

3D Numerical Analysis of Two-Phase Immersion Cooling for Electronic Components

Xudong An*, Manish Arora, Wei Huang, William C. Brantley, Joseph L. Greathouse
AMD Research
Advanced Micro Devices, Inc.
*Email: xudong.an@amd.com

ABSTRACT

This paper presents a three-dimensional numerical analysis, using ANSYS Fluent, of a two-phase immersion cooling solution for high-powered processor designs. The primary electronic component, a CPU package, is modeled as a bare $5\text{cm} \times 5\text{cm}$ flat plate heat source. The remainder of the model is based on the structure of typical two-phase immersion-cooled servers. Two arrangements are investigated, both fully immersing the heat-producing components in liquid coolant. The first has two vertically mounted heat sources to achieve higher packing density of the server, while the second only has a single heat source. This study considers 3M Novec7000 as the phase change coolant, which is a dielectric with a low (34°C) boiling temperature. We validate our numerical model against published results.

Our simulations show: (1) when the two heat sources are in series, the upper source runs significantly hotter than the lower source because bubbles from lower source reduce the coolant contact area on the upper source. (2) Novec7000 can support cooling a $5\text{cm} \times 5\text{cm}$ heat source in a vertical orientation with power as high as 225W (heat flux $9\text{W}/\text{cm}^2$). However, if two such sockets are thermally coupled, the power of the upper socket must be lower than 185W . If socket power exceeds that limit, a heat transfer enhancement layer should be applied to the coupled sockets to increase cooling area and reduce heat flux.

KEY WORDS: Two-phase, Immersion cooling, Supercomputer, CPU Cooling, CFD

INTRODUCTION

The power density of high-performance computing systems and supercomputers has been rapidly increasing in the past decades, a process that is expected to continue in the foreseeable future. By the early 2020s, supercomputers are expected to approach exascale level of performance (10^{18} floating point operations per second). Removing dissipated heat of such large systems will be a critical design constraint. The highest-powered components in these systems, such as CPUs and GPUs, must be prevented from exceeding temperature limits to avoid component failure. In some situations, temperatures should be kept even lower so that the components can operate more efficiently. These constraints imply increasing demands from heat dissipation mechanisms.

Two-phase immersion cooling is a recent approach for handling such demanding electronic cooling needs. In this approach, components are submerged into a dielectric liquid. Heat then transfers from components into the surrounding liquid, causing the liquid to vaporize. The high latent heat of vaporization in liquid helps rapidly extract heat.

Two-phase immersion cooling has been studied in various aspects and applications. Wada, et al. [1] discussed the feasibility of phase change in cooling electronic devices such as microwave transmitters. Campbell and Tuma [2] demonstrated that two-phase immersion cooling offers more than 10% better cooling performance than cold plate liquid cooling. Wagner et al. [3] indicated that two-phase immersion cooling can effectively cool computing systems with higher packing density compared to air, cold-plate liquid, or single-phase mineral oil cooling solutions. This is because two-phase immersion cooling does not require complicated and bulky cooling devices like heatsinks or water plates, which also indicates that systems using two-phase immersion cooling can dedicate more space to electronics instead of cooling components. Based on experimental results, Parker and El-Genk [4] and Amir and El-Genk [5] studied a set of parameters that may impact two-phase immersion cooling, such as type of liquid, material of boiling surface, and ambient temperature.

Due to the presence of vaporization and condensation in two-phase immersion cooling systems, there are special requirements in assembling the computing rack. Tuma [6] provided rack and data center solutions so that there is neither significant vapor-phase coolant loss nor over-pressure issues in the rack. 3M has presented information about deployed data centers using two-phase immersion cooling [7]. Results show that when using two-phase immersion cooling, computing power per rack can be as high as 250kW and the data center power usage effectiveness (PUE) can be as low as 1.02.

