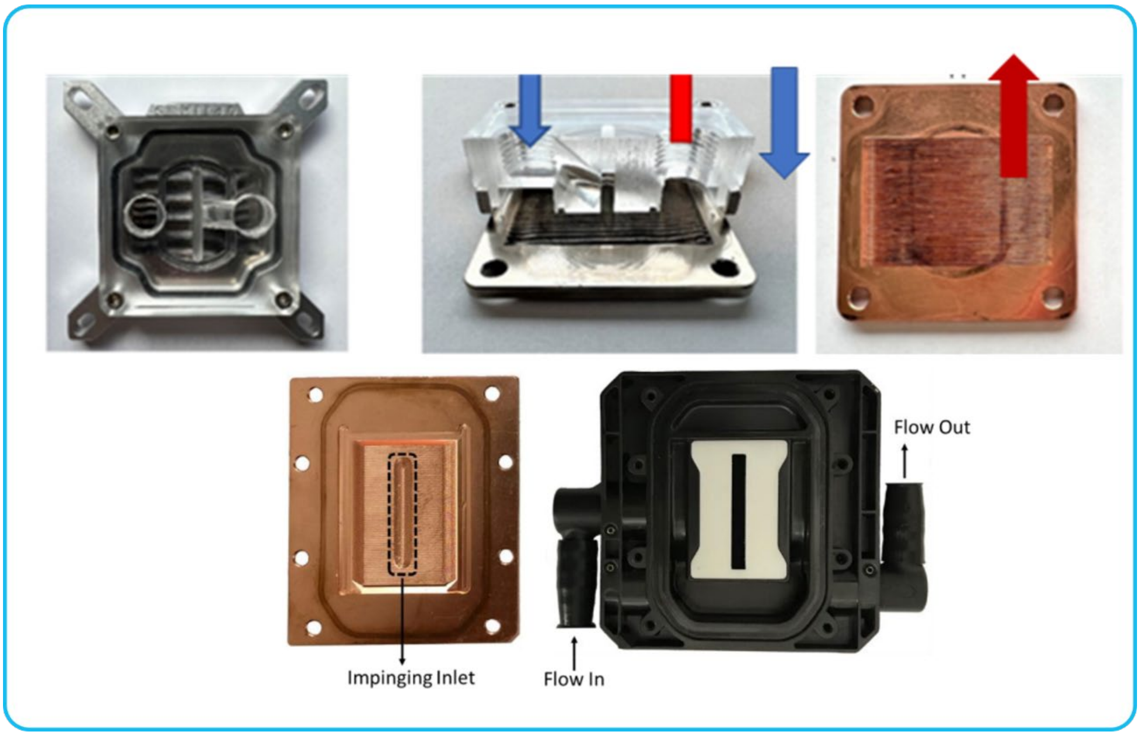

Figure 8: Skived microchannel heat sink with slot jet inlet manifold (Top) from Radmard et al. [53], (Bottom) from Hadad et al. [56]

Perhaps due to the skiving process being so fast and inexpensive, it seems that most commercial microchannel heat sinks use the same basic design, with only small differences between them. The split-flow design is depicted in Figure 8. Here, fluid enters through an inlet port and is routed to a slot jet that runs along the width of the channel. The liquid flows into the channels from the slot jet and then bifurcates into the right-side and left-side array of channels, exiting at the far ends. The base of the slot jet manifold acts as the upper confining wall of the microchannels, containing the fluid flow and ensuring it is unidirectional. The spent flow from the ends of each channel array are then accumulated and routed to the exit port of the manifold. Since the microchannel heat exchanger must cover a reasonable surface area (commensurate with the CPU IHD size, for example), this is one of the best microchannel array designs since it mitigates the severe pressure drop and temperature rise that would be associated with a unidirectional channel array. Here, the flow is split into two array circuits (left/right), which are hydraulically in parallel, and thus each channel has half the velocity and length compared to channels running the full length of the heat sink. Furthermore, fresh coolant is introduced in the centre, pushing sensible heating influences to the outer region, away from the die hot spot (presuming it is centrally located). Further to this, there is likely some local enhancement where the slot jet impinges, which is situated above the region of the die hot spot.

Regardless, this design suitably addresses the main pitfalls of microchannel heat exchanger design, and is straight forward to optimize using modern numerical tools [56,57]. It is thus not surprizing that so many commercial microchannel heat sinks use this split-flow microchannel design. Figure 9 shows a few of these, and illustrates that there is in fact very little to distinguish the overarching design between many commercial microchannel coldplates, apart from styling, form factor, fluid routing/manifolding. Considering the maturity of this technology and ease of CFD simulation, it is unlikely that there is substantial difference in the thermal-hydraulic performance between different commercial heat sinks of this type, nor should it be expected that a marked improvement in performance is possible since these should be optimized.

On the topic of performance optimization, it is worth noting that there is antagonistic tension between what is optimal for thermal-hydraulic performance, and what is practical in terms of high-volume and economically viable manufacturing. Low-cost skiving limits the parameter space for optimization, which on one hand is beneficial since it simplifies the optimization due to the uncomplicated geometry. On the other hand, it is limiting because there exists broad scope for alternative geometries and patterns that are well known to outperform linear channels (e.g. offset interrupted fins [58], zigzag [59]), though manufacturing at the microscale constrains these. For electronics packages that contain a small die fixed to a heat spreader, as is the case for commercial CPUs, the additional concept of *hotspot targeting* should also be addressed. As discussed in Robinson et al. [60], localized heating from a die can create heat fluxes so high that an IHS is required. The IHS ostensibly acts as a horizontal fin to increase the area available for heat transfer to the coldplate heat exchanger. However, this creates a hotspot (high heat flux zone) in the region of the die [61]. This then becomes a conjugate heat transfer problem, where the cooling intensity influences the heat conduction in the solid phases. If cooling intensity can be focussed, it is a substantially different optimization problem compared with linear microchannels, where the heat transfer intensity is relatively uniform over its cooling area. This underpins the emerging field of *targeted cooling*, where either or both surface structuration and flow mechanics are designed in such a way as to generate higher cooling intensity in the region of the hotspot. Importantly, doing so can target the minimization of multiple objective functions, such as maximum temperature, average temperature, and temperature gradient (important for reduction of thermal stresses) and described recently by Elliot et al. [61]. In the context of liquid cooled coldplates, this is likely the next development phase of electronics liquid cooling research, since the hydraulic limitations (flow rate, pressure drop, pumping power) demand that the available liquid should be used to its maximum cooling potential.

