The Merits of Open Bath Immersion Cooling of Datacom Equipment

Phillip E. Tuma
3M Company
3M Center, 236-2B-01
St. Paul, MN 55144-1000
petuma@mmm.com

Abstract

This paper discusses the economic and environmental merits of passive 2-phase immersion in semi-open baths of dielectric fluid for cooling datacom equipment such as servers. The technique eliminates the need for hermetic connectors, pressure vessels, seals and clamshells typically associated with immersion cooling and the connectors, plumping, pumps and cold plates associated with more traditional liquid cooling techniques. A board level power density of 11.7W/cm^2 can be sustained with 100cm^3 of fluid per kW. The modular 80kW baths modeled can eject 130kW per m^2 of floor space via water-cooled condensers. It is estimated that 28°C water at 15gpm could maintain average CPU junction temperatures, T_j<60°C and 62°C water at 30gpm could maintain T_j<85°C, maximizing the availability of the heat for other purposes. Alternatively, the heat can be transferred directly to ambient air without water as an intermediate. The costs and greenhouse gas emissions associated with conservative annual fluid emission estimates are found to be less than those associated with the electrical power required for traditional chassis fans and liquid pumps. Since these fugitive losses occur at one point, more efficient capture techniques can be easily applied.

Keywords

data center, datacenter, datacom, cooling, immersion, open bath, fluoroketone, evaporative bath, passive, 2-phase.

Nomenclature

a	constant
C	specific heat [J/kg-K]
h	heat transfer coefficient [W/m^2-K] or [W/cm^2-K]
k	thermal conductivity [W/m-K]
K	Henry's Law constant [Pa-mol/mol]
$\dot{m}$	mass flow rate [kg/s]
n	moles
ν	kinematic viscosity [cSt]
P	pressure [Pa]
Q	power or heat [W]
Q''	heat flux [W/cm^2]
R	ideal gas constant = 8.314 J/mol-K or thermal resistance [°C/W], [°C-cm^3/W]
ρ	density [kg/m^3]
T	temperature [°C] or [K]
V	volume [m^3]

Subscripts

a	ambient
air	air
atm	atmospheric
b	boiling or boiling point
c	chip
$cond$	condenser or condensation
f	fluid
H	headspace
i	initial or inlet
j	junction
o	final or outlet
s	sink
sat	fluid saturation
t	cold trap
v	vapor
w	water

1. Introduction

1.1. Air Cooling and its Limitations

Among the causes of inefficiency in traditional data center air cooling schemes are: 2nd Law irreversibility resulting from multiple heat transfer processes; mixing of warm and cool airstreams; power consumption of cooling hardware such as chillers, computer room air conditioners (CRACs), fans, blowers and pumps; and the reliance on air as a heat transfer media. The technologies being implemented target one or more of the aforementioned causes of inefficiency.

Water-cooled rear door heat exchangers [1], ducted rack exhaust plena and closed forced air racks, for example, limit the mixing of airstreams. Raising the facility air temperature can, in some instances, increase overall efficiency [2]. These and other technologies may enable a facility to operate without a chiller substituting cooling-tower-only operation with on-demand economizers when whether permits [3, 4] Facilities that use full-time economizers are simpler still and can achieve a Power Usage Effectiveness (PUE) <1.3 [5]. This is limited by economizer blowers, air filters [6] and fans within the servers that draw power, particularly when temperatures are high. Furthermore, facilities like these must be located in relatively cool climates [7].

There are other inherent economic and environmental impacts. Managing airflow at the chassis, rack or facility level adds significant engineering cost during the development of each new server, datacenter or datacenter expansion. The number of publications on the subject and the growth in companies supplying software, sensors, components and services aimed at refining air cooling are testament to its limitations. Air-cooled facilities are usually large buildings as required to accommodate airflow and evaporative cooling towers require a large volume of water [8]. The fabrication of fans, heat sinks, CRACs, economizers, filters, etc. requires refined natural resources [9, 10] that may end up in landfills.

