A Learning-Based Thermal Simulation Framework for Emerging Two-Phase Cooling Technologies

Zihao Yuan[1], Geoffrey Vaartstra[2], Prachi Shukla[1], Zhengmao Lu[2], Evelyn Wang[2], Sherief Reda[3], and Ayse K. Coskun[1]

[1]Boston University, Boston, MA - {yuan1z, prachis, acoskun}@bu.edu
[2]Massachusetts Institute of Technology, Cambridge, MA - {gvaartst, zhengmao, enwang}@mit.edu
[3]Brown University, Providence, RI - {sherief_reda}@brown.edu

***Abstract*—Future high-performance chips will require new cooling technologies that can extract heat efficiently. Two-phase cooling is a promising processor cooling solution owing to its high heat transfer rate and potential benefits in cooling power. Two-phase cooling mechanisms, including microchannel-based two-phase cooling or two-phase vapor chambers (VCs), are typically modeled by computing the temperature-dependent heat transfer coefficient (HTC) of the evaporator or coolant using an iterative simulation framework. Precomputed HTC correlations are specific to a given cooling system design and cannot be applied to even the same cooling technology with different cooling parameters (such as different geometries). Another challenge is that HTC correlations are typically calculated with computational fluid dynamics (CFD) tools, which induce long design and simulation times. This paper introduces a learning-based temperature-dependent HTC simulation framework that is used to model a two-phase cooling solution with a wide range of cooling design parameters. In particular, the proposed framework includes a compact thermal model (CTM) of two-phase VCs with hybrid wick evaporators (of nanoporous membrane and microchannels). We build a new simulation tool to integrate the proposed simulation framework and CTM. We validate the proposed simulation framework as well as the new CTM through comparisons against a CFD model. Our simulation framework and CTM achieve a speedup of 21× with an average error of 0.98°C (and a maximum error of 2.59°C). We design an optimization flow for hybrid wicks to select the most beneficial hybrid wick geometries. Our flow is capable of finding a geometry-coolant combination that results in a lower (or similar) maximum chip temperature compared to that of the best coolant-geometry pair selected by grid search, while providing a speedup of 9.4×.**

I. Introduction

Over the last few decades, processor performance has grown tremendously following the down-scaling of transistors. High power densities that reach 1-2 kW/cm^2 [8] caused by the performance boost can occur in future high performance chips and result in amplified localized hot spots. Existing cooling solutions such as forced air cooling via fans or traditional pin-fin heat sinks are often not sufficient to mitigate these high power densities efficiently and can lead to over/under-cooling, affecting system design cost and power. Therefore, designing new cooling solutions for high-performance processors has started to gain traction.

Recently, researchers have developed several emerging cooling solutions such as liquid cooling via microchannels [4], thermoelectric coolers (TECs) [2], two-phase cooling (e.g., microchannel-based two-phase cooling, two-phase vapor chambers (VCs), or heat pipes) [1], [14], and hybrid cooling (e.g., hybrid of liquid cooling via microchannels and TECs [16]). Among these emerging cooling solutions, two-phase cooling is particularly attractive as it removes heat effectively and has the potential to limit the cooling power compared to the cooling power of single-phase cooling techniques. For example, VCs and heat pipes leverage capillary pumping to passively (without active cooling power on the evaporator side) supply the evaporator [1], [14]. Two-phase cooling thermal models are usually designed by calculating the temperature-dependent heat transfer coefficient[1] (HTC) [8], [12], [14], [15]. In these models, pre-computed temperature-dependent HTC correlations embedded in the simulation framework are functions of temperature and cooling parameters (e.g., coolant type, flow velocity, saturation temperature, and structural parameter such as microchannel width and height, micropillar height, diameter, and pitch [12], [15], [17]). The HTC correlations are derived either based on in-house prototypes or using computational fluid dynamics (CFD) modules in COMSOL and ANSYS [1], [15]. HTC correlations based on prototypes are generally not applicable to the same cooling method with different cooling parameters, or will likely result in accuracy loss [12]. Commercial simulation tools (e.g., COMSOL and ANSYS), on the other hand, are computationally expensive and experience long design and simulation times as well as large memory requirements. To enable cooling design exploration and optimization, there is a need for a fast and generalized temperature-dependent HTC simulation framework that can be applied for a wide range of cooling parameters for the same two-phase cooling technology while maintaining the desired accuracy.