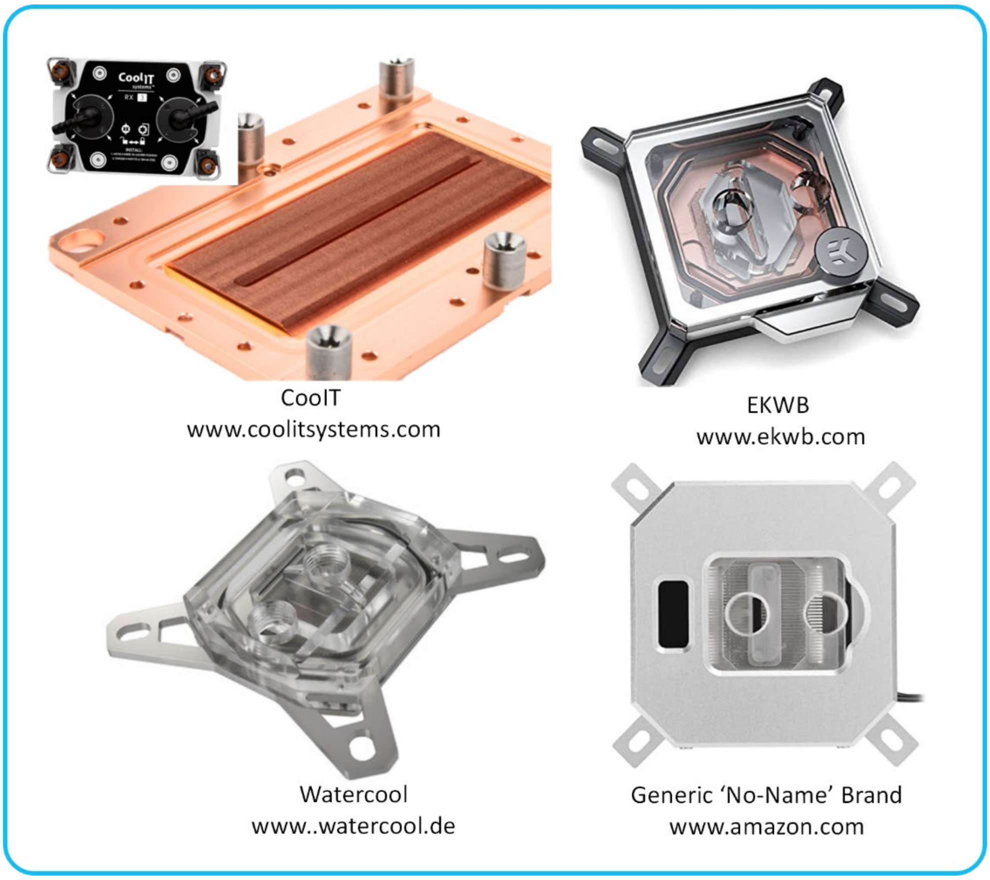

Figure 9: Examples of commercial microchannel coldplates using ostensibly the same split-flow microchannel technology

Another limitation of microchannel heat sinks, and this is true for any coldplate that uses surface features to increase convective heat transfer and/or surface area intensity, is that it necessarily requires a solid (metal) base onto which the features are machined. This has the benefit of sealing the coolant in the liquid flow loop, isolating it from the electronics. However, there is a penalty associated with the requirement of a thermal interface material, here termed TIM2, between the package and the coldplate. Compared with air cooled scenarios, where the air-side heat sink occupies so much of the thermal budget that the TIM2 resistance is a small proportion, this is not true in high performance liquid coldplate solutions. TIM2 now becomes quite important in the context of its percentage of the source-to-sink thermal resistance. Eliminating TIM2 is technically feasible if the liquid coolant can directly contact the exposed metallic face of the IHS. This of course is possible, since it only requires that a reliable seal be created between the IHS and the coldplate manifold. However, the IHS would be unstructured, which essentially eliminates microchannels and other surface feature enhancement designs. Direct-to-IHS integrated cooling is one of the most promising design approaches for some next generation processor cooling, where the TDPs and associated heat fluxes may strain cooling demand to the extent that TIM2 may need to be eliminated. However, this will require a design that optimizes the heat transfer on the flow delivery side, opposed to using surface features. Importantly, by focussing on the liquid flow dynamics, direct-to-IHS is also an ideal scenario for optimized targeted cooling.

Fouling, corrosion and clogging are areas of concern for microchannel heat exchangers. This is particularly relevant to DLC deployment in data centres since they can influence performance, maintenance and downtime. Fouling refers to the scenario where small particles entrained in the fluid become attracted to the surface causing a layer to form and grow. This can cause a reduction in the channel cross sectional area, influencing the hydrodynamics (flow distribution, overall pressure drop & heat transfer) as well as creating a barrier to heat transfer (additional thermal resistance) [62,63]. In fluid networks, alterations in the hydraulic resistance of components can also interfere with the flow distribution to the servers in a cabinet. In the worst case, particle built up will progress to the point of entirely clogging the channel [64], and this concern intensifies in microchannels owing to the high *size-ratio* (size of the particle diameter to the channel size) [65]. Corrosion refers to the chemical interaction of the fluid and the solid that causes dissolution of the solid material into solution and formation of brittle oxide layers. Cathodic corrosion is of particular concern for copper-based microchannel heat exchangers [53]. This can cause build up of particulates in the fluid and cause *corrosion fouling* [63] and/or *simple size exclusion* [65] (occurs when the particle size is larger than the channel size) in coldplate channels, as well as other hydraulic hardware in the system cooling loop [66]. Avoiding fouling, corrosion and clogging requires a holistic approach since is depends on several factors, including but not limited to the liquid quality and chemistry, the surface quality and chemistry, and the fluid-structure interaction. In terms of those specific to microchannel coldplates, at least the following may be considered as problem areas;

- Manufacturing of microchannels:
    - Rough and or burred surfaces promote fouling [63]
    - Exposed copper is prone to corrosion [53]
- Design of microchannels:
    - Required channel sizes on the order of 200 μm – 400 μm for optimal thermal-hydraulic performance, creates high size-ratio and are more prone to simple size exclusion

- Fluid flow in microchannels:
    - Low channel velocities and low vorticity flows (i.e. developed laminar flow) have lower shear stress and mixing capability, and thus less propensity to avoid adherence of impurities [63]
    - Stagnation regions in headers and manifolds are known to cause particulate agglomeration on fin ribs and/or surface fouling [67]

Key Take-Aways

- Best practice DLC coldplate technology appears to be the split-flow microchannel heat sink design. They can produce very low thermal resistance at moderate pressure drop.
- Split-flow microchannel heat sinks are mature, to the point of being the common underlying design in a host of commercial coldplates. They are at or near the end of their optimization design cycle.
- Microchannel heat sinks, in particular with exposed copper, are particularly susceptible to thermal-hydraulic performance degradation due to fouling, corrosion and/or clogging.

