

Two-Phase Vapor Chambers with Micropillar Evaporators: A New Approach to Remove Heat from Future High-Performance Chips

Zihao Yuan¹, Geoffrey Vaartstra², Prachi Shukla¹, Mostafa Said³, Sherief Reda³, Evelyn Wang², and Ayse K. Coskun¹

¹Boston University, Boston, MA - {yuan1z, prachis, acoskun}@bu.edu

²Massachusetts Institute of Technology, Cambridge, MA - {gvaartst, enwang}@mit.edu

³Brown University, Providence, RI - {mostafa_said, sherief_reda}@brown.edu

ABSTRACT

High power densities lead to thermal hot spots in modern processors. These power densities are expected to reach kW/cm^2 scale in future high-performance chips and this increase may significantly degrade performance and reliability, if not handled efficiently. Using two-phase vapor chambers (VCs) with micropillar wick evaporators is an emerging technique that removes heat through the evaporation process of a coolant and has the potential to remove high heat fluxes. In this cooling system, the coolant is supplied passively to the micropillar wick via capillary pumping, eliminating the need for an external pump and ensuring stable thin-film flow. Evaluation of such an emerging cooling technique on realistic chip power densities and micropillar geometries necessitates accurate and fast thermal models. Although multi-physics simulators based on either finite-element or finite-volume methods are highly accurate, they have long design and simulation times. This paper introduces a novel compact thermal model capable of simulating two-phase vapor chambers with micropillar wick evaporators. In comparison to COMSOL, our model shows a competitively low error of 1.25°C and a $214\times$ speedup. We also present a comparison of the cooling performance of different cooling techniques such as a conventional heat sink, liquid cooling via microchannels, hybrid cooling using thermoelectric coolers and liquid cooling via microchannels, and two-phase VCs with micropillar wick evaporators for the first time. Based on our observations, two-phase VCs and microchannel-based two-phase cooling show better cooling performance for hot spot power densities of less than $1500 \text{ W}/\text{cm}^2$, while hybrid cooling achieves lower hot spot temperature and thermal gradients for hot spot power densities between 1500 and $2000 \text{ W}/\text{cm}^2$.

INTRODUCTION

High power densities have already become major challenges for modern processors and they are expected to surpass $1 \text{ kW}/\text{cm}^2$ in future systems [1]. These power densities cause localized on-chip hot spots, which in turn result in performance and reliability degradation. Existing cooling solutions (e.g., forced air cooling via fans) are not able to solve these high temperature issues efficiently. Therefore, cooling is becoming more crucial for future processors.

Several emerging cooling techniques such as phase change materials (PCMs) [2], liquid cooling via microchannels [3], thermoelectric coolers (TEC) [4], two-phase cooling [5], [6], and hybrid cooling [7] have been explored by researchers to mitigate thermal problems. Among these new cooling




Fig. 1: Vapor chamber structure view.

techniques, two-phase cooling with vapor chambers (VCs) is a good candidate for future processor cooling owing to its advantages of better cooling performance¹ and no pumping power² (in contrast to liquid cooling via microchannels and microchannel-based two-phase cooling) [5]. The schematic of a VC is shown in Figure 1. In this technique, a capillary-driven flow conducts thin-film evaporation through a porous wick placed on the bottom surface of the VC [5]. We specifically focus on a porous wick that consists of a micropillar array with a periodic and precisely defined geometry. The selection of the micropillar array structure has a significant impact on the overall cooling performance of the evaporators as well as the VCs.

Fast and accurate thermal models for cooling techniques are essential for early-stage design exploration and optimization. There are a number of compact thermal models (CTMs) for various cooling techniques, such as heat sinks [8], liquid cooling via microchannels [9], [10], [11], hybrid cooling (TEC and liquid cooling via microchannels) [10], and microchannel-based two-phase cooling [12]. However, there is no existing CTM for two-phase VCs that can facilitate experiments with any processor and coolant of interests. Two-phase VCs simulations are typically carried out using computational fluid dynamics (CFD) modules with a finite-element solver in COMSOL or a finite-volume solver in ANSYS [13], [14]. Nevertheless, such tools are time-consuming and have large memory requirements; as a result, the simulation time of a realistic chip can take days. Therefore, these CFD tools are not suitable for co-design and optimization of realistic processor

¹Higher/better cooling performance refers to achieving lower maximum temperatures and thermal gradients.

²While there is no pumping power on the evaporator side, VCs may still need cooling power to reject the heat using a fan or other cooling methods.

architectures and cooling techniques.

In this paper, we present a CTM of two-phase VCs with micropillar wick evaporators to enable fast and accurate steady-state thermal analysis. This CTM is modular and can be applied to various chip designs and power profiles. The contributions of this paper are as follows:

- We introduce a steady-state CTM of two-phase VCs with micropillar wick evaporators and integrate it in a commonly used CTM simulator.
- We design a COMSOL model of two-phase VCs with micropillar wick evaporators and validate the proposed CTM against this COMSOL model. Our CTM achieves a maximum error of 1.25°C with a speedup of 214x.
- To analyze the cooling performance of two-phase VCs with micropillar wick evaporators and select the best possible cooling solution with respect to various hot spot power densities, we compare the cooling performance of different cooling techniques including a conventional heat sink, liquid cooling via microchannels, hybrid cooling (TEC and liquid cooling via microchannels), microchannel-based two-phase cooling, and two-phase VCs with micropillar wick evaporators. We observe that the two-phase cooling techniques including two-phase VCs and microchannel-based two-phase cooling provide the best cooling performance when the hot spot power density is less than 1500 W/cm². For hot spot power densities between 1500 and 2000 W/cm², hybrid cooling shows superior cooling performance.

THERMAL MODELING OF TWO-PHASE VCS

In this section, we first review the structure and the operating mechanism of the VCs. Next, we present our COMSOL model and CTM for two-phase VCs with micropillar wick evaporators.