In this work, we introduce a learning-based temperature-dependent HTC simulation framework to model two-phase cooling technologies for processor cooling. The proposed simulation framework enables fast and accurate simulations with a wide range of cooling parameters for the same two-phase cooling technology. To demonstrate the simulation speedup and accuracy of our proposed simulation framework, we build a compact thermal model (CTM) for two-phase VCs with hybrid wick evaporators (of nanoporous membrane and microchannels) and integrate it into our proposed simulation framework. In addition, we design an optimization flow to select the nanoporous and microchannel geometries that minimize the hot spot temperature. The main contributions of our work are as follows:

- We propose a learning-based temperature-dependent HTC simulation framework for two-phase cooling solutions and embed our proposed steady-state two-phase VCs with hybrid wick evaporators CTM in the simulation framework (Sec. III).
- We build a CTM tool and integrate our proposed learning-based temperature-dependent HTC simulation framework and the CTM[2] (Sec. III-E). We validate our proposed simulation framework and CTM against a COM-

[1]Heat transfer coefficient is a parameter that determines the rate of heat transfer per unit temperature difference.

[2]We will open source our modeling tool along with paper publication.

SOL model. Our proposed learning-based temperature-dependent HTC simulation framework achieves a 21× speedup with an average error of 0.98°C when compared to the COMSOL model (Secs. III and V).
- We design an optimization flow that maximizes the cooling performance of two-phase VCs with hybrid wick evaporators by selecting the best nanoporous membrane and microchannel geometry. Our proposed optimization flow is capable of finding better (or similar) hybrid wick geometry and coolant combinations than grid search with an average speedup of 4× (Secs. IV and V).

II. Related Work

High power densities have already become major concerns in processors. To mitigate the hot spots, a group of work has focused on creating design-time and run-time thermal management policies such as thermal-aware floorplanning, dynamic voltage and frequency scaling, task scheduling/migration, and others. Another body of work focuses on designing emerging cooling technologies such as liquid cooling via microchannels, TECs, two-phase VCs, hybrid cooling (e.g., of liquid cooling via microchannels and TEC), and microchannel-based two-phase cooling, and integrating these technologies as heat sinks or inter-layer cooling methods [1], [2], [4], [14], [16]. In addition, recent work has introduced fast and accurate CTMs for various emerging cooling methods [3], [7], [12], [13].

Sridhar *et al.* introduced STEAM, a CTM simulator for microchannel-based two-phase cooling [12]. STEAM is a temperature-dependent HTC simulation framework integrated with a set of predetermined HTC correlations. STEAM's average error ranges from 10.2% to 23.7% depending on the HTC correlation it uses. The reason for this accuracy loss is that these correlations are derived from various prototypes with different microchannel height, width, shape, and mass flow velocity values, and they are not sufficiently modular for application to any other microchannel structures and flow velocities. Yuan *et al.* introduced a CTM for two-phase VCs with micropillar wick evaporators [17], [18]. The HTC values of micropillar wick geometries are extracted from a finite element model and stored in a lookup table. However, their model does not enable simulation for arbitrary geometries that are not stored in the lookup table. For other valid geometries that have not been stored in the lookup table, one has to run CFD simulations first to generate the HTC data.

The key innovation in our work is that we introduce a learning-based temperature-dependent HTC simulation framework that is capable of running fast and accurate thermal simulations for two-phase cooling methods with a wide range of cooling design parameters. We build a tool and integrate the simulation framework with a CTM for two-phase VCs with hybrid wick evaporators to show the accuracy and speedup of our model (in comparison to COMSOL). We also introduce a hybrid wick evaporator optimization flow to rapidly select the best geometry to maximize cooling performance.