The objective of this paper is to investigate configurations for passive two-phase immersion cooling of supercomputing systems. We develop a model of a computing rack in ANSYS Fluent [8], and sweep the power ranges of high-power components, such as CPUs and GPUs, from 50W to 300W . In order to study the thermal impact among processors in the rack, we investigate two board arrangements: the first has two vertically mounted heat sources, while the second has only one heat source. Similar arrangements have been considered by Matsuoka et al. [9], though that study focused on single-phase mineral oil cooling. The boiling liquid in two-phase cooling causes turbulence, so the upper and lower sockets in our study see significantly different cooling situations than those studied in [9]. Our studies model 3M Novec7000 [10], a dielectric liquid, as the coolant because of its low boiling point (34°C) and high latent heat of vaporization ($142\text{kJ}/\text{kg}$). Both ambient and liquid temperature are kept at 34°C , hence there is no sub-cooling in our simulation configuration. Finally, we collect thermal maps and boiling behaviors in different system setups, and analyze this data to find the maximum power allowed on each socket.

MODEL CONFIGURATION

For our studies, we developed our 3D numerical models in the ANSYS Fluent CFD simulator. We considered two arrangements, as shown in Figure 1(a) and 1(b). In each case, four computing boards, labelled 1-4 from left to right, are combined in a single immersion cooling tank. In Figure 1(a), two sockets are arranged vertically on each computing board, labelled A and B for the lower and upper sockets, respectively. We refer to this as a two-socket arrangement. In Figure 1(b), a single socket is installed on each computing board. We refer to this as a single-socket arrangement. Sockets are indexed by the board they belong to and their locations on that board. For example, in the two-socket arrangement, the lower socket on second computing board is socket 2A; in the single-socket arrangement, the only socket on second computing board is socket 2.

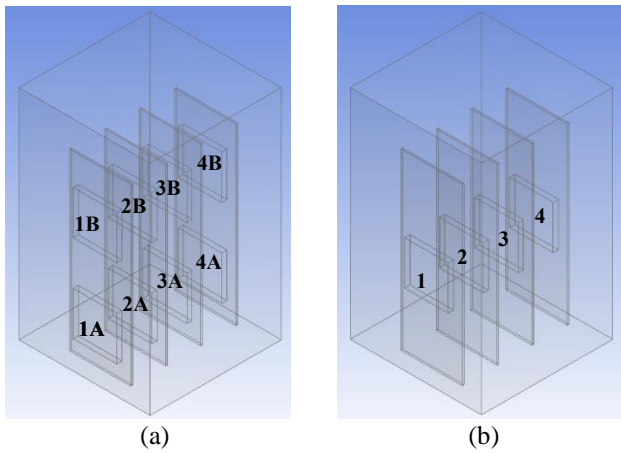


Fig. 1. 3D view of two-phase immersion cooling tank in ANSYS Fluent (a) two sockets per PCB (b) one socket per PCB

We model the immersion cooling tank as a cuboid of 15cm width (W) \times 15cm length (L) \times 25cm height (H) that is filled with Novec7000 coolant liquid. The boundary conditions and dimensions for the cooling tank are shown in Figure 2. The flow outlet is at the top of the tank, where the vaporized Novec7000 rises and escapes. The size of each computing board is 7cm width (W_b) \times 20cm height (H_b), and the size of each socket is 5cm \times 5cm. Square heat sources are arranged on each socket to mimic the processor power. No heatsink or cooling enhancement are attached to the socket, but each socket is fully immersed in the Novec7000 liquid. The power on all sockets is assumed to be the same, except for the sockets on board 1 which are perfectly power gated (0W). We consider a range of processor power consumption, from 50W to 300W per socket. To leave enough space for boiling on the socket lid, we set a gap (L_{gap}) of 4cm between each board.

In ANSYS Fluent 18.0, we used Eulerian models to solve multiphase flow and non-equilibrium boiling equations for boiling phenomenon. In turbulent settings, we applied standard k -epsilon equations. For solution methods, we set a phase coupled SIMPLE scheme. We used a transient model with a

time interval of 0.05s for surface boiling, since surface boiling is a transient flow. Finally, we collected results starting at 5 seconds into the simulation, when flow is in a steady state.