Overview of The Vapor Chamber

As we can see in Figure 1, the evaporator, which consists of a porous wick, is placed directly on the heat source, which in this case is the chip power. The evaporator conducts thin-film evaporation and helps the coolant transform from liquid to vapor. On the top side of VC, a condenser chills down the heated vapor and condenses it back to liquid state. The condensed liquid is then driven by the wicking structures along the side walls of the VC back to the evaporator. There are two metrics that impact the cooling performance of the VCs: (i) heat transfer coefficient (HTC), and (ii) dry-out heat flux. HTC is a parameter that determines the rate of heat transfer per unit temperature difference of the evaporator. In other words, a higher HTC will lead to better cooling performance. Dry-out heat flux is the thermal limit of a two-phase device. If the hot spot power density of the chip is higher than the dry-out heat flux, the coolant will no longer remain in the two-phase state and possibly cause over-heating and damage to the chip. Therefore, a higher dry-out heat flux means more heat can be removed before the coolant dries out. Micropillar wick evaporators have been shown to improve the heat transfer rate and enhance the dry-out heat flux owing to their high capillary pumping budget and large evaporation area [15], [16]. In this

work, we focus on the VCs that use micropillar wicks as the evaporators.

Thermal modeling of two-phase VCs is complicated because of the fact that the coolant exists in two phases. This makes temperature a function of many parameters such as mass flux, pressure, saturation temperature of the coolant, or surface tension [5]. Most of the existing two-phase VCs thermal models are implemented using CFD tools (e.g., ANSYS) [17], [14]. These tools provide high accuracy but are computationally expensive with large design and simulation times. Compact thermal modeling is another popular thermal modeling methodology that uses the duality between electrical and thermal properties to model temperature. In this approach, the chip is represented as a network of thermal nodes, and the chip temperature is modeled based on an equivalent resistor-capacitor (RC) network of these thermal nodes, where R and C correspond to the thermal resistance and the thermal capacitance, respectively. The equivalent RC network is solved using differential solvers to get the temperature of each node. A CTM enables fast thermal modeling of cooling techniques on target chips with a reasonable tradeoff in accuracy.

The Proposed COMSOL Model

Our proposed COMSOL model solves the coupled fluid flow and heat transfer within the micropillar wick to accurately resolve the temperature distribution in the chip. In steady-state operation, the liquid is continuously supplied via capillary pumping as it evaporates. Since the driving pressure gradient (∇p) is a result of capillary forces, we can relate the pressure at any point on the wick to the local mean curvature of the liquid-vapor interface ($H(x)$) using the Young-Laplace equation (Eq. (1)):

$$p_0 - p(x) = 2\sigma H(x), \quad (1)$$

where p_0 is the VC ambient pressure and σ is the liquid-vapor surface tension of the coolant. The pressure distribution and, thus, the local mean curvature must be calculated to solve the heat transfer problem as the shape of the liquid-vapor interface influences the local HTC. Therefore, fluid flow and heat transfer are coupled through the influence of local pressure on HTC and the response of evaporative flux to the distribution of HTC. The fluid flow domain is modeled as a uniform porous media using Darcy's Law (Eq. (2)):

$$\mathbf{u} = -\frac{\kappa}{\mu} \nabla p, \quad (2)$$

where $\mathbf{u}$ is the flow velocity, κ is the permeability of the wick, and μ is the liquid's dynamic viscosity. By using Darcy's Law, we avoid meshing the exact micropillar geometry, which would be computationally expensive due to the vast number of pillars on the evaporator ($\sim 10^6$). Rather, geometry is lumped into the permeability term. The wick permeability was first solved parametrically for a range of micropillar geometries and interfacial curvatures using a CFD method identical to that described by recent work [13]. We then use these results as a lookup table for permeability as a function of micropillar geometry and local pressure. Reference pressure ($p = p_{sat}$) is assigned to the boundaries contacting the liquid supply while the symmetry condition ($\frac{\partial p}{\partial n} = 0$) is set on the two lines that




Fig. 2: (a) Processing layer grid cell, (b) VC layer grid cell. bisect the wick. We account for the evaporative flux through a volumetric mass loss term that relates to the local heat flux (Eq. (3)):

$$\nabla \cdot (\rho \mathbf{u}) = -\frac{q(x)}{h_{lv}h}, \quad (3)$$

where ρ is the density of the liquid, h_{lv} is the latent heat of vaporization and h is the height of micropillar. $q(x)$, the heat flux distribution at the evaporator surface, is an input from the heat transfer domain, which consists of the solid chip with the input heat flux at the bottom boundary and the evaporative HTC prescribed at the top.

In our work, we solve the Laplace equation for steady state conduction and consider edges of the chip to be thermally insulated. Here, the distribution of HTC is determined from the pressure distribution via a lookup table. We construct this lookup table from a parametric sweep of micropillar geometry and local pressure (or curvature of the liquid-vapor interface) using a finite element solver to calculate the HTC. Our method to parametrically solve HTC is similar to recent work [15], but we additionally introduce local pressure as a parameter to capture the effects of interfacial curvature. We employ an iterative approach to solve the two coupled domains described in the preceding paragraphs. First, to solve the fluid domain, we assume a uniform input heat flux. The output pressure distribution is converted into a distribution of HTC using the lookup table and is passed as an input to the heat transfer domain. In the heat transfer domain, we resolve the temperature distribution and the distribution of heat flux at the evaporator surface is returned to the fluid domain for the next iteration. We output the final temperature distribution once the Euclidian norms of the pressure and heat flux distributions converge to a relative tolerance of 10^{-6} .