III. Modeling of Two-Phase VCs with Hybrid Wick Evaporators

In this work, we focus on two-phase VCs because of their high heat transfer rate with no additional cooling power on the evaporator side. A hybrid wick evaporator is particularly interesting because it enhances both HTC and dry-out heat flux (dry-out heat flux is the thermal limit of a two-phase device). We design a CTM as well as a learning-based temperature-dependent HTC simulation framework to enable fast and accurate thermal simulation for this cooling technology with a wide range of valid hybrid wick geometries.

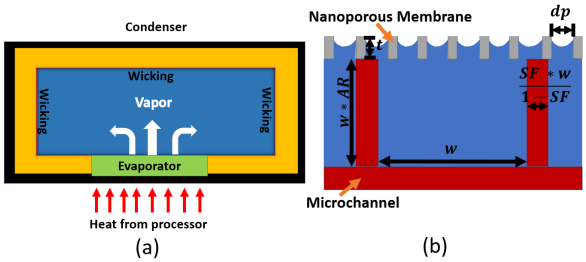

Fig. 1: (a) A vapor chamber structure view, and (b) a hybrid wick evaporator cross-section view.

Fig. 1 (a) shows the schematic of a VC. The coolant absorbs heat from the processor when it vaporizes at the evaporator. Heat is rejected to the ambient at the condenser and the condensate flows back to the evaporator driven by the wicking structure along the side walls. The heat removal ability of the VC is often dominated by the evaporator [1]. An evaporator with a higher HTC is desired to reduce the thermal resistance of the VCs. However, such high-HTC evaporators often suffer from low critical dry-out heat flux [1], [6]. These two metrics are typically conflicting with each other and it is challenging to maximize HTC while enhancing dry-out heat flux [1]. In this work, we focus on a hybrid wick evaporator (of nanoporous membrane and microchannels) as shown in Fig. 1 (b) that improves both HTC and dry-out heat flux. The microchannel and membrane geometries can be varied independently so as to enhance the permeability of the microchannels and the heat transfer from highly conductive solids (the substrate, microchannels, and nanoporous membrane) to the liquid-vapor interface.

A. CTM for Two-Phase VCs with Hybrid Wick

To build a CTM, we abstract both the nanoporous membrane and the microchannel layers into a hybrid wick layer. The whole chip stack is shown in Fig. 2 (a). We divide the whole chip into grids as in prior work [17]. For the processing layer, the grid cell structure is shown in Fig. 2 (b). Thermal resistance along the north, south, east, west, and vertical directions are represented using silicon properties, i.e., $R_{silicon}$. As for the hybrid wick layer, the grid cell is shown in Fig. 2 (c). We represent lateral thermal resistances using $R_{silicon}$ and the vertical thermal resistance, R_{hybrid}, stands for the inverse of the heat conduction from the hybrid wick to the saturated vapor. We add an additional virtual temperature node on top of the hybrid wick grid cell to represent the saturated vapor. We consider steady-state and a predetermined VC pressure. In this way, we do not need to model the heat sink on top of the VC. Instead, we use a previously established relationship between HTC and thermal resistance to define R_{hybrid} [13], [17]. In addition, we assume the VC itself only contains saturated vapor at a constant temperature [8], [15]. From the COMSOL model, we extract HTC correlations for various nanoporous membranes and microchannel geometries. For a specific hybrid wick geometry, we use its corresponding HTC correlation to determine the R_{hybrid} value.

B. Building a COMSOL Model

We determine the resistance of the hybrid wick layer by solving for its effective HTC via finite element calculations using COMSOL. The simulated domain consists of a nanoporous

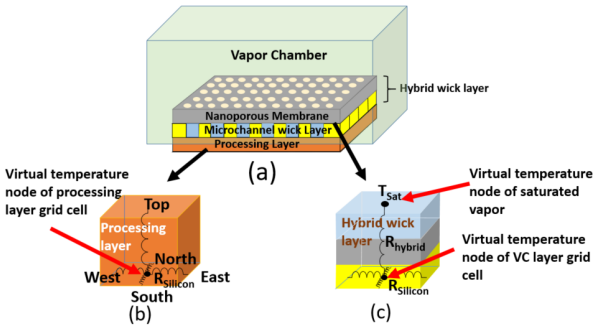

Fig. 2: (a) The chip stack of the processing layer and two-phase VCs with hybrid wick evaporators, (b) processing layer grid cell, and (c) hybrid wick layer grid cell.