Fig. 2. Model configurations and boundary conditions of two-phase immersion cooling tank

Model validation. Some experimental data of a similar cooling setup was published by Parker and El-Genk [4]. In those experiments, a 1cm \times 1cm heat source is oriented vertically and immersed in FC72 liquid, which is also frequently used for two-phase boiling studies. The boiling point of FC72 is higher than Novec7000, and its latent heat of vaporization is lower. The main properties of FC72 [11] and Novec7000 [10] are listed in Table 1.

Table 1. Properties of 3M FC72 and Novec7000 fluid

	FC72	Novec7000
Molecular Weight (g/mol)	338	200
Boiling point @ 1atm ($^{\circ}$ C)	56	34
Liquid density (kg/m^3)	1680	1400
Kinematic viscosity (cSt)	0.38	0.32
Latent heat of vaporization (kJ/kg)	88	142
Specific heat ($\text{J}\cdot\text{Kg}^{-1}\text{K}^{-1}$)	1100	1300
Thermal conductivity ($\text{W}\cdot\text{m}^{-1}\text{K}^{-1}$)	0.057	0.075
Critical temperature ($^{\circ}$ C)	176	165
Critical pressure (MPa)	1.83	2.48

The effect of heat flux and temperature on the heat source were investigated by Parker and El-Genk [4]. To validate our numerical model, we compared critical heat flux (CHF) and corresponding lid temperature between our models and the previous publication's experiments. Figure 3 shows the curve obtained from Parker and El-Genk [4] (shown with the \times symbol) as well as the results from our simulations. On the x -axis, ΔT_{sat} is the temperature difference between T_{avg_lid} of the heat source and T_{sat} of the FC72 liquid. T_{avg_lid} is the average

temperature on lid, and T_{sat} is the boiling point of a liquid. The range of lid temperature in experiment and simulation is close, when heat flux is near critical point. The experiment results indicate that the CHF in current condition is $15.7\text{W}/\text{cm}^2$ and the lid temperature is 25.2°C higher than boiling point of FC72. In our model, when ΔT_{sat} is 25.9°C , heat flux is $14\text{W}/\text{cm}^2$ and it is reaching critical point. Therefore, discrepancy of CHF is smaller than 10% and that of corresponding temperature is smaller than 3%. Given the nature of boiling complexity, uncertainties on such scale are acceptable.

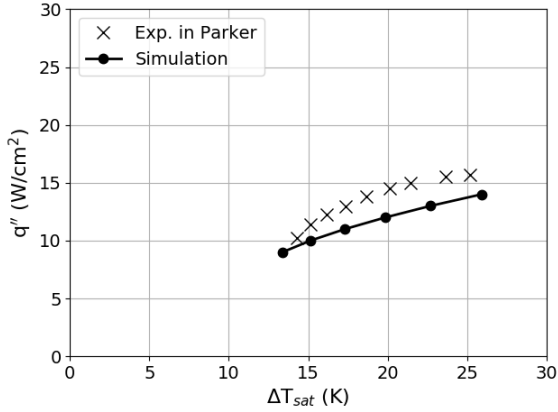


Fig. 3. Validation: comparison of our ANSYS Fluent model to the data described by Parker and El-Genk [4]

RESULTS AND DISCUSSION

In this paper, we study two-phase immersion cooling configurations for a square heat source using Novec7000 as our immersion liquid. The ambient temperature and temperature of the liquid are both 34°C , which is the boiling point of Novec7000. Therefore, with this configuration, there is no sub-cooling in our simulations. The ambient pressure of the entire system is 1atm. A variety of power values are considered, ranging from 50W to 300W per socket. A few parameters are analyzed, including rack arrangement, socket location, socket thermal map and maximum allowed power per socket.