The Proposed Compact Thermal Model

In this section, we discuss the methodology for the proposed CTM for two-phase VCs with micropillar wick evaporators. We use a *grid mode* to model the proposed CTM, which means the whole chip system is divided into small grids [8]. To apply our proposed CTM to a processor chip, we model the system as two vertically stacked layers. The bottom one is the processor layer which generates the heat flow. The upper layer is the VC layer which incorporates a micropillar wick evaporator. For processing layer, we use silicon grid cells as shown in Figure 2 (a). Each grid cell has 5 thermal resistances that represent the heat conduction from 5 directions and an additional *virtual temperature node* placed on the bottom surface to represent the temperature of the grid cell.

The VC layer grid cell is shown in Figure 2 (b). Similar to the processing layer grid cell, the virtual temperature node of a

VC layer cell is also placed on the bottom surface. We assign silicon resistivity to lateral thermal resistances to represent the heat conduction from north, south, east, and west directions. In order to represent the heat conduction of the evaporation process, we create an additional virtual temperature node on the top surface of the VC layer grid cell to mimic the saturated vapor inside the VC. The vertical thermal resistance of the grid cell can be interpreted as the heat conduction from the micropillar wick evaporators to the saturated vapor. We assume that inside the VC, there is only saturated vapor that has a constant temperature, T_{Sat} , and all coolants are contained inside wicking structures and evaporators. Based on this assumption, we do not need to model the condenser. Since thermal resistance is inversely proportional to the cross-section area and HTC [9], the vertical thermal resistance can be represented using Eq. (4):

$$R_{VC} = \frac{1}{HTC \cdot w \cdot l}, \quad (4)$$

where w and l are the width and length of the grid cell, respectively. To simplify the evaporation process, instead of a menisci evaporation surface, we assume a flat surface, which results in uniform HTC across the micropillar wick evaporators. Compared to the menisci surface, a flat surface has a relatively smaller heat conduction surface, which in turn underestimates the overall cooling ability of the VCs. The uniform HTC is extracted from the proposed COMSOL model and stored in a lookup table.

To evaluate the cooling performance of the micropillar wick evaporators for different chip architectures, coolants, and power profiles, we also incorporate the dry-out heat flux correlation of figure of merit, micropillar geometric parameter, and contact angle from recent work [16] and design an optimization flow to select the best possible micropillar wick geometry. Algorithm 1 shows the optimization flow. For each micropillar geometry and input combination, we first calculate the dry-out heat flux. Then, if the geometry fails to provide a dry-out heat flux greater than or equal to the maximum chip power density, we proceed to the next geometry. Finally, among all the geometries that satisfy the dry-out heat flux, we pick the one that has the highest HTC.

VALIDATION OF THE PROPOSED COMPACT THERMAL MODEL

In this section, we validate our proposed CTM against a COMSOL CFD model to show the accuracy and speedup of our proposed CTM. We implement and integrate the proposed CTM in HotSpot-6.0 thermal simulator [8]. HotSpot is a CTM-based thermal simulator for modeling processor temperature. This thermal simulator supports a layer configuration feature that allows the user to define layers with desired properties (e.g., silicon, thermal interface material). It also provides the user with packaging layers, such as heat sink (with fan) and heat spreader. To run a simulation in HotSpot, the user needs to provide the physical geometry of the chip stack, floorplan of each layer, thermal properties of the materials, and the power profile of the chip. Based on the input parameters, HotSpot constructs a thermal RC network to solve the temperature

Algorithm 1: Micropillar Geometry Optimization Flow

Input : $q_{dry-out}$ Dry-out heat flux
Input : Q_{max} Hot spot power density on-chip
Input : HTC Heat transfer coefficient
Input : ψ Dimensionless function of micropillar geometry
Input : M Figure of merit for the coolant
Input : θ_{rec} Contact angle
Input : MG Micropillar geometry
Output: MG_{best} Optimized micropillar geometry

```
1  $MG_{best} = None$ 
2 for each  $MG$  in lookup table do
3    $q_{dry-out} = (40/3)\psi M \cos\theta_{rec}$ 
4   if  $q_{dry-out} > Q_{max}$  then
5     if  $HTC_{MG} > HTC_{MG_{best}}$  then
6        $MG_{best} = MG$ 
7     else
8        $next\ MG$ 
9   else
10     $next\ MG$ 
```

for each individual node. This thermal simulator supports both steady-state and transient simulation modes. In steady-state, the thermal network only consists of thermal resistances. HotSpot uses SuperLU-based matrix solver to solve this thermal R network and get the steady-state temperature profile of the simulated chip. The SuperLU-based solver enables fast simulation at the cost of higher memory usage than an iterative solver.

To validate the accuracy of our proposed CTM, we model a specific chip separately in COMSOL and HotSpot. We use the proposed COMSOL model with a menisci evaporation surface to validate our proposed CTM model. The processing layer is a 2×2 mm² square chip with a thickness of 0.1 mm². The VC is placed directly above the processing layer. We run two sets of simulations by using our proposed CTM and COMSOL model: (i) processing layer with uniform power density, and (ii) processing layer with non-uniform power density with a 0.5×0.5 mm² hot spot placed at the center. Each simulation set contains three different micropillar wick geometries and three types of coolant: water, R134a, and R141b. Compared to R134a and R141b, since water has a higher latent heat of evaporation, in order to maintain a safe on-chip temperature ($< 85^\circ\text{C}$), we assign higher power densities to water and relatively lower power densities to the other types of coolant. In the uniform power density cases, we set the power density to 100 W/cm² for water and 20 W/cm² for the other coolants. For the non-uniform power density cases, we assign 50 W/cm², 20 W/cm², and 20 W/cm² as background power densities for water, R134a, and R141b, respectively. The hot spot power densities for simulations with water are set to 100, 200, and 300 W/cm². For simulations using R134a and R141b, hot spot power densities are 25, 50, and 75 W/cm². To avoid extremely high temperatures, we set the T_{Sat} of all three coolants to 50°C , and the corresponding pressure for water,