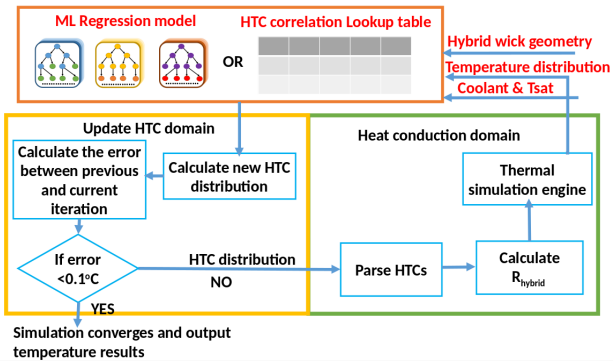

Fig. 3: Temperature-dependent HTC simulation framework and our proposed learning-based temperature-dependent HTC simulation framework.

silicon membrane, a silicon microchannel, and the coolant. Due to device symmetry, we only need to simulate one microchannel and we can additionally limit the domain to two dimensions by considering the channels to be infinitely long. We neglect convection in the liquid phase since the Peclet number is small [8]. Constriction resistance between the substrate and the hybrid wick is accounted for by including 1 μm of the silicon substrate in the simulation domain. The thermal conductivities of the silicon and the working fluid, as well as the resistance to evaporation posed by the liquid-vapor interface, are temperature-dependent properties. Thus, the effective HTC of the hybrid wick is dependent on temperature in addition to geometry and the vapor conditions.

We calculate the HTC by imposing an inward heat flux (q'') at the bottom of the domain and an evaporative boundary condition at the top of the membrane, except where the membrane is supported by the microchannel, which is set as an insulated boundary. The evaporative boundary condition is modeled using a numerical solution to the Boltzmann transport equation, which governs the flux of vapor molecules from the liquid-vapor interface to the far-field vapor [11]. We utilize the numerical Direct Simulation Monte Carlo (DSMC) data and prescribed boundary conditions from a recent work to model the evaporative boundary condition [9]. We extract the average temperature at the substrate-hybrid wick interface (T_b) from the temperature distribution determined by COMSOL, from which the effective HTC of the hybrid wick is calculated by Fourier's Law ($\text{HTC} = q''/(T_b - T_{Sat})$), where T_{Sat} is the temperature of the far-field vapor. We obtain HTC as a function of T_b by imposing a range of heat fluxes on the hybrid wick for each fixed set of geometries and vapor conditions. This COMSOL model has been validated against the experimental results presented in a recent work [5].

C. A Temperature-Dependent HTC Simulation Framework

The HTC of the hybrid wick evaporator is highly dependent on the temperature distribution [8]. To precisely calculate the temperature distributions of the hybrid wick layer and the processing layer, we implement a temperature-dependent HTC simulation framework for the proposed CTM. The simulation flow is shown in Fig. 3. We divide the simulation flow into two domains: (i) the "update HTC" domain, which takes the HTC correlation from a lookup table generated by the COMSOL model for a specific nanoporous membrane and microchannel geometry, and then updates the HTC value for each hybrid wick layer grid cell, based on both the HTC correlation and temperature distribution, and (ii) the "heat conduction" domain, which takes the HTC distribution, calculates the R_{hybrid} for the corresponding hybrid wick layer grid cell, and then carries out the thermal simulation to generate a new temperature distribution to pass it to the "update HTC" domain. The simulation framework iteratively solves for the HTC and temperature distributions until the temperature distribution converges (temperature difference of $< 0.1°\text{C}$).