3.1.3 Performance

The key engineering motives for Direct Liquid Cooling are as follows;

- Coldplates offer such low thermal resistance that they significantly reduce the temperature difference between the CPU die and the coolant, to the extent that elevated coolant temperatures are feasible.
- This extra headroom in the thermal budget means that above-ambient coolant temperatures can be used, which eliminates the need to chill the coolant; the coolant can dissipate heat directly to the ambient. This is why it has been termed *hot water cooling*.
- This simplifies the infrastructure to circulate and condition the coolant and reduces the energy required to do so , improving the PUE.
- The increase in coolant temperature due to sensible heating of the liquid can be to the extent that it is at high enough quality that it can be reused.

Historically, IBM have been champions of hot water DLC technology. Iyengar et al. [68] discussed some of the first system-level performance data for a 15 kW fully populated liquid-cooled rack (38 1U IBM X3550 M3 servers, max. 400 W server node power). DLC was provided to the CPUs (TDP~130 W) via coldplates as well as the DIMM cards (~6 W each). All other components, like hard drives, power supplies etc. were air cooled. The system was designed for a pre-rack (inlet) coolant temperature of 45°C. The cooling system itself included only 3 power consuming components in two closed loops. The inner loop comprised a water circulation pump and the outer loop a circulation pump and dry cooler. The two loops combined at the CDU which housed a liquid-liquid heat exchanger. The main results are given in Figure 10, which shows that not only was the cooling power only a fraction (~3.4%) of the IT load, but the CPU and DIMM temperatures were well within safe operating temperature limits despite it being a relatively hot day over the 22 hour test interval. On the same facility, the IBM team also performed one of the first impact studies on varying operating conditions, such as flow rates, flow configuration in CDUs, dry-cooler fan speed, and addition of propylene glycol conditions [69]. They were able to determine ideal operating conditions for cooling energy minimization for a 13.2 kW IT load with 40°C outdoor temperature, using only 1.6% IT power for cooling, which is an impressive COP=64, illustrating that control and optimization is a significant part of data centre energy dynamics [70]. More recently, Yuksel et al. [71] presented IBMs hybrid air-water cooled IBM POWER AC922, touted as the fastest supercomputer in the world at the time, with coldplate-cooled IBM POWER9 processors and multi-component and NVIDIA SXM2 GPUs.

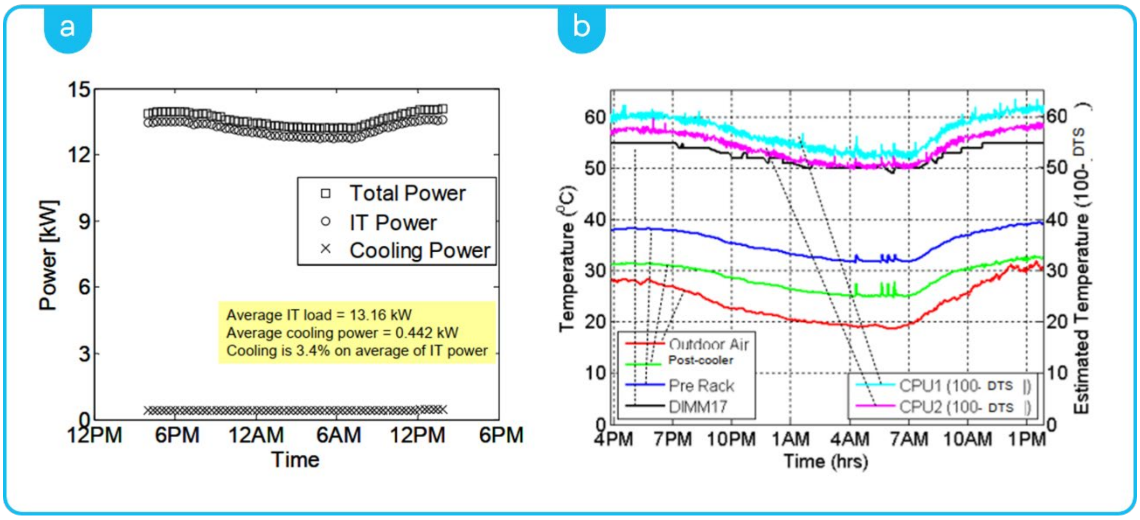

Figure 10: Results of a 22-hour run documenting (a) thermal/power data, (b) Temperature data on a hot summer day [68]

Druzhinin et al. [72], in partnership with Intel, tested a single benchmark DLC dual socket Intel ServerBoard S2600KPF (two Intel Xeon E5-2697 v3 -14 cores-2.6 GHz-TDP=145 W, 64 GiB (8×8 GiB) of DDR4-2133 memory, 120 GB solid state drive) as well as a 24 node 13 kW cluster. Tests were performed for varying water inlet temperatures from 19°C to 65°C at a flow rate of just under 2 L/min per server. For the cluster-level experiments, the heat absorbed in the servers was dissipated directly to the outdoor environment via a dry cooler. Interestingly, for cooler outdoor temperatures, free cooling was possible for coolant temperatures to about 50°C. Consistent with the Iyengar et al. [68] study, very low PUEs of between 1.04-1.08 were achieved, with the lower level being theoretically feasible for a full rack of 150 nodes. At elevated coolant temperatures, there was a modest decrease in power efficiency (~10%) due to leakage current, and this was largely from chipset components and voltage regulators, opposed to CPUs and memory. As noted, this would be more than offset by overall efficiency gains and the potential for energy reuse. This point was reiterated by Shoukourian et al. [73], where a small drop in power efficiency was measured in the 2 MW DLC SuperMUC Petascale supercomputer. Using a glycol water mix as cooling fluid ,the chiller-less high-temperature DLC supercomputer achieved COPs in the region of 20, which is impressive though associated PUE values were not reported. An important point brought up in this work, which is one of the few full scale data centre studies published, was the need for better control and optimization of the DLC cooling infrastructure, in the sense that efficiency is not only about the cooling equipment, but how they are controlled and managed for varying IT loads and outdoor conditions. This echoes the earlier findings of David et al. [69] and was investigated in more detail by Lucchese et al. [74] who illustrated that lower order modelling of the thermal-hydraulics can be leveraged to dynamically control the coolant supply to regulate component temperature with minimum coolant supply cost, and even maximize liquid outflow temperature for heat recovery applications.