R134a, and R141b are equal to 0.12, 13.2, and 1.83 bar, respectively. We use 64x64 grids to simulate the steady-state temperature profile of the chip and the accuracy results are shown in Figure 3. For both uniform and non-uniform power density simulations, our CTM demonstrates high accuracy with maximum error and average error of less than 1.25°C and 1.05°C , respectively. The maximum simulation time across the above experiments for a 2×2 mm² chip using COMSOL CFD model is 45s, while it only takes 0.21s to simulate our proposed CTM. Thus, our proposed CTM achieves a 214x speedup in simulation time. In addition, typical accuracies for CTMs of various cooling technologies range from 89.9% to 97.3% [10], [9], [12]. Our proposed CTM provides a 98.5% accuracy which is similar to the approaches mentioned above. The errors are mainly due to the different evaporation surface assumptions in our proposed CTM and the COMSOL model. In COMSOL model, we assume a menisci evaporation surface, which results in a higher HTC than our proposed CTM. These simulation results show that our model can achieve fast simulation time with only a small tradeoff in accuracy.

COMPARISON OF EMERGING COOLING TECHNIQUES

In this section, we compare the cooling performance among five different cooling techniques including two-phase VCs with micropillar wick evaporators, a conventional heat sink, liquid cooling via microchannels, hybrid cooling (TEC and liquid cooling via microchannels), and microchannel-based two-phase cooling. We first discuss the experimental methodologies and then compare the cooling performance of these techniques. Finally, we discuss the additional cooling power related to each cooling technique.

Experimental Methodology

The chip we use in our experiments is 20×20 mm² large and 0.75 mm thick with a 0.5×0.5 mm² hot spot placed at the center. The background power density is set to 50 W/cm² and the hot spot power density is set to 100, 200, 300, 500, 1000, 1500, and 2000 W/cm². We assume the thermal threshold of our chip is 85°C , which means if the maximum temperature surpasses this limit, thermal throttling will be enabled. For a fair comparison, we use water as a coolant with a T_{Sat} of 31.1°C (pressure = 0.045 bar).

Two-Phase VCs with Micropillar Wick Evaporators: We apply the proposed CTM to the experimental chip with different hot spot power densities. The steady-state temperature simulation results for each power profile are shown in Table I. In these simulations, we use the micropillar wick geometries selected by the optimization flow and the optimal geometries are also shown in Table I. PD_{HS} stands for the hot spot power density. T_{Max} and ΔT represent the hot spot temperature and the thermal gradients, respectively. As we can see from the results, when $PD_{HS} = 100, 200, 300, 500, 1000,$ and 1500 W/cm², using two-phase VCs with micropillar wick evaporators satisfies the chip thermal constraint. However, when $PD_{HS} = 2000$ W/cm², we can only reduce the hot spot temperature down to 87°C .




Fig. 3: Accuracy results compared to the COMSOL model for (a) uniform power densities, and (b) non-uniform power densities with a hot spot. The power densities shown on the x-axis in (b) are the hot spot power densities. The x-axis also shows the micropillar geometries in terms of height, diameter, and pitch (μm). PD and PD_{bg} stand for uniform power density and background power density, respectively.

Conventional Heat Sink: We model the same chip in HotSpot and apply a default heat sink as the cooling method. The length, width, and thickness of the heat sink are 6, 6, and 0.69 cm, respectively. The thermal conductivity of the heat sink is 400 W/(mK) and the convection resistance is 0.1 K/W.

Liquid Cooling via Microchannels: We build a CTM for liquid cooling via microchannels introduced in recent work [10]. We model the same chip in HotSpot with a 0.1 mm thick liquid microchannel layer on top of it as a heat sink. An additional 0.04 mm thick silicon bulk layer (packaging layer) is placed on top of the liquid microchannel layer to seal the channels. We set the microchannel width to 0.05 mm (same as the wall width) and the channel thickness to 0.1 mm. At the top of the silicon bulk layer, we assign air convection HTC to 0.01 W/cm² to represent minimal convection to the ambient. We set the liquid flow velocity to 0.5, 1, 1.5, and 2.6 m/s [9].

Hybrid Cooling: For the chip system described in the liquid cooling via microchannels section, we add an additional TEC layer between the processing layer and the liquid microchannel layer to build a hybrid cooling system. The chip stack is shown in Figure 4. A CTM of the TEC unit is described in recent work [10]. The TEC device is a superlattice-based thin-film TEC made of Bi_2Te_3 as the bulk material. It is composed of an array of 7×7 p-n thermocouples. In the TEC layer, a $3.5 \times 3.5 \text{ mm}^2$ TEC unit is placed directly above the hot spot in the processing layer. The rest of the TEC layer is covered in




Fig. 4: Hybrid cooling chip stack with different layers thickness.

silicon. We vary the liquid flow velocities from 0.5 m/s to 2.6 m/s and the TEC currents from 1 A to 7 A [4].

Microchannel-based Two-Phase Cooling: Inspired by the microchannel-based two-phase cooling CTM in STEAM [12], we build another CTM for the same technique in HotSpot. We adopt the most accurate HTC correlation [18] reported by STEAM. In order to validate the accuracy of the microchannel-based two-phase cooling CTM, we model a chip similar to the one in liquid cooling via microchannels section. A 0.1 mm thick microchannel layer is placed on top of a $0.5 \times 0.25 \text{ mm}^2$ chip. We place an additional silicon bulk layer to seal the channels and minimize the air convection. Both channel width and wall width are set to 0.05 mm. We use R245fa as the coolant with an input mass flux of 700 kg/m²s. T_{Sat} of the coolant is set to 31.1°C and the corresponding pressure is 1.84 bar. We run steady-state thermal simulations using uniform power profiles of 25, 100, and 200 W/cm². We model the same chip in STEAM [12], which has been validated using prototypes. Among all the simulations, when compared to the STEAM model, our microchannel-based two-phase CTM achieves a maximum error of less than 0.5°C.