Temperature vs. socket location

Our first set of simulations investigate the impact of socket location on system temperature. In these studies, we set the power of all sockets on boards 2-4 to 125W, while there is no power on computing board 1. Figure 4(a) shows that sockets on the same vertical level (socket 2B, 3B, and 4B) have almost the same temperature distribution. It is also worth noting that the thermal maps of lower sockets, 2A, 3A and 4A in the two-socket arrangement are very similar to the ones of sockets 2, 3 and 4 in the single-socket arrangement.

As this data shows, the upper socket in a two-socket arrangement will become much hotter than the lower socket. The temperature gradient across the lower socket is not remarkable. Additionally, the lower socket thermal characteristics are quite similar to the single-socket scenario, indicating that thermal coupling from upper socket to lower socket is minimal. In contrast, the upper sockets have large hotspots in their top-central region. These issues arise because

vapor generated from the lower socket rises and prevents sufficient liquid contact with the upper socket. Especially in the top-central region, the amount of liquid on the lid is greatly reduced.

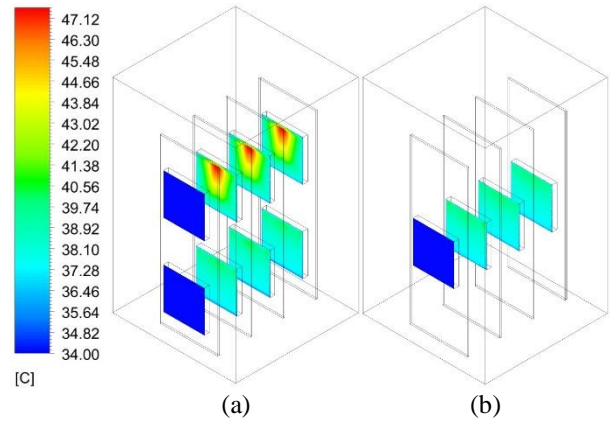


Fig. 4. Processor lid thermal maps when socket power is 125W (a) two sockets per PCB (b) one socket per PCB (No sub-cooling, $T_{\text{sat}} = 34^\circ\text{C}$, ambient pressure is 1atm)

Quantitative values from this study, including the average and maximum temperature on the lid of each socket, are shown in Figure 5. The average temperature on the lids of sockets 2A, 3A, 4A and socket 2, 3, 4 are all close to 38.3°C (blue cross and blue circle), and their maximum temperature are between $40.9^\circ\text{C} - 41.3^\circ\text{C}$ (red cross and red circle). For the upper sockets, 2B, 3B and 4B, the average lid temperatures are all about 40.8°C (blue square) and the maximum temperatures are between $47.3^\circ\text{C} - 47.6^\circ\text{C}$ (red square).

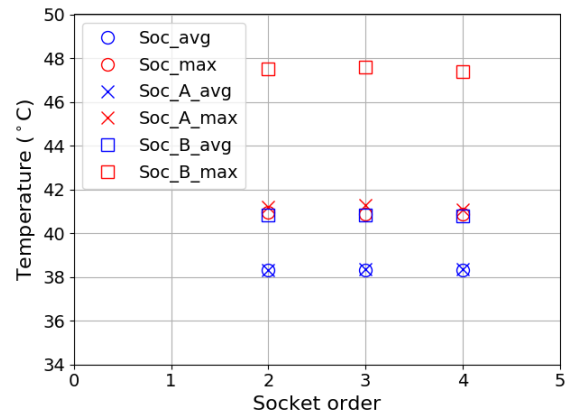


Fig. 5. Maximum temperature and average temperature on each socket when socket power is 125W

Figure 6 shows the volume fraction of vapor at the mid-plane inside the cooling tank. The dark blue color represents liquid, while red represents vapor. Vapor is generated on the surface of sockets and rises to the top of the tank. This figure shows that vapor created at the lower sockets rises to cover the upper sockets, and the vapor density at the upper sockets is thus much higher than that on lower sockets and those on the single-

socket model. Figure 6 also shows that the boiling phenomenon on boards 2, 3, and 4 are almost the same. Because of the similarity between these results, we focus the analysis on board 2.