Optimization of energy efficiency eventually leads beyond considerations of how best to eject a facility's waste heat to how best to utilize it. However, the feasibility and cost of recovering its waste heat at any distance from an air-cooled datacenter are limited by the heat's low thermodynamic availability or exergy and the large volumetric flow rate of air.

1.2. Traditional Liquid Cooling and its Limitations

Liquid cooling can reduce each of the aforementioned causes of inefficiency; facilitate waste heat recovery; and increase the thermodynamic availability of the heat removed [11]. Ellsworth and Iyengar [12] compared the facility level cooling performance of an air/liquid hybrid supercomputing cluster to its air-cooled equivalent with traditional chilled water and CRACs. They also predicted the efficiency gains achievable with an all-liquid technology and that same technology operating in a chiller-less or cooling-tower-only (water economizer) mode (Figure 1). The latter would result in a 90% reduction in cooling energy consumption versus the air-cooled cluster. While the efficiency of an all-liquid cluster was not shown to be dramatically higher than that of an air/liquid hybrid, the elimination of facility level air cooling infrastructure would likely reduce facility construction cost.

However, implementation of traditional liquid cooling schemes, be they single- or two-phase, indirect-contact or immersion, is complicated by the number and variety of heat generating devices on a server and the requirement that each server within a rack be "hot swappable," meaning that it can be removed and replaced without disturbing the operation of others. This makes it challenging to capture *all* of the heat generated on a printed circuit board (PCB) and move it to an external liquid stream. As a result, hybrid air-liquid systems bear inherent costs for design [13] and fabrication of cold plates, redundant pumps, plumbing, quick disconnects (QDs), controls, and heat exchangers [14] (Table 1). *All*-liquid systems (Table 2) are often more complicated, requiring additional or more complex cold plates, clamshells, or hermetic electrical connectors. Performance for many is limited by secondary even ternary thermal interfaces and fluid temperature glide. Hydrofluorocarbon (HFC) and perfluorocarbon (PFC) systems are prone to leakage and resultant global warming emissions at intractable sites like couplings.

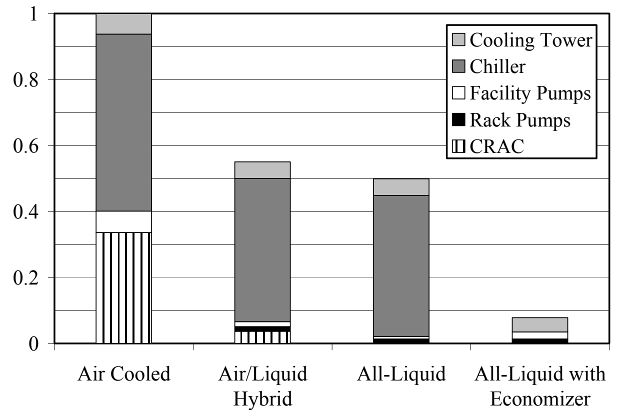

Figure 1: Normalized cooling power consumption of different components for various technologies [12]

#	Description
224	Machined/brazed copper cold plates, clamps and TIMs
168	Inter cold plate brazed copper manifold assemblies
28	Brazed copper header assemblies
28	Centrifugal blowers
14	EMI gasket shield assemblies
28	High performance quick disconnects (QDs)
Rack Level	
2	Large brazed copper manifold assemblies
1	Rear door heat exchanger
2	Plate heat exchangers
2	Reservoir tanks
2	Proportional metering valves
2	Pressure relief valves
2	Check valves
2	Solenoid isolation valves
12	Temperature sensors
6	Liquid level sensors
2	Magnetically-coupled centrifugal pumps
8	High performance large bore QDs
1	Pressure relief valves

TABLE 1: Partial list of cooling components in a 72kW air/liquid hybrid rack [14]

Technology	See List
1. Pump water or HFC refrigerant onto each node and through cold plates[12].	a,b,c,d,(e),(f),g
2. Cold plates in contact with, not connected to server/node. Transfer heat to interface via heat pipes etc[15, 16]	a,b,d2,(e)
3. Pump dielectric fluid onto each server/node enclosed in direct immersion cooling clamshell [17].	a,c,(e),(f),g

Shortcomings ()=depends on fluid technology

a. Cost and environmental impact of cold plates, pumps, etc.
b. Difficult to capture *all* heat without complex cold plate assemblies.
c. QD leakage risk.
d. Secondary thermal interface(s) (TIM2) impacts performance.
e. Fluid glide effects performance (single-phase)
f. Global warming emissions from fluid loss at intractable sites
g. Risk of leakage at clamshell, QD or electrical vias.