For the cooling performance experimental setup of microchannel-based two-phase cooling, We model the same chip as the one in the liquid cooling via microchannels section and select water as the coolant with $T_{Sat} = 31.1^\circ\text{C}$ (pressure = 0.045 bar). We vary the input mass flux from 100 kg/m²s to 560 kg/m²s [6].

TABLE I: Optimal geometries (height, diameter, pitch), steady-state hot spot temperatures, and thermal gradients for various power profiles. MG_{best} stands for optimal micropillar geometry.

PD_{HS} (W/cm ²)	100	200	300
MG_{best} (μm)	20,10,5	25,10,5	30,10,5
T_{Max} (°C)	46	48	50
ΔT (°C)	1	3	5
500	1000	1500	2000
30,10,5	35,10,5	45,10,5	45,10,5
55	66	77	87
10	20	31	42




Fig. 5: Comparison of cooling performance between hybrid cooling and liquid cooling via microchannels, when liquid flow velocity = 1.5 m/s.

Cooling Performance Comparison

Liquid Cooling via Microchannels vs Hybrid Cooling:

When flow velocity = 1.5 m/s, the comparison results of liquid cooling via microchannels and hybrid cooling are shown in Figure 5. When TEC current = 1 A, compared to liquid cooling via microchannels, hybrid cooling results in higher T_{Max} s and ΔT s, which indicates liquid cooling via microchannels has better cooling performance in these scenarios. The reason is that adding an additional TEC layer will increase the vertical thermal resistance (due to increased thickness) from the processing layer to the liquid cooling heat sink. In addition, Joule heating in the TEC device also contributes to the increasing hot spot temperature. When TEC current = 5 and 7 A, and PD_{HS} = 100, 200, 300, and 500 W/cm², liquid cooling via microchannels results in better cooling performance than hybrid cooling. In these cases, the high TEC currents make Joule heating dominant and therefore increases the hot spot temperature. However, when TEC current = 5 and 7 A, and PD_{HS} = 1000, 1500, and 2000 W/cm², hybrid cooling results in better cooling performance.

Vapor Chambers vs Single-Phase Cooling and Conventional Heat Sink: We compare the cooling performance of two-phase VCs to the single-phase cooling techniques (liquid cooling via microchannels and hybrid cooling) and the conventional heat sink. We assume the condenser is a high air convection HTC heat sink. In such a case, two-phase VCs with micropillar wick evaporators do not require any cooling power at all. As we discussed above, liquid cooling via microchannels has better cooling performance when PD_{HS} = 100, 200, 300, and 500 W/cm², while hybrid cooling provides better cooling performance when PD_{HS} = 1000, 1500, and 2000 W/cm². We evaluate the cooling performance of the two-phase VCs by comparing to liquid cooling via microchannels when PD_{HS} = 100, 200, 300, and 500 W/cm², and comparing to hybrid cooling when PD_{HS} = 1000, 1500, and 2000 W/cm². Figure 6 shows the results for PD_{HS} = 100, 200, 300, and 500 W/cm². Only when flow velocity = 2.6 m/s and PD_{HS} = 300 and 500 W/cm², liquid cooling via microchannels shows lower T_{Max} s than two-phase VCs. In all other experiments, two-phase VCs provide better cooling performance when compared to liquid cooling via microchannels and the conventional heat sink. Figure 7 shows the results for PD_{HS} = 1000, 1500, and 2000 W/cm². As the figure suggests, in some low flow velocity




Fig. 6: Comparison of cooling performance among the conventional heat sink, two-phase VCs with micropillar evaporators, and liquid cooling via microchannels.




Fig. 7: Comparison of cooling performance among the conventional heat sink, two-phase VCs with micropillar evaporators, and hybrid cooling when TEC current = 7 A.

cases (0.5 and 1 m/s), two-phase VCs show lower T_{max} s or ΔT s, but generally, hybrid cooling provides better cooling performance compared to two-phase VCs. Moreover, when the hot spot power density increases, the cooling performance of two-phase VCs start decreasing. When PD_{HS} = 2000 W/cm², the conventional heat sink has a relatively lower ΔT than two-phase VCs. The reason is that as we increase the maximum chip power density, the optimization flow selects the micropillar geometry that has a higher micropillar height. Increasing the height of the micropillar array can enhance the dry-out heat flux. However, it also decreases the HTC because of the thicker liquid film. For the conventional heat sink, one can observe that except for PD_{HS} = 100, 200, and 300 W/cm² cases, all other experiments result in hot spot temperatures higher than 85°C, which means that the conventional heat sink is not an effective cooling solution for high power density chips.

Microchannel-based Two-Phase Cooling vs Single-Phase Cooling:

We further compare the cooling performance between the microchannel-based two-phase cooling and the single-phase cooling techniques, namely hybrid cooling and liquid-cooling via microchannels. For hybrid cooling and liquid cooling via microchannels, we set flow velocity = 2.6 m/s and TEC current = 7 A. For microchannel-based two-phase cooling, we set mass flux = 560 kg/m²s. The comparison results are shown in Figure 8. Across all the experiments, microchannel-based two-phase cooling always provides better cooling performance than liquid cooling via microchannels. Hybrid cooling shows lower T_{Max} s and ΔT s than microchannel-based two-phase cooling when PD_{HS} = 1500 and 2000 W/cm², while in all other cases, microchannel-based two-phase cooling provides better cooling performance than hybrid cooling. Theoretically, latent heat evaporation of fluid can result in a higher HTC than single-phase cooling such




Fig. 8: Comparison of cooling performance among hybrid cooling, microchannel-based two-phase cooling, and liquid cooling via microchannels.