An important attribute of implementing DLC at the rack-level is the flow distribution to the server nodes. This was investigated by Kadhim [66] and Kadhim et al. [75]. In these works, split-flow microchannel coldplates were used to cool the CPUs in a rack with thirty 1 U Sun Fire V20z servers (dual core 2 GHz AMD Opteron 64-bit CPUs, hard disc, eight DIMMs), whilst the rest of the components were air cooled with chilled facility air (Carel spray evaporative cooling system). It is first noteworthy that during the initial long-term testing, Kadhim [66] reported significant corrosion fouling that caused the flow rate to decrease to levels that caused CPUs to exceed tolerated operating temperatures. This is depicted in Figure 11, which shows the flow rate deterioration after about 3 months of operation and the degraded split-flow microchannel coldplate, highlighting points made earlier with regard to potential pitfalls with copper microchannel coldplates. With regard to flow distribution, the studies found that flow maldistribution, with the top servers receiving about 30% higher flow rate than the lowest server, and this caused the lower CPU temperatures to be higher. This was caused by a poorly designed distribution manifold, and simply increasing its diameter so that it behaved closer to constant pressure plenum improved the flow and temperature distribution markedly.

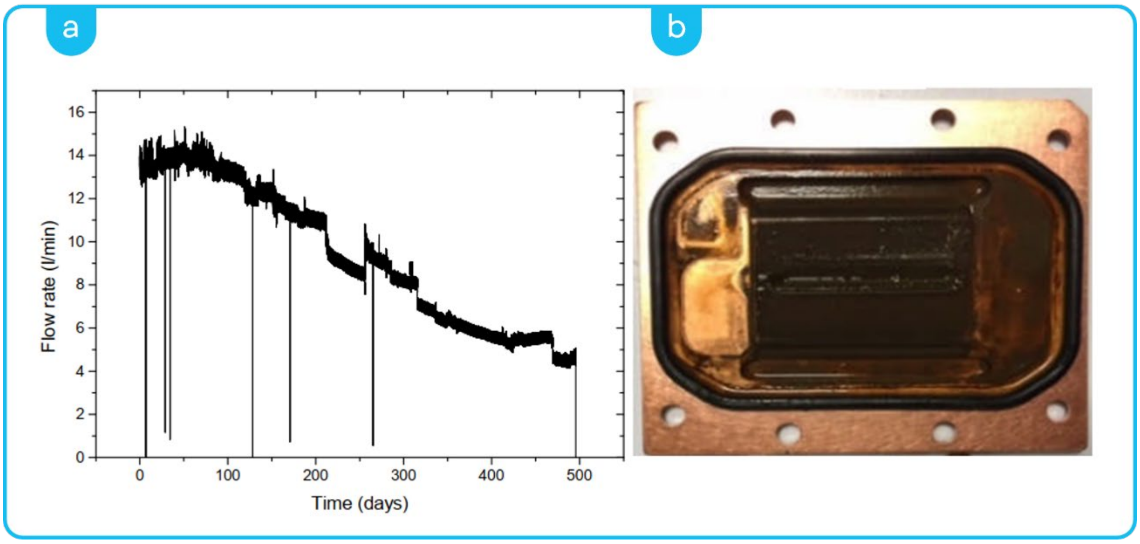

Figure 11: (a) Flow rate degradation over long interval testing, and (b) degradation of microchannel coldplate [66]

In terms of energy performance, Kadhim [66] investigated the influence of several parameters, including flow rate, inlet temperature and workload, with the later from idle to 100% usage. In their setup, a rear door heat exchanger managed the cooling of the air, and a bespoke Air handling Unit (AHU) dissipated all of the server heat to ambient. The measured PUE was weak function of the liquid flow rate, and this was because the AHU dominated the cooling power budget. However, liquid inlet temperature greatly influenced the PUE. For example, at 100% IT load operation it decreased between 1.2 to 1.025 for inlet temperatures between 24°C and 42°C, and this is simply because the AHU had to work harder to achieve lower liquid temperatures. The use of the rear door heat exchanger in the Kadhim [66] study raises an interesting question regarding the benefit of DLC compared with this alternative liquid cooling scenario. This was explored in Ovaska et al. [76] who investigated both rear door liquid cooling (RDLC) only and hybrid RDLC-DLC of a 47 node ~20 kW HPC operating at 100% load. They found that the RDLC-DLC not only reduced the cooling power consumption by 9.4%-14.4%, depending on inlet air temperature, but also improved the cluster performance per Watt (i.e. MFLOPS/W) considerably, by between 10.9% to 18.2%, since the DLC CPUs were cooler. This is an interesting point when considering hot water cooling and leakage current, in the sense that it is true that server power efficiency can increase with liquid supply and thus CPU temperature, but this impact should be compared against the air-cooling alternative.

Some other points worth noting are as follows:

- For the stated DLC PUEs:
  - Includes server fan power as part of the IT load
  - Should be considered partial as they do not include lighting, UPS, fire suppressions etc.
- Water leakage: this has not been reported to be an issue, so may be a perceived problem opposed to a real problem.
- Air infrastructure: Some hybrid DLC system still requires facility air conditioning and handling systems which will influence cost and PUE. Others use rack-level air-water heat exchangers, like rear-door heat exchangers.
- Noise: The hybrid DLC still includes internal fans and through-flow of air, so fan noise is likely still to be problematic.
- Dust/Moisture: The servers are still susceptible to dust and moisture ingress. This may be particularly limiting in edge computing applications.

Key Take-Aways

- Conventional hybrid DLC systems that use coldplates for the higher-powered server electronics and air to cool peripheral components can achieve exceptionally low (partial) PUEs.
- In the context of liquid cooling solutions discussed thus far (SPLIC and TPUC), DLC is consistently and by far the lowest PUE system.
- DLC can provide low enough thermal resistance that hot water cooling is possible, eliminating the need of sub-ambient 'chilled' coolant, even with hot outdoor conditions.
- Dynamic control of cooling system can improve already low PUEs.
- Flow distribution to server nodes in a rack requires special attention to ensure even flow distribution and subsequent operating temperatures.
- Small drop in server power efficiency will occur at elevated temperatures, though must be considered in the context of overall reduction in energy use and potential heat reuse.
- Unlike SPLIC and TPIC systems, hybrid DLC systems retain some of the less positive attributes of conventional air-cooled systems since air is still used as a cooling medium.