Temperature vs. socket power

To investigate the relation of lid temperature and socket power in two-phase immersion cooling, we studied a wide range of socket powers between 50W and 300W. We kept the power of the sockets on boards 2, 3 and 4 identical while using no power on board 1. As shown, the temperature distribution on computing board 2, 3 and 4 are very close, and so we focus our analysis on board 2.




Fig. 6. Volume fraction of vapor on the middle plane of the tank when socket power is 125W for (a) the two-socket arrangement (b) the single-socket arrangement

The thermal maps for the socket lids as power is varied from 50W to 300W are shown in Figure 7 (note that the temperature ranges are not the same for each subfigure). There are a few common features among the thermal maps.

First, there is significant hotspot on the upper socket of each thermal map. When power increases, temperature gradient on the lid of the upper socket becomes larger. For example, in Figure 7(a), the average temperature is 36.8°C and the maximum is 0.5°C hotter. In Figure 7(b), the average temperature at the same location is 38.9°C and the maximum is 3.3°C higher. In the extreme, Figure 7(f) shows an average temperature of 67.9°C and a hotspot temperature of 123.2°C, which the difference rises to 55.3°C. This data shows that the processors (which generally must operate below 105°C) could not safely run within the upper socket if both cores were burning 300W.

Second, our data shows that when power increases, the temperature gradient on the lower socket lid is no longer trivial. For example, when the power on the lower socket is 50W, its average lid temperature is 36.6°C and its maximum temperature is 36.8°C. When we increase the power to 300W, the average temperature is 52.4°C, but maximum temperature reaches 85.8°C. This is because higher power socket generates more vapor, which prevents the upper part of the chip from contacting enough liquid coolant.