Fig. 9: Comparison of cooling performance between two-phase VCs with micropillar evaporators and microchannel-based two-phase cooling of various input mass fluxes: (a) 100 kg/m²s, (b) 560 kg/m²s.

as liquid cooling via microchannels. Nevertheless, adding an additional TEC unit with a high bias current can significantly reduce the hot spot temperature. The above results show that for hot spot power densities of 1500 and 2000 W/cm², hybrid cooling with a high bias current can achieve better cooling performance.

Vapor Chambers vs Microchannel-based Two-Phase Cooling: Figure 9 shows the cooling performance comparison results between the microchannel-based two-phase cooling and two-phase VCs. As we can see from the results, compared to microchannel-based two-phase cooling, two-phase VCs show better cooling performance when $PD_{HS} = 100, 200,$ and 300 W/cm² and input mass flux = 100 kg/m²s. When the input mass flux increases to 560 kg/m²s, microchannel-based two-phase cooling always shows better cooling performance. A higher mass flux can remove the heat more effectively, but similar to single-phase cooling, it also requires additional cooling power. We will discuss this in the following section.

Next, we compare the cooling performance of two-phase VCs to microchannel-based two-phase cooling by using a realistic chip architecture and power profile. We select a 256-core Intel SCC processor scaled to 22nm presented in a recent work [19]. The core architecture is based on IA-32 core [20]. The power profiles of this processor are generated by running various multi-threaded benchmarks using Sniper [21] and McPAT [22]. These power profiles are then used to conduct steady-state thermal simulations using HotSpot simulator [19]. We choose the power profile that results in maximum hot spot temperature and thermal gradients. The selected power profile has a maximum power density of 216.6 W/cm². We then conduct steady-state thermal simulations




Fig. 10: Intel SCC thermal maps with different cooling technologies: (a) two-phase VCs with micropillar wick evaporators, and (b) microchannel-based two-phase cooling.

for this processor using our proposed model as well as the microchannel-based two-phase cooling CTM. We select R141b as a coolant because it is dielectric and more realistic than water, with a mass flux of 2000 kg/m²s [6]. T_{Sat} is set to 50°C (pressure = 1.83 bar). The simulated thermal maps are shown in Figure 10. We observe that two-phase VCs with micropillar wick evaporators achieve up to 16°C lower peak temperatures and 11°C lower thermal gradients.

Comparison of Additional Cooling Power

Table II shows the estimated additional cooling power for the above active cooling techniques (liquid cooling via microchannels, hybrid cooling and microchannel-based two-phase cooling). For liquid cooling via microchannels, the total additional cooling power is the same the pumping power of the hybrid cooling technique that has the same liquid flow velocity. The estimated pumping power is calculated based on Eq. (5) [11]:

$$P_{pump} = \frac{\Delta P \cdot V}{\eta}, \quad (5)$$

where ΔP is the pressure drop across the channel, V is the volumetric flow rate, and η is the pump efficiency.

TEC power consumption is calculated as follows:

$$P_{TEC} = Q_{hot} - Q_{cold}, \quad (6)$$

where Q_{hot} and Q_{cold} are the heat absorbed and rejected on the cold and hot sides, respectively.

As we can see from the Table II, hybrid cooling with TEC current = 7 A and liquid flow velocity = 2.6 m/s results in the highest cooling power of 15.6 W. Meanwhile, it also provides the best cooling performance for hot spot power densities of 1500 and 2000 W/cm². Using two-phase VCs with micropillar wick evaporators is attractive because it shows the favorable cooling ability for power densities less than 500 W/cm² without additional cooling power on the evaporator side.

Comparison Results Summary

Based on the above observations, we conclude the comparison results as follows:

- The conventional heat sink is not capable of cooling high power density chips (e.g., $500, 1000, 1500,$ or 2000 W/cm²).

W/cm² and results in the worst cooling performance among the aforementioned cooling techniques.

- Two-phase cooling techniques (two-phase VCs with micropillar wick evaporators and microchannel-based two-phase cooling) have better cooling performance for hot spot power densities of 100, 200, 300, 500, and 1000 W/cm², when compared to other cooling techniques.
- When the hot spot power density is high, such as 1500 and 2000 W/cm², hybrid cooling shows superior cooling performance than two-phase cooling techniques.
- Since the coolant in VCs is passively driven by the capillary force generated by the wicking structure, two-phase VCs are also expected to save cooling power.

RELATED WORK

This section presents a review of existing models for two-phase cooling, liquid cooling via microchannels, TEC cooling, and hybrid cooling with TEC and liquid cooling via microchannels. It then discusses the distinguishing aspects of our proposed model for two-phase VCs with micropillar wick evaporators.

Two-phase cooling: Sridhar *et al.* [12] build a CTM-based simulator to model the liquid-vapor phase change inside microchannels. Chen *et al.* [14] introduce a numerical model for two-phase cooling in vapor chambers. PCMs are another example of solid-liquid phase change. These are explored as a passive cooling technique to boost performance by computation sprinting [23]. Many PCM thermal models have been built to evaluate its thermal management capabilities and energy savings [24], [25].

Liquid cooling via microchannels: Liquid cooling via microchannels can provide inter-tier cooling in 3D-stacked architectures. Coskun *et al.* [11] build a CTM for this technique in 3D multicore systems and integrate it into the HotSpot-4.01 simulator. Sridhar *et al.* [9] introduce a CTM-based simulator, 3D-ICE, that can model the temperature gradient between the inlet and outlet of the microchannels. They validate the accuracy of their model with a CFD tool.

TEC: TECs have gained attraction due to their ability to mitigate high-density hot spots [4]. Paterna *et al.* [26] investigate TECs in conjunction with dynamic voltage and

frequency scaling and a number of workload threads, to maximize performance under the given power and thermal constraints on multicore processors. A recent work [10] builds a CTM for this technique to investigate its heat removing capabilities.