3.1.4 Heat Recovery

One of the earliest mentions of direct reuse of heat energy from data centres was in Brunschwiler et al. [77]. They recognized that the advancement of very low thermal resistance coldplate technology lowered the junction to ambient temperature to such a degree that it facilitated the use of 'hot' water whilst still maintaining safe chip operating temperatures. At such high coolant temperatures (>50°C), the heat energy could be used directly, without the need to raise the temperature with a heat pump (like conventional air-cooled systems), potentially leading to sub-unity PUEs. Apart from direct use of heat 'as heat', which may not be on demand continuously, other uses of the low-grade heat, such as adsorption refrigeration, have been shown to be feasible [e.g. 78,79], which broadens the potential for overall efficiency improvement and related CO2 reduction.

The first system-level demonstration of hot water-cooled data centres for heat recovery application was by Zimmermann et al. [80], which was closely related to the more foundational theoretical work performed in Zimmermann et al. [81]. Here, the IBM Aquasar is described, where just under 7 kW of blade server capacity was retrofitted to enable DLC of all components over 3 W. The hydraulic system comprised of three separate loops. The inner pumped loop used de-ionized water with a small percentage of corrosion inhibiter and dual filtration to supply the servers with high purity coolant, since the CPU coldplates were microchannel-based. The middle pumped loop transferred heat from the coolant in the primary loop via a liquid-liquid heat exchanger and to the heating grid of the ETH Zurich campus. Various control valves were used to maintain a constant heat recovery temperature with varying loading on the servers, as well as to mitigate component overheating. The main results of this study were that a PUE of 1.15 could be sustained over a range of water temperatures up to 60°C. Energy Reuse Efficiency, which is a modified PUE which subtracts the energy reused from the total IT and cooling power (lower is better i.e. if no heat reuse, ERE=PUE), was also calculated at about 0.4. It was noted that the heat recovery is more useful at higher water temperatures, though this tended to increase parasitic heat losses which could be improved with insulation. Overall, about 80% of the heat was recovered at the highest water temperature of 60°C at 100% IT load.

Meyer et al. [82] detail the IBM iDataCool hot water cooled HPC cluster at the University of Regensburg (IBM System x iDataPlex architecture, three racks with 72 compute nodes each, each node with either two four-core Intel Xeon E5630 (44 in total) or two six-core Intel Xeon E5645 (388 in total, and 24 GB of memory arranged in six 4 GB DDR3 DIMMs). Here, and impressive 70°C water temperature was achieved, which was sufficient to run an adsorption chiller. Although a minor increase in the power consumption of the nodes was observed with increasing coolant temperature, this was more than offset by the energy recovered in the adsorption chiller, and was not compared against the air-cooled equivalent. Again, poor insulation and unwanted heat losses influenced the overall heat recovery performance, and could theoretically achieve 50% for this adsorption refrigeration systems if these losses were mitigated. It is also noted that, at the time of publication, the system had been operational for over a year with no apparent negative effects, providing some positive indicators with respect to DLC and data centre system reliability. For completeness, the CoolMUC hot water cooled HPC [83] has also published positive results on heat recovery for use in adsorption refrigeration, with cluster feed temperatures as low as 50°C and refrigeration COP of up to 18.

Some other points worth noting in the context of heat recovery are as follows:

- High levels of heat recovery from the servers require almost full direct contact liquid coolers on the electronic components, which increases complexity and pumping power.
- Most of the published research is on HPC servers and may not be the same for IT servers which may contain different hardware (power supplies, hard drives etc) and power partitioning.
- In hybrid air-DLC heat recovery systems, a significant portion of the available energy can be missed unless additional water-cooling systems are deployed, like rear-door heat exchangers.

Key Take-Aways

- Heat recovery has been demonstrated for both direct use and adsorption refrigeration.
- Coolant temperatures as high as 70°C have been reached without notable penalty on computing efficiency and reliability.
- Overall PUEs and EREs are promising and can be much improved with mitigation of unwanted heat losses from uninsulated equipment and piping and full heat recovery from server components.
- Available data is for HPC servers, which may or may not be transferable to IT data centres.

References

[1] M. Saini and R. L. Webb, Heat rejection limits of air cooled plane fin heat sinks for computer cooling, IEEE Transactions on Components and Packaging Technologies 26(1) (2003) 71-79.

[2] A.C. Kheirabadi, D. Groulx, Cooling of server electronics: A design review of existing technology, Applied Thermal Engineering 105 (2016) 622–638.

[3] Kheirabadi and Groulx (2016)-roadmap towards the inevitable use of liquid cooling fin data centres from a thermal perspective – so its not just PUE (190 W tap-out)

[4] S.I., Guggari, Analysis of Thermal Performance Metrics-Application to CPU Cooling in HPC Servers, IEEE Transactions on Components, Packaging and Manufacturing Technology 11(2) (2021) 222-232.

[5] N. Jones, Data centres are chewing up vast amounts of energy - so researchers are trying to make them more efficient, Nature Magazine 561 (2018) 163-167.

[6] E. Masanet, A. Shehabi, N. Lei, S. Smith, J. Koomey, Recalibrating global data center energy-use estimates, Science 367 (2020) 984-986.

[7] Uptime Institute, 2021 Data Center Industry Survey Results.

[8] H. Chen, D. Li, Current Status and Challenges for Liquid-Cooled Data Centers, Frontiers in Energy Research, 10 (2022) 952680.

[9] Chhetri, D. Kashyap, A. Mali, C. Agarwal, C. Ponraj, N. Gobinath, Numerical simulation of the single-phase immersion cooling process using a dielectric fluid in a data server, Materials Today: Proceedings 51 (2022) 1532–1538.

[10] S. Vishal, P. Vikas, Transformer's history and its insulating oil, 2011.

[11] N.A. Pambudi, A. Sarifudin, R.A. Firdaus, D.K. Ulfa, I.M. Gandidi, R. Romadhon, The immersion cooling technology: Current and future development in energy saving, Alexandria Engineering Journal (2022) 61, 9509–9527.

[12] S. Oktay, Multi-liquid Heat Transfer, 555,730, 1968

[13] B.B. Kanbur, C. Wu, S. Fan, F. Duan, System-level experimental investigations of the direct immersion cooling data center units with thermodynamic and thermoeconomic assessments, Energy 217 (2021) 119373.