TABLE 2: Shortcomings of some commonly cited all-liquid cooling technologies

Clearly a simple, compact, inexpensive liquid cooling technique is needed that minimizes natural resource consumption and global warming emissions. It must capture all heat while minimizing the junction-to-water temperature difference. It should be modular, scalable and should easily accommodate evolving hardware [18].

1.3. Immersion Cooling History

Passive 2-phase (evaporative) immersion cooling has been used for decades to cool high value electronics like transformers, traction inverters (Figure 2), specialized computers and klystrons. This technology is still in use today, being favored for its simplicity, reliability, power density and performance.

These systems generally use sealed pressure vessels with hermetic electrical connections. They are evacuated and filled much like refrigeration systems and as such do not lend themselves to field service. It can be costly and complex to create such a hermetic enclosure for commodity computational electronics with their myriad of swappable components and electrical connections. For this reason, many people dismiss the idea of immersion in the context of datacom equipment. As will be shown, these measures are by no means necessary for static systems that remain at nearly constant temperature and power output.

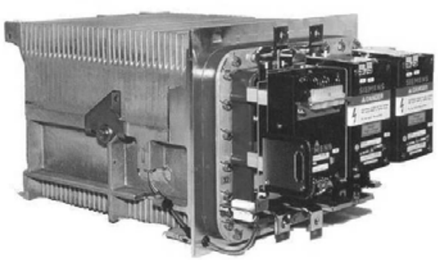

Figure 2: Immersion-cooled traction inverter from mining haul truck (photo courtesy of Siemens)

1.4. The Open Bath Vapor Degreaser

The management of fluid in vapor degreasers [19] is directly applicable to the concept discussed in this work. These ubiquitous machines are commonly used to clean parts ranging from screws and bearings to printed circuit boards, orthopedic implants and diesel engines. Degreasers comprise a tank, open at the top and fitted around its interior periphery with primary and secondary cooling coils (Figure 3). The tank is partitioned below a certain level to create two baths or "sumps" that are filled with a volatile solvent. The *boil sump* is heated from below causing the solvent within it to boil. Vapor rises to the height of the primary coils creating a saturated vapor zone beneath them. The condensed vapor flows back to the *rinse sump*, usually through a water separator. Because the sub-cooled solvent condensate is distilled, it is quite free of dissolved contaminants.

The sequence of cleaning steps is also shown in Figure 3. The part to be cleaned is placed in a wire basket and lowered into the saturated vapor zone (1). Vapor condenses on the part beginning the cleaning process. The part is next immersed in the rinse sump where ultrasonics may be used to displace particulate (2). It is next moved from the rinse sump back to the vapor where clean condensing vapor again rinses the part (3). After a brief pause above the saturated vapor zone, during which solvent remaining on the part quickly evaporates (4), it leaves the machine dry.

Vapor degreasers can clean thousands of parts per day with little fluid consumption despite the fact that they are open most of the time and only loosely lidded when not in use. Low loss rates are due in part to the secondary coils that often operate below 0°C. These set the solvent partial pressure that drives diffusion out of the bath through the freeboard region during use.