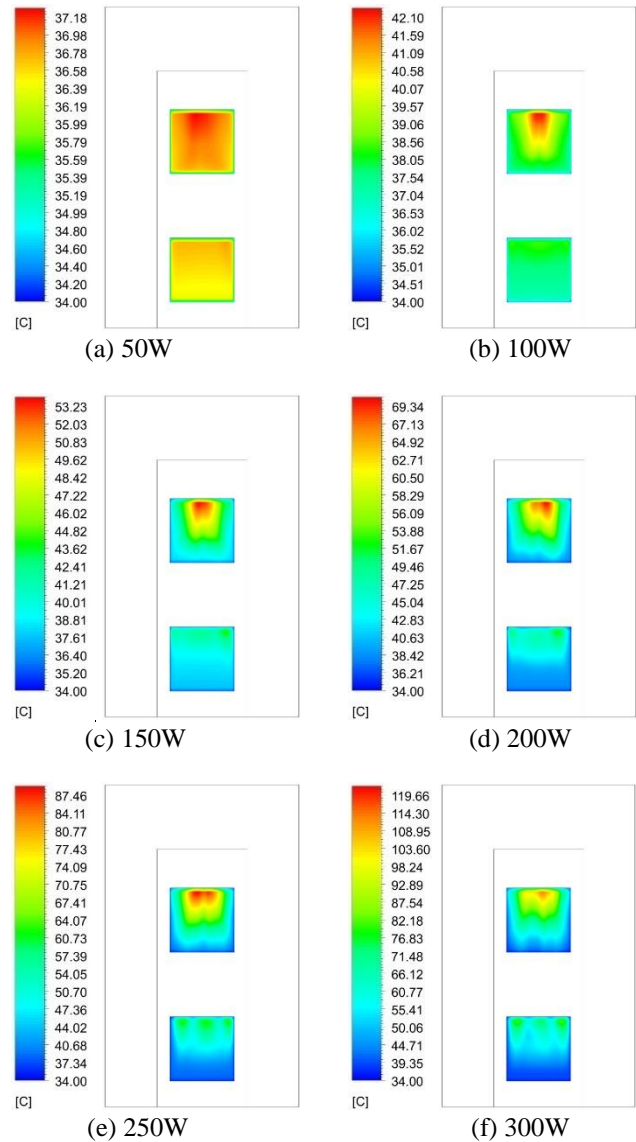


Fig. 7. Thermal map on lid of sockets on 2nd computing board when socket power varies from 50W to 300W (No sub-cooling, $T_{\text{sat}} = 34^\circ\text{C}$, ambient pressure is 1atm)

Figure 8 shows the relation of power and maximum lid temperature on both the upper and lower sockets of board 2. We also show the values for the single-socket arrangement for reference. The temperature curve for socket 2 in the single-socket arrangement almost completely overlaps the results for socket 2B in the two-socket arrangement. As seen, the maximum temperature on the upper socket rises much faster due to the vapor-induced hotspot. This data implies that the upper socket should be assigned lower power, or attached under a boiling enhancement layer to keep all processors running within a thermally safe range. Additionally, temperature rises super-linearly with power increases in this phase-change immersion cooling. This is because of loss of contact between the liquid and cooling surface because of vapor. On the other hand, conventional pure liquid cooling and pure air cooling

usually have a linear relationship between temperature rise and power, i.e. constant convection thermal resistance.

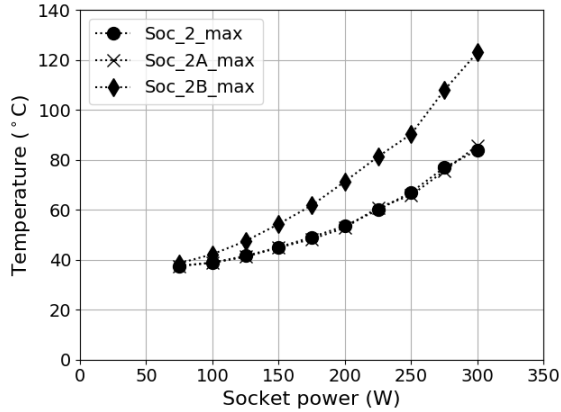


Fig. 8. Relation of socket power and lid temperature

Power limitation analysis

Using results from a conjunction of our numerical model and detailed processor die and floorplan analysis from HotSpot [12], we developed a methodology to estimate the maximum on-chip silicon temperature. This analysis of on-chip temperature is plotted in Figure 9, with Novec7000 liquid and 34°C ambient. Considering that the maximum allowed on-chip temperature is 105°C, this data shows that the maximum allowed socket power is 225W in our single-socket server arrangement. When adding another socket above the original one, the maximum power for each is 40W lower.

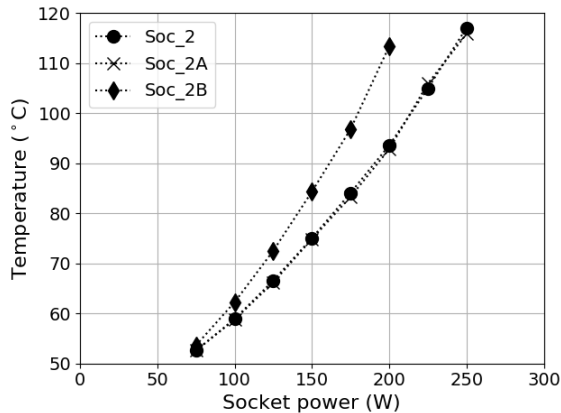


Fig. 9. Relation of socket power and maximum on chip temperature

CONCLUSIONS

This work presents a thermal analysis for various configurations of two-phase immersion cooled electronics components. A numerical model was developed in the ANSYS Fluent CFD simulation tool and validated against existing experiment data. Two models of cooling tanks and computational board arrangements were compared and analyzed. From our simulation results, we found that when

there are two sockets arranged vertically, significant hotspots appear on the upper socket due to the tremendous thermal coupling impact; in addition, this coupling causes significant on-chip temperature gradients for both sockets. This is mainly caused by the vapor bubbles in a two-socket arrangement. Moreover, even the lower socket (or the single socket on each board in a one-socket arrangement) can have significant temperature gradients across the lid.