[14] R. Eiland, E.J. Fernandes, M. Vallejo, A. Siddarth, D. Agonafer, V. Mulay, Thermal Performance and Efficiency of a Mineral Oil Immersed Server Over Varied Environmental Operating Conditions, ASME Journal of Electronic Packaging 139(4) (2017) 041005.

[15] J.M. Shah, C. Bhatt, P. Rachamreddy, R. Dandamudi, S. Saini, D. Agonafer, Computational Form Factor Study of a 3rd Generation Open Compute Server for Single-Phase Immersion Cooling." Proceedings of the ASME 2019 International Technical Conference and Exhibition on Packaging and Integration of Electronic and Photonic Microsystems. ASME 2019 International Technical Conference and Exhibition on Packaging and Integration of Electronic and Photonic Microsystems. Anaheim, California, USA. October 7–9, 2019.

[16] F. P. Incropera, D. P. DeWitt, T. L. Bergman, A. S. Lavine, Fundamentals of Heat and Mass Transfer, 6th Ed., John Wiley & Sons, New York, NY, USA, 2007.

[17] C.C. Cheng, P.C. Chang, H.C. Li, F. Hsu, Design of a single-phase immersion cooling system through experimental and numerical analysis, International Journal of Heat and Mass Transfer 160 (2020) 120203.

[18] Y.Q. Chi, J. Summers, P. Hopton, K. Deakin, A. Real, N. Kapur, H. Thompson, Case Study of a Data Centre Using Enclosed, Immersed, Direct Liquid-cooled Servers, Case study of a data centre using enclosed, immersed, direct liquid-cooled servers," 2014 Semiconductor Thermal Measurement and Management Symposium (SEMI-THERM), 2014, pp. 164-173.

[19] K. Haghshenas, B. Setz, Y. Bloch, M. Aiello, Enough Hot Air: The Role of Immersion Cooling, submitted May 2022.

[20] V. Mulay, "Open Compute Project – Server and Data Center Design" Panel presentation at ASME InterPack 2011, Portland, Oregon, 2011.

[21] Y. Ngoko, N. Saintherant, C. Cerin, D. Trystram, How future buildings could redefine distributed computing. In 2018 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW), (2018) 1232–1240.

[22] J.M. Shah, R. Eiland, P. Rajmane, A. Siddarth, D. Agonafer, V. Mulay, Reliability Considerations for Oil Immersion-Cooled Data Centers, ASME Journal of Electronic Packaging 141(2) (2019) 021007.

[23] S. Ramdas, P. Rajmane, T. Chauhan, A. Misrak, D. Agonafer, Impact of Immersion Cooling on Thermo-Mechanical Properties of PCB's and Reliability of Electronic Packages, Proceedings of the ASME 2019 International Technical Conference and Exhibition on Packaging and Integration of Electronic and Photonic Microsystems. ASME 2019 International Technical Conference and Exhibition on Packaging and Integration of Electronic and Photonic Microsystems. Anaheim, California, USA. October 7–9, 2019. V001T02A011.

[24] G.A. González-Reyes, S. Bayo-Besteiro, J.V. Llobet, J.A. Anel, Environmental and Economic Constraints on the Use of Lubricant Oils for Wind and Hydropower Generation: The Case of NATURGY, Sustainability 12 (2020) 4242.

[25] X. Lina, X. Yuping, L. Bin, Research on key issues of coolant in immersion liquid cooling of data centers, Information and Communications Technology and Policy 48(3) (2022) 40–46.

[26] J.M. Shah, S.H.I. Rizvi, I.S. Kota, S.R. Nagilla, D. Thakkar, D. Agonafer, Design Considerations Relating to Non-Thermal Aspects of Oil Immersion Cooling, Proceedings of the ASME 2016 International Mechanical Engineering Congress and Exposition. Volume 10: Micro- and Nano-Systems Engineering and Packaging. Phoenix, Arizona, USA. November 11–17, 2016. V010T13A054.

[27] R.J. MacNamara, T.L. Lupton, R. Lupoi, A.J. Robinson, Enhanced Nucleate Pool Boiling on Copper-Diamond Textured Surfaces, Applied Thermal Engineering 162 (2019) 1141452.

[28] P. Warrier, A. Sathyanarayana, D.V. Patil, S. France, Y. Joshi, A.S. Teja, Novel heat transfer fluids for direct immersion phase change cooling of electronic systems, International Journal of Heat and Mass Transfer 55 (2012) 3379–3385.

[29] I.L. Pioro, W. Rohsenow, S.S. Doerffer, Nucleate pool-boiling heat transfer II: assessment of prediction methods, International Journal of Heat and Mass Transfer 47 (2004) 5045–5057.

[30] M.S. EL-Genk, H. Bostanci, Saturation boiling of HFE-7100 from a copper surface, simulating a microelectronic chip, International Journal of Heat and Mass Transfer 46 (2003) 1841–1854.

[31] P.F. Alvarino, M.L.S. Simon, M. dos Santos Guzella, J.M.A. Paz, J.M.S. Jabardo, L.C. Gomez, Experimental investigation of the CHF of HFE-7100 under pool boiling conditions on differently roughened surfaces, International Journal of Heat and Mass Transfer 139 (2019) 269–279.

[32] Bar-Cohen, T.W. Simon, Wall Superheat Excursions in the Boiling incipience of Dielectric Fluids, Heat Transfer Engineering 9(3) (1988) 19-31.

[33] U. Sajjad, A. Sadeghianjahromi, H.M. Ali, C.C. Wang, Enhanced pool boiling of dielectric and highly wetting liquids - a review on enhancement mechanisms, International Communications in Heat and Mass Transfer 119 (2020) 104950.

[34] M.S. El-Genk, H.H. Saber, J.L. Parker, Efficient spreaders for cooling high-power computer chips, Applied Thermal Engineering 27 (2007) 1072–1088.

[35] G. Liang, I. Mudawar, Review of pool boiling enhancement by surface modification, International Journal of Heat and Mass Transfer 128 (2019) 892–933.

[36] C. Wu, W. Tong, B. B. Kanbur and F. Duan, "Full-scale Two-phase Liquid Immersion Cooing Data Center System in Tropical Environment," 2019 18th IEEE Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm), 2019, pp. 703-708.