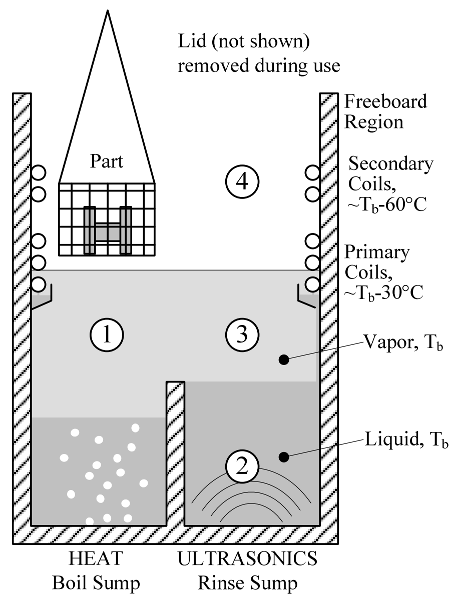

Figure 3: Schematic of an open bath vapor degreaser.

2. Open bath Immersion Cooling Concept

The concept discussed in this work is based on the premise that electronics can be immersion cooled in semi-open baths similar in many ways to a vapor degreaser (Figure 4). The term "semi" denotes a bath that is closed when access is not needed much like a chest-type food freezer. Like a freezer, it operates at atmospheric pressure and has no specialized hermetic connections for electrical inputs and outputs.

In this concept, each server or node plugs into a backplane in the bottom of a tank (versus the back of a rack) that is partially filled with a volatile dielectric working fluid. Electrical connections from the backplane enter a conduit beneath the liquid level and exit the top of the tank. A vapor condenser integrated into the tank is cooled by tower water or water used at some distance for comfort heating. If desired, the vapor can flow passively to an outdoor natural draft cooling tower to transfer its heat passively to outdoor air without water as an intermediate. This open bath concept has a multitude of advantages over the other liquid cooling schemes that have been proposed.

Principal among them (Table 3) is the fact that most of the aforementioned air and liquid cooling hardware are eliminated as are considerations relating to their integration, reliability and power consumption. Power density and efficiency are very high and fire protection is intrinsic to the technology. Of course, there are other considerations, such as fugitive fluid emissions. Since these occur at one point rather than at countless seals and junctions, they are easily quantified and mitigated with simple techniques as will be discussed.

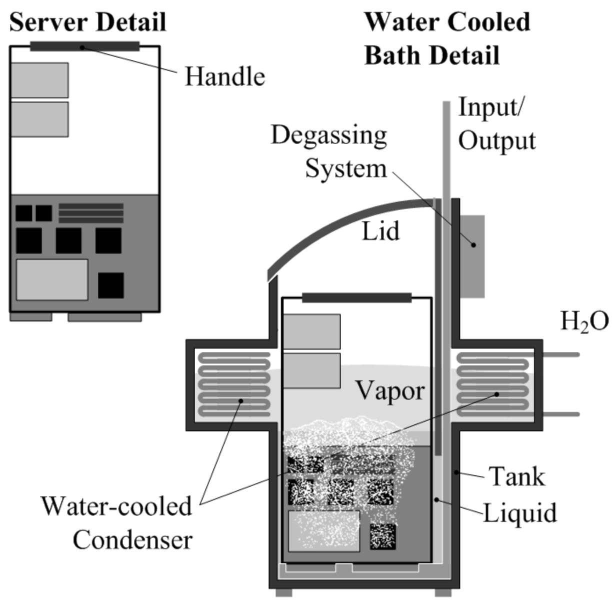

Figure 4: Water-cooled open bath immersion concept

Attribute	See list
1. No quick disconnects (QDs), clamshells, hermetic connectors.	a,c,d,(e)
2. No pumps, fans, economizers, compressors	b,c,d,e,h
3. No server/rack-level cold plates and plumbing	a,c,d
4. Fluid losses at one point	a, e
5. Intrinsic fire protection	h
6 ΔT_{jf} is low, no fluid glide	g
7. High power density	h
Advantage	
a. Less risk due to leakage	e. Reduced greenhouse gas emission
b. Reduced power consumption	f. Uniform device temperatures
c. Uses less natural resources	g. More efficient capture of heat
d. Reduced cost and complexity	h. Reduced facility construction cost

TABLE 3: Advantages of open bath immersion cooling compared with other liquid and hybrid techniques