Hybrid of TECs and liquid cooling via microchannels: Hybrid cooling with TECs and liquid cooling via microchannels has recently been proposed to efficiently remove high-density hot spots [10], [7], [27]. Kaplan *et al.* [10] introduce a CTM for this hybrid technique to show its superior cooling performance on high density hot spots when compared to liquid cooling via microchannels. They further integrate their model into the HotSpot-6.0 simulator.

The key distinguishing aspects of our work are as follows:

- 1) This work introduces a steady-state CTM for two-phase VCs with micropillar wick evaporators.
- 2) While there are many emerging cooling solutions to target localized hot spots, our work is the first to compare these techniques against each other for future processors and discuss their tradeoffs.

CONCLUSION

In this paper, we design a CTM for two-phase VCs with micropillar wick evaporators for fast and accurate steady-state thermal evaluation of this cooling technique. We have integrated our CTM in HotSpot-6.0 thermal simulator and provided a comparison against a COMSOL CFD model. Our proposed CTM provides a maximum error of less than 1.25 °C in comparison to COMSOL. In addition, our model also achieves a 214x speedup of the simulation time. We carry out a comprehensive study of the cooling performance for various cooling techniques including a conventional heat sink, liquid cooling via microchannels, hybrid cooling (TEC and liquid cooling via microchannels), microchannel-based two-phase cooling, and two-phase VCs with micropillar wick evaporators. Using two-phase VCs with micropillar wick evaporators shows better cooling performance than single-phase cooling when the hot spot power density is less than 500 W/cm². Two-phase VCs also provide better cooling performance than microchannel-based two-phase cooling when the mass flux is under 100 kg/m²s and the hot spot power density is less than 300 W/cm². In addition, no additional cooling power on the evaporator side is needed for two-phase VCs with micropillar wick evaporators.

ACKNOWLEDGEMENT

This project has been partially funded by the NSF CRI (CI-NEW) grant #1730316/1730003/1730389.

REFERENCE

- [1] M. Schultz, F. Yang, E. Colgan, R. Polastre, B. Dang, C. Tsang, M. Gaynes, P. Parida, J. Knickerbocker, and T. Chainer, "Embedded two-phase cooling of large three-dimensional compatible chips with radial channels," *Journal of Electronic Packaging*, vol. 138, no. 2, p. 021005, 2016.
- [2] L. Shao, A. Raghavan, L. Emurian, M. C. Papaefthymiou, T. F. Wenisch, M. M. Martin, and K. P. Pipe, "On-chip phase change heat sinks designed for computational sprinting," in *IEEE Proc. of 30th Annual Semiconductor Thermal Measurement and Management Symposium (SEMI-THERM)*, 2014, pp. 29–34.

TABLE II: Additional cooling power for hybrid cooling and microchannel-based two-phase cooling.

	Parameter	Pumping	TEC	Total
Hybrid	1 A, 0.5 m/s	0.17 (W)	0.16 (W)	0.33 (W)
	1 A, 1.0 m/s	0.66 (W)	0.16 (W)	0.82 (W)
	1 A, 1.5 m/s	1.5 (W)	0.16 (W)	1.66(W)
	1 A, 2.6 m/s	4.5 (W)	0.16 (W)	4.66(W)
	5 A, 0.5 m/s	0.17 (W)	5.4(W)	5.57(W)
	5 A, 1.0 m/s	0.66 (W)	5.4(W)	6.06(W)
	5 A, 1.5 m/s	1.5 (W)	5.4(W)	6.9(W)
	5 A, 2.6 m/s	4.5 (W)	5.4(W)	9.9(W)
	7 A, 0.5 m/s	0.17 (W)	11.1 (W)	11.27(W)
	7 A, 1.0 m/s	0.66 (W)	11.1 (W)	11.76(W)
	7 A, 1.5 m/s	1.5 (W)	11.1 (W)	12.6(W)
	7 A, 2.6 m/s	4.5 (W)	11.1 (W)	15.6(W)
Two-phase	100 kg/m ² s (0.1m/s)	0.24 (W)	N/A	0.24 (W)
	200 kg/m ² s (0.2m/s)	0.49 (W)	N/A	0.49 (W)
	560 kg/m ² s (0.56m/s)	1.36 (W)	N/A	1.36 (W)