[37] A.F. Ali, M.S. El-Genk, Spreaders for immersion nucleate boiling cooling of a computer chip with a central hot spot, Energy Conversion and Management 53 (2012) 259–267.

[38] J. L. Gess, S. H. Bhavnani, R. W. Johnson, Experimental Investigation of a Direct Liquid Immersion Cooled Prototype for High Performance Electronic Systems, IEEE

Transactions on Components, Packaging and Manufacturing Technology 5(10) (2015) 1451-1464.

[39] C. Liu, H. Yu, Evaluation and Optimization of a Two-Phase Liquid-Immersion Cooling System for Data Centers Energies 14 (2021) 1395.

[40] C. Liu, H. Yu, Experimental Investigations on Heat Transfer Characteristics of Direct Contact Liquid Cooling for CPU, Buildings 12 (2022) 913.

[41] B. Ramakrishnan, H. Alissa , I. Manousakis, R. Lankston., CPU Overclocking: A Performance Assessment of Air, Cold Plates, and Two-Phase Immersion Cooling, IEEE Transactions on Components, Packaging and Manufacturing Technology 11(10) (2021). 1703-1715.

[42] B.B. Kanbur, C. Wu, S. Fan, W. Tong , F. Duan, Two-phase liquid-immersion data center cooling system: Experimental performance and thermoeconomic analysis, International Journal of Refrigeration 118 (2020) 290–301.

[43] B.B. Kanbur, C. Wu, S. Fan, F. Duan, System-level experimental investigations of the direct immersion cooling data center units with thermodynamic and thermoeconomic assessments, Energy 217 (2021) 119373.

[44] N. Karwa, S.Y. Motta, Low-pressure heat transfer fluids for pumped two-phase cooling, 20th IEEE Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems (iTherm) 2021.

[45] B.B. Kanbur˙, C. Wu, F. Duan, Combined heat and power generation via hybrid data center cooling-polymer electrolyte membrane fuel cell system, International Journal of Energy Research 44 (2020) 4759–4772.

[46] B.B. Kanbur¨, C. Wu, F. Duan, Multi-criteria thermoeconomic and thermodynamic assessments of the desalination-integrated two-phase liquid-immersion data center cooling system, International Journal of Energy Research 44 (2020) 10453–10470.

[47] E. Oróa, P. Taddeob, J. Salom, Waste heat recovery from urban air cooled data centres to increase energy efficiency of district heating networks, Sustainable Cities and Society 45 (2019) 522–542.

[48] A.C. Kheirabadi, D. Groulx, Cooling of server electronics: A design review of existing technology, Applied Thermal Engineering 105 (2016) 622–638.

[49] D. B. Tuckerman, R.F.W. Pease, High Performance Heat Sinking for VLSI, IEEE Electron Device Letters EDL-2(5) (1981) 126-129.

[50] Z.Y. Guo, Z.X. Li , Size effect on single-phase channel flow and heat transfer at microscale, International Journal of Heat and Fluid Flow 24 (2003) 284–298.

[51] H. Herwig, O. Hausner, Critical view on ''new results in micro-fluid mechanics'':an example, International Journal of Heat and Mass Transfer 46 (2003) 935–937.

[52] P.S. Lee, S.V. Garimella, D. Liu D, Investigation of heat transfer in rectangular microchannels, International Journal of Heat and Mass Transfer 48 (2005) 1688–1704.

[53] V. Radmard, A.R. Gharaibeh, M.I. Tradat, C.H. Hoang, Y.Y. Manaserh, K. Nemati, S.N. Schiffres, B.G. Sammakia, Performance Analysis of Corrosion Resistant Electroless Nickel-Plated Impinging Computer Numerical Control Manufactured Liquid Cooling Cold Plate, ASME Journal of Electronic Packaging 145(2) (2023) 021001.

[54] M.E. Steinke, V. Kamath, Experimental Investigation of a Cost Effective Skived Microchannel Cold Plate for Mass Production, Proceedings of the ASME 2015 13th International Conference on Nanochannels, Microchannels, and Minichannels collocated with the ASME 2015 International Technical Conference and Exhibition on Packaging and Integration of Electronic and Photonic Microsystems. ASME 2015 13th International Conference on Nanochannels, Microchannels, and Minichannels. San Francisco, California, USA. July 6–9, 2015.

[55] A.M. Sahar, J. Wissink, M.M. Mahmoud, T.G. Karayiannis, M.S. Ashrul Ishak, Effect of hydraulic diameter and aspect ratio on single phase flow and heat transfer in a rectangular microchannel, Applied Thermal Engineering 115 (2017) 793–814.

[56] Y. Hadad, N. Fallahtafti, L. Choobineh, C.H. Hoang, V. Radmard, P.R. Chiarot, B. Sammakia*et.*, Performance Analysis and Shape Optimization of an Impingement Microchannel Cold Plate," in *IEEE Transactions on*

Components, Packaging and Manufacturing Technology 10(8) (2020) 1304-1319.

[57] S.I., Guggari, Analysis of Thermal Performance Metrics-Application to CPU Cooling in HPC Servers, IEEE Transactions on Components, Packaging and Manufacturing Technology 11(2) (2021) 222-232.

[58] X. Shi, S. Li, Y. Mu, B. Yin, Geometry parameters optimization for a microchannel heat sink with secondary flow channel, International Communications in Heat and Mass Transfer 104 (2019) 89–100.

[59] S. Di Bari, J.S. Cotton, A.J. Robinson, Numerical Simulation of Exhaust Gas Cooling in Channels with Periodic Elbows for Application in Compact Heat Recovery Systems, Journal of Physics: Conference Series 395 (2012) 01204.

[60] A.J. Robinson, J. Colenbrander, R. Kempers, R. Chen, Solid and Vapour Chamber Heat Spreaders: Which to Choose and Why, IEEE Transactions on Components and Packaging Technologies 8(9) (2018) pp. 1581-1592.

[61] J.W. Elliott, M.T. Lebon, A.J. Robinson, Optimising integrated heat spreaders with distributed heat transfer coefficients: A case study for CPU cooling, Case Studies in Thermal Engineering 38 (2022) 102354.

[62] J. Bucko, M. Mayer, W. Benzinger, W. Augustin, R. Dittmeyer, S. Scholl, Influence of Deposits by Crystallization Fouling in Microchannels on the Heat Transfer Performance of Micro Heat Exchangers, Proceedings of the ASME 2011 9th International Conference on Nanochannels, Microchannels, and Minichannels. ASME 2011 9th International Conference on Nanochannels, Microchannels, and Minichannels, Volume 2. Edmonton, Alberta, Canada. June 19–22, 2011. pp. 71-75.