3. Thermal Performance

The system thermal performance has two components. The first is at the board level and is device dependent. It is quantified by considering the junction-to-fluid temperature difference of the most difficult to cool component, the central processing unit (CPU). The second is the temperature difference from the fluid to the facility water. The fluid temperature, T_f, is the fluid's atmospheric boiling point,

$$T_f = T_b = T_{sat}(P_{atm})\,. \qquad (1)$$

3.1. Junction-to-Fluid Performance

The performance capabilities of passive 2-phase heat transfer with dielectric coolants are well documented [20]. Despite its simplicity and passive nature, it is not an inferior technology. Passive 2-phase immersion has been used to cool power semiconductors dissipating over Q_c''=1100W/cm^2 with $\Delta T_{j\text{-}f}$=45°C, a performance level competitive with the best emerging pumped water technologies [21]. A typical CPU package configuration with its integrated heat spreader (IHS) is almost ideally suited for passive 2-phase immersion cooling. In most cases, it requires only the addition of a 100µm thick porous metallic boiling enhancement coating (BEC). These coatings produce boiling heat transfer coefficients, H>10 W/cm^2-K at heat fluxes exceeding Q''=30W/cm^2. Incorporating this technology directly onto the IHS during package assembly eliminates the secondary thermal interface common to many liquid cooling schemes without altering the package assembly process.

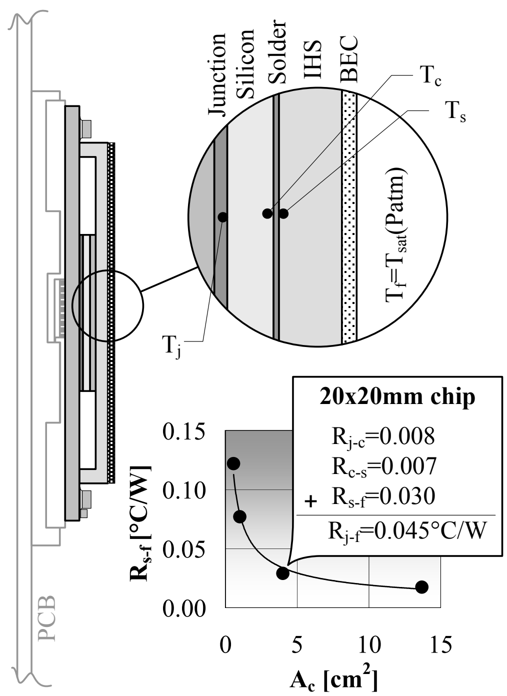

Figure 5: Sink-to-fluid performance as a function of die size and total junction-to-fluid performance for 20x20mm die in a hydrofluoroether (HFE) liquid

The resultant sink-to-fluid resistance, $R_{s\text{-}f}$, is dependent on the chip size (Figure 5). For a typical 20x20mm chip, $R_{s\text{-}f}$ =0.03°C/W. The additional resistances from sink-to-junction based on a 20x20mm thinned die and solder interface total 0.015°C/W[23]. With $R_{j\text{-}f}$=0.045°C/W, a 200W processor has an average junction-to-fluid temperature difference,

$$\Delta T_{j-f} = R_{j-f} Q_{CPU} \quad , \tag{2}$$

of about 9°C. It is likely this could be improved by optimizing a secondary heat transfer path through the package substrate.

It is worth noting that although this passive technique requires a first level interface (TIM1), the resultant chip-to-fluid thermal resistance is lower than that achievable with direct-die-contact spray or jet impingement schemes based on dielectric coolants. These *active* techniques achieve heat transfer coefficients of H=1-3W/cm^2-K [23,24] resulting in $R_{c\text{-}f}$ >0.09°C/W, for the 20x20mm die discussed earlier.