- [3] B. Dang, M. S. Bakir, D. C. Sekar, C. R. King Jr, and J. D. Meindl, "Integrated microfluidic cooling and interconnects for 2D and 3D chips," *IEEE Transactions on Advanced Packaging*, vol. 33, no. 1, pp. 79–87, 2010.
- [4] I. Chowdhury, R. Prasher, K. Lofgreen, G. Chrysler, S. Narasimhan, R. Mahajan, D. Koester, R. Alley, and R. Venkatasubramanian, "On-chip cooling by superlattice-based thin-film thermoelectrics," *Nature nanotechnology*, vol. 4, no. 4, p. 235, 2009.
- [5] M. Bulut, S. G. Kandlikar, and N. Sozbir, "A review of vapor chambers," *Heat Transfer Engineering*, pp. 1–23, 2018.
- [6] S. G. Kandlikar, "Fundamental issues related to flow boiling in minichannels and microchannels," *Experimental Thermal and Fluid Science*, vol. 26, no. 2–4, pp. 389–407, 2002.
- [7] K. Yazawa, A. Ziabari, Y. R. Koh, A. Shakouri, V. Sahu, A. G. Fedorov, and Y. Joshi, "Cooling power optimization for hybrid solid-state and liquid cooling in integrated circuit chips with hotspots," in *IEEE Proc. of 13th Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm)*, 2012, pp. 99–106.
- [8] K. Skadron, M. R. Stan, W. Huang, S. Velusamy, K. Sankaranarayanan, and D. Tarjan, "Temperature-aware microarchitecture," in *IEEE Proc. of 30th Annual International Symposium on Computer Architecture (ISCA)*, 2003, pp. 2–13.
- [9] A. Sridhar, A. Vincenzi, M. Ruggiero, T. Brunschweiler, and D. Atienza, "3D-ICE: Fast compact transient thermal modeling for 3D ICs with inter-tier liquid cooling," in *IEEE Proc. of the International Conference on Computer-Aided Design (ICCAD)*, 2013, pp. 463–470.
- [10] F. Kaplan, S. Reda, and A. K. Coskun, "Fast thermal modeling of liquid, thermoelectric, and hybrid cooling," in *IEEE Proc. of Intersociety Conf. on Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm)*, 2017, pp. 726–735.
- [11] A. K. Coskun, J. L. Ayala, D. Atienza, and T. S. Rosing, "Modeling and dynamic management of 3D multicore systems with liquid cooling," in *IFIP/IEEE 17th International Conference on Very Large Scale Integration (VLSI-SoC)*, 2009, pp. 35–40.
- [12] A. Sridhar, Y. Madhour, D. Atienza, T. Brunschweiler, and J. Thome, "STEAM: A fast compact thermal model for two-phase cooling of integrated circuits," in *IEEE Proc. of International Conference on Computer-Aided Design (ICCAD)*, 2013, pp. 256–263.
- [13] Y. Zhu, D. S. Antao, Z. Lu, S. Somasundaram, T. Zhang, and E. N. Wang, "Prediction and characterization of dry-out heat flux in micropillar wick structures," *Langmuir*, vol. 32, no. 7, pp. 1920–1927, 2016.
- [14] Y.-S. Chen, K.-H. Chien, T.-C. Hung, C.-C. Wang, Y.-M. Ferng, and B.-S. Pei, "Numerical simulation of a heat sink embedded with a vapor chamber and calculation of effective thermal conductivity of a vapor chamber," *Applied Thermal Engineering*, vol. 29, no. 13, pp. 2655–2664, 2009.
- [15] M. Wei, B. He, Q. Liang, S. Somasundaram, C. S. Tan, and E. N. Wang, "Optimization and thermal characterization of uniform silicon micropillar based evaporators," *International Journal of Heat and Mass Transfer*, vol. 127, pp. 51–60, 2018.
- [16] S. Adera, D. Antao, R. Raj, and E. N. Wang, "Design of micropillar wicks for thin-film evaporation," *International Journal of Heat and Mass Transfer*, vol. 101, pp. 280–294, 2016.
- [17] X. Wei and K. Sikka, "Modeling of vapor chamber as heat spreading devices," in *IEEE Proc. of 10th Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronics Systems (ITherm)*, 2006, pp. 578–585.
- [18] S. S. Bertsch, E. A. Groll, and S. V. Garimella, "A composite heat transfer correlation for saturated flow boiling in small channels," *International Journal of Heat and Mass Transfer*, vol. 52, no. 7–8, pp. 2110–2118, 2009.
- [19] F. Eris, A. Joshi, A. B. Kahng, Y. Ma, S. Mojumder, and T. Zhang, "Leveraging thermally-aware chiplet organization in 2.5D systems to reclaim dark silicon," in *IEEE Proc. of Design, Automation & Test in Europe Conference & Exhibition (DATE)*, 2018, pp. 1441–1446.
- [20] J. Howard, S. Dighe, S. R. Vangal, G. Ruhl, N. Borkar, S. Jain, V. Erraguntla, M. Konow, M. Riepen, M. Gries *et al.*, "A 48-core IA-32 processor in 45 nm CMOS using on-die message-passing and DVFS for performance and power scaling," *IEEE Journal of Solid-State Circuits*, vol. 46, no. 1, pp. 173–183, 2011.
- [21] T. E. Carlson, W. Heirman, and L. Eeckhout, "Sniper: Exploring the level of abstraction for scalable and accurate parallel multi-core simulation," in *ACM Proc. of International Conference for High Performance Computing, Networking, Storage and Analysis*, 2011, p. 52.
- [22] S. Li, J. H. Ahn, R. D. Strong, J. B. Brockman, D. M. Tullsen, and N. P. Jouppi, "McPAT: An integrated power, area, and timing modeling framework for multicore and manycore architectures," in *IEEE/ACM Proc. of 42nd Annual International Symposium on Microarchitecture (MICRO)*, 2009, pp. 469–480.
- [23] A. Raghavan, Y. Luo, A. Chandawalla, M. Papaefthymiou, K. P. Pipe, T. F. Wenisch, and M. M. Martin, "Computational sprinting," in *IEEE Proc. of 18th International Symposium on High Performance Computer Architecture (HPCA)*, 2012, pp. 1–12.
- [24] F. Kaplan, C. De Vivero, S. Howes, M. Arora, H. Homayoun, W. Burleson, D. Tullsen, and A. K. Coskun, "Modeling and analysis of phase change materials for efficient thermal management," in *IEEE Proc. of 32nd International Conference on Computer Design (ICCD)*, 2014, pp. 256–263.
- [25] D.-w. Yoo and Y. K. Joshi, "Energy efficient thermal management of electronic components using solid-liquid phase change materials," *IEEE Transactions on Device and Materials Reliability*, vol. 4, no. 4, pp. 641–649, 2004.
- [26] F. Paterna and S. Reda, "Mitigating dark-silicon problems using superlattice-based thermoelectric coolers," in *IEEE Proc. of Design, Automation & Test in Europe Conference & Exhibition (DATE)*, 2013, pp. 1391–1394.
- [27] V. Sahu, A. G. Fedorov, Y. Joshi, K. Yazawa, A. Ziabari, and A. Shakouri, "Energy efficient liquid-thermoelectric hybrid cooling for hot-spot removal," in *IEEE 28th Annual Conference on Semiconductor Thermal Measurement and Management Symposium (SEMI-THERM)*, 2012, pp. 130–134.