[63] M. Awais, A.A. Bhuiyan, Recent advancements in impedance of fouling resistance and particulate depositions in heat exchangers, International Journal of Heat and Mass Transfer 141 (2019) 580–603.

[64] H. Christophe, J.P. Minier, G. Lefèvre, Towards a description of particulate fouling: from single particle deposition to clogging."Advances in Colloid and Interface Science 185 (2012) 34-76.

[65] Y. Shao, X. Ruan, S. Li, Mechanism for clogging of microchannels by small particles with liquid cohesion, AIChE Journal 67(7) (2019) e17288.

[66] M.A. Kadhim, Holistic Study of Thermal Management in Direct Liquid Cooled Data Centres: from the Chip to the Environment, PhD Thesis, The University of Leeds, 2018.

[67] J. Perry, Fouling in silicon microchannel designs used for IC chip cooling and its mitigation, PhD Thesis, Rochester Institute of Technology, 2007.

[68] M. Iyengar, M. David, P. Parida, V. Kamath, B. Kochuparambil, D. Graybill, M. Schultz, M. Gaynes, R. Simons, R. Schmidt, T. Chainer, Extreme energy efficiency using water cooled servers inside a chiller-less data center," 13th InterSociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems, 2012, pp. 137-149.

[69] M. P. David *et al.*, "Impact of operating conditions on a chiller-less data center test facility with liquid cooled servers," *13th InterSociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems*, 2012, pp. 562-573

[70] P.R. Parida, T.J. Chainer, M.D. Schultz, M.P. David, Cooling Energy Reduction During Dynamically Controlled Data Center Operation, Proceedings of the ASME 2013 International Technical Conference and Exhibition on Packaging and Integration of Electronic and Photonic Microsystems. Volume 2: Thermal Management; Data Centers and Energy Efficient Electronic Systems. Burlingame, California, USA. July 16–18, 2013.

[71] Yuksel, V. Mahaney, C. Marroquin, S. Tian, M. Hoffmeyer, M. Schultz, T. Takken, Thermal and Mechanical Design of the Fastest Supercomputer of the World in Cognitive Systems: IBM POWER AC 922, Proceedings of the ASME 2019 International Technical Conference and Exhibition on Packaging and Integration of Electronic and Photonic Microsystems. ASME 2019 International Technical Conference and Exhibition on Packaging and Integration of Electronic and Photonic Microsystems. Anaheim, California, USA. October 7–9, 2019.

[72] E. Druzhinin, A. Shmelev, A. Moskovsky, Y. Migal, V. Mironov and A. Semin, High temperature coolant demonstrated for a computational cluster, 2016 International Conference on High Performance Computing & Simulation (HPCS), 2016, pp. 814-817.

[73] H. Shoukourian, T. Wilde, H. Huber, Arndt Bode, Analysis of the efficiency characteristics of the first High-Temperature Direct Liquid Cooled Petascale supercomputer and its cooling infrastructure, J. Parallel Distrib. Comput. 107 (2017) 87–100.

[74] R. Lucchese, D. Varagnolo, A. Johansson, Controlled Direct Liquid Cooling of Data Servers, IEEE Transactions on Control Systems Technology 29(6)(2021) 2325-2338.

[75] M.A. Kadhim, N. Kapur, J.L. Summers, H. Thompson, Rack Level Study of Hybrid Liquid/Air Cooled Servers: The Impact of Flow Distribution and Pumping Configuration on Central Processing Units Temperature, Heat Transfer Engineering 41 (19-20) (2020) 1683-1698.

[76] S.J. Ovaska, R.E. Dragseth, S.A. Hanssen, Direct-to-chip liquid cooling for reducing power consumption in a subarctic supercomputer centre, International Journal of High Performance Computing and Networking 9(3) (2016) 242-249.

[77] T. Brunschwiler, B. Smith, E. Ruetsche and B. Michel, Toward zero-emission data centers through direct reuse of thermal energy, IBM Journal of Research and Development 53(3) (2009) 11:1-11:13.

[78] Y.J. Kim, M. Gonzalez, Exergy analysis of an ionic-liquid absorption refrigeration system utilizing waste-heat from datacenters, International Journal of Refrigeration 48 (2 0 1 4) 2 6 -3 7.

[79] Y.I. Aristov, Adsorptive conversion of ultralow-temperature heat: Thermodynamic issues, Energy 236 (2021) 121892.

[80] S. Zimmermann[1], I. Meijer, M.K. Tiwari, S. Paredes, B. Michel, D. Poulikakos, Aquasar: A hot water cooled data center with direct energy reuse, Energy 43 (2012) 237-245.

[81] S. Zimmermann[2], M.K. Tiwari, I. Meijer, S. Paredes, B. Michel, D. Poulikakos, Hot water cooled electronics: Exergy analysis and waste heat reuse feasibility, International Journal of Heat and Mass Transfer 55 (2012) 6391–6399.

[82] N. Meyer, M. Ries, S. Solbrig, T. Wettig, iDataCool: HPC with Hot-Water Cooling and Energy Reuse, Supercomputing (2013) 383–394.

[83] Auweter, H. Huber, Direct warm water cooled Linux Cluster Munich:CoolMUC, inSiDE 10 (2012) 26

www.nexalus.com

CONTACT US:

Nexalus Ltd,
Unit 13, South Bank,
Crosse's Green,
Cork, Ireland,
T12 XT71

Nexalus Labs,
Water St,
Loughanalla,
Castlepollard,
Co. Westmeath, Ireland,
N91 EP90

Nexalus Research,
Trinity Research & Innovation,
O'Reilly Institute, Trinity College,
Dublin 2, Ireland,

info@nexalus.com
www.nexalus.com

Registered in Ireland No. 628880,
VAT No. IE3573607SH

IN PARTNERSHIP WITH:

Munters · DANA · mergon · arteco · 8PACK APPROVED

WITH ENDORSEMENTS AND SUPPORT FROM:

Enterprise Ireland · Science Foundation Ireland · Trinity College Dublin · CONNECT · Irish Research Council

QUALITY CERTIFICATIONS:

IATF 16949 / Automotive Standard
ISO 13485 / Medical Device Standard
ISO 50001 / Energy Management Standard
ISO 45001 / Occupational Heath & Safety Standard
ISO 14001 / Environmental Standard

NEXALUS
COOLING THE CLOUD