We used Novec7000 as the phase-change coolant in this study, and we studied a typical socket size of 5cm × 5cm. No cooling enhancements are attached, so boiling happens on the bare socket. Under such conditions, the maximum allowed power for one socket with no thermal impact from others is 225W. When two sockets are arranged vertically, to ensure both sockets safely operate under their on-chip temperature constraints, the power for each should be no more than 185W. If exceeding the power limitation, special enhancement layer, which incurs higher cost, should be applied for the sockets, such as simple metal pin fin plate to enlarge the boiling area, or special coating to increase surface boiling.

ACKNOWLEDGMENT

AMD, the AMD Arrow logo, and combinations thereof are trademarks of Advanced Micro Devices, Inc. Other product names used in this publication are for identification purposes only and may be trademarks of their respective companies. © 2018 Advanced Micro Devices, Inc. All rights reserved.

References

- [1] Mizuki Wada, Arihiro Matsunaga, Mahiro Hachiya, Masaki Chiba, Kunihiro Ishihara, and Minoru Yoshikawa. "Feasibility study of two-phase immersion cooling in closed electronic device." *Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm)*, 2017 16th IEEE Intersociety Conference on. IEEE, 2017.
- [2] Levi Campbell and Phillip Tuma. "Numerical prediction of the junction-to-fluid thermal resistance of a 2-phase immersion-cooled IBM dual core POWER6 processor." *Semiconductor Thermal Measurement and Management Symposium (SEMI-THERM)*, 2012 28th Annual IEEE. IEEE, 2012.
- [3] Guy R. Wagner, Joseph R. Schaadt, Justin Dixon, Gary Chan, William Maltz, Kamal Mostafavi, and David Copeland. "Test results from the comparison of three liquid cooling methods for high-power processors." *Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm)*, 15th IEEE Intersociety Conference on. IEEE, 2016.
- [4] Jack L. Parker, and Mohamed S. El-Genk. "Effect of surface orientation on nucleate boiling of FC-72 on porous graphite." *Journal of heat transfer* 128.11: 1159-1175, 2006.
- [5] Amir F. Ali, and Mohamed S. El-Genk. "Numerical analysis of spreaders with an enhancing nucleate boiling surface for immersion cooling of chips with central hot spots." *Thermal and Thermomechanical Phenomena in*

Electronic Systems (ITherm), 13th IEEE Intersociety Conference on. IEEE, 2012.

- [6] Phillip E. Tuma. "The merits of open bath immersion cooling of datacom equipment." Semiconductor Thermal Measurement and Management Symposium, 2010. SEMI-THERM. 26th Annual IEEE. IEEE, 2010
- [7] 3M, "Two-phase immersion a revolution in data center efficiency", whitebook, 2015
- [8] ANSYS Fluent, "ANSYS Fluent theory guide 18.0." Ansys Inc., 2017
- [9] Morito Matsuoka, Kazuhiro Matsuda, and Hideo Kubo. "Liquid immersion cooling technology with natural convection in data center." Cloud Networking (CloudNet), 2017 IEEE 6th International Conference on. IEEE, 2017.
- [10] 3M, "3M Novec7000 Engineered Fluid", 2014
- [11] 3M, "3M Electronic Liquid FC-72", 2014
- [12] Wei Huang, Shougata Ghosh, Sivakumar Velusamy, Karthik Sankaranarayanan, Kevin Skadron, and Mircea R. Stan. "HotSpot: A compact thermal modeling methodology for early-stage VLSI design." IEEE Transactions on Very Large Scale Integration (VLSI) Systems 14.5: 501-513, 2006