3.2. Fluid-to-Water Performance

Previous research [21] has shown that a high density, water-cooled condenser can achieve a volume-specific, fluid-to-water inlet resistance,

$$R_{f-w} = \frac{V_{cond}\left(T_f - T_w\right)}{Q_{cond}} \quad , \tag{3}$$

of 1.4ºC-cm^3/W under conditions of low water temperature glide. This number can be used in a log mean temperature difference (LMTD) analysis to predict condenser performance when the water temperature rises significantly:

$$LMTD = \frac{Q_{cond} R_{f-w}}{V_{cond}} \tag{4}$$

$$Glide = \Delta T_w = T_{w,i} - T_{w,o} = \frac{Q_{cond}}{\dot{m} C_w} \tag{5}$$

The resultant inlet water temperature, $T_{w,i}$, is

$$T_{w,i} = T_f - \frac{a}{1-a} \cdot glide \tag{6}$$

where

$$a \equiv exp\left(\frac{Glide}{LMTD}\right). \tag{7}$$

4. Power Density

By eliminating the heat sinks, fans and airspace normally required for air cooling, an immersion cooled server can, in principle, be quite compact. Current off-the-shelf air-cooled hardware, spread out as it is to facilitate air cooling, is not well suited for immersion. Determining how densely a server could be packaged is beyond the scope of this work. However, it is worthwhile exploring what power density could be cooled. Previous research of immersion cooling of power electronics [21] suggests that power densities are limited more by the electrical bus than by the capabilities of immersion. The authors estimated that 100cc of fluid are required to cool a 1kW module if density limitations are exercised.

Additional experiments were conducted with a form factor more consistent with a high density server. The simulated printed circuit board (PCB) shown in Figure 6 holds 20 heater assemblies comprised of 19x19mm 200W ceramic heaters epoxy bonded on one side to 30x30x3mm copper heat spreaders enhanced on the opposite side with a boiling enhancement coating (BEC). A thermocouple in the fluid, T_f, and one within each heat spreader, T_s, permit calculation of the individual thermal resistances and ensure that the heaters have not passed into the film boiling regime.

The simulated PCB was immersed in a confined vertical channel of the same area as the board with 4 and 7mm gaps between the boiling surface and the adjacent wall. This assembly was able to dissipate 4kW (200W per heater assembly) for a 4mm gap at atmospheric pressure when the bath was filled with $C_3F_7OCH_3$, a hydrofluoroether working fluid. Incipience overshoot did not exceed 7°C for any heater assembly. The average $R_{s\text{-}f}$ are shown in Figure 6, bracketed by ½ standard deviation on each side. 4kW equates to a PCB level heat flux of Q''=11.7W/cm^2 versus 1.7W/cm^2 for the Cray X1E spray-cooled supercomputer [24]. These data suggest that 1kW/cm of bath normal to the PCB and 100cc of fluid per kW are certainly attainable, if only from a thermal point of view. Also, the potential for reduction of materials and waste associated with PCB manufacture is significant.

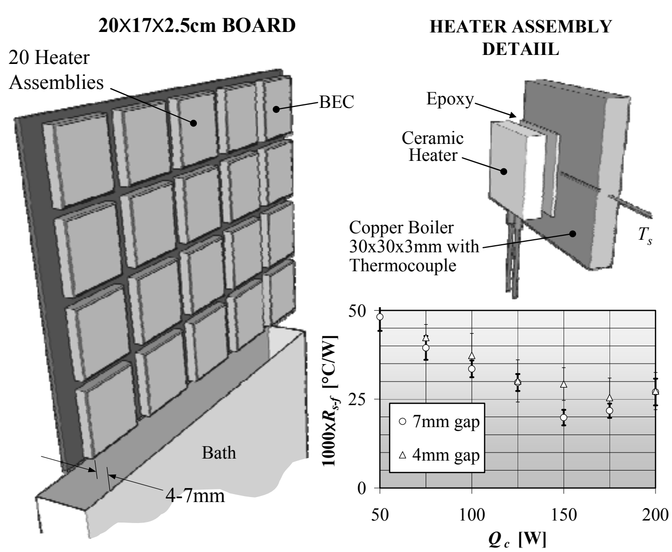

Figure 6: Experimental apparatus to demonstrate power density capabilities of immersion in a vertical PCB orientation

DOCKET ALARM

Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts

Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research

With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips

Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS

Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS

Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS

Sync your system to PACER to automate legal marketing.