D. A Learning-Based Temperature-Dependent HTC Simulation Framework

In the original temperature-dependent HTC simulation framework, the HTC correlation of each hybrid wick geometry within a valid range is generated from COMSOL simulations. The naming convention and valid range of hybrid wick geometry parameters are shown in Table I [8]. Fig. 1 (b) shows the structure of a hybrid wick. We store these HTC correlations in a lookup table. Generating a 4096-entry lookup table using COMSOL takes more than 24 hours. If we select 10 cases for each parameter (total 1 million entries), the generation time is 786 days. In this case, the HTC correlation lookup table pre-computing time for the hybrid wick geometry is the bottleneck of this temperature-dependent HTC simulation framework. In addition, even if we have a finer granularity HTC lookup table for one specific geometry range, we cannot run thermal simulations for a comprehensive range of the hybrid wick geometries. To enable thermal simulations for a wide range of valid hybrid wick geometries, we propose a learning-based temperature-dependent simulation framework as shown in Fig. 3. Compared to the original temperature-dependent simulation framework, we add additional machine learning regression models for different coolants to replace the HTC lookup table. The selection of the machine learning regression model depends on the two-phase cooling method as well as the cooling parameters. For two-phase VCs with hybrid wick evaporators, we select the random forest regression model. The inputs of the machine learning regression model are $\{t, dp, \phi, AR, SF, w, T_{Sat} - T_{Cur}, coolant\}$ (see Table I). The output of the machine learning model is an HTC value. During each iteration, the thermal simulation engine generates a new temperature distribution and passes it to the machine learning regression model to predict the HTC value based on the temperature distribution, hybrid wick geometry parameters, coolant saturation temperature, and coolant type.

E. Proposed Thermal Simulation Tool

We build a CTM simulator that models emerging cooling techniques integrated on a silicon chip. Our CTM simulator

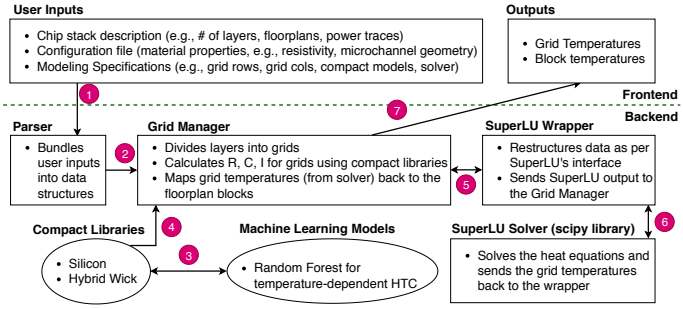

Fig. 4: Proposed Simulation Tool Flow.

works with the same *RC* modeling principle in popular CTM simulators such as HotSpot and 3D-ICE [10], [13]. We support only steady-state simulations in this first version. Our tool has well-defined interfaces that allow easy integration of CTMs for different materials or technologies as libraries, and also includes interfaces for the machine learning models. Fig. 4 shows a high-level tool flow. The user first passes information of the chip stack (including silicon and cooling layers, floorplans, and power traces), material properties (such as thermal resistivity or microchannel geometry), and modeling specifications (e.g., # of grid rows and columns, compact models, solver name) to the tool. Next, the tool parses the input data, performs sanity checks and then bundles the user inputs into various data structures and passes them to a grid manager. The grid manager systematically carries out 3 functions: (i) forming a 3D network of grids in the chip, (ii) redistributing the power traces of the floorplan blocks to the grids, and (iii) calculating the lateral and vertical thermal resistances of the grids using compact libraries. Our proposed hybrid wick CTM (that uses a learning-based temperature-dependent HTC simulation framework) is also integrated with the tool. The grid manager creates resistance and power arrays for all the grids and passes them to the SuperLU solver to solve for grid temperatures.

IV. Optimization of Hybrid Wick Evaporators

The ultimate goal of the optimization flow is to find a hybrid wick geometry that minimizes the hot spot temperatures while satisfying the dry-out constraint. We adopt the dry-out heat flux formula from recent work and use this formula to define the dry-out limit [8]. For each hybrid wick geometry parameter, there is a range of values we can select. It is time-consuming and inefficient to use the grid search to find the optimal geometry and coolant from a fine-grained geometry and coolant solution space. To speed up the searching time for the optimal hybrid wick geometry and coolant, we propose a multi-start simulated annealing (MSA) approach.

Our proposed MSA algorithm is shown in Algorithm 1. The algorithm randomly selects a hybrid wick geometry G and checks whether it satisfies the dry-out constraint by comparing it to the maximum power density PD_{Max} (lines 2-6). If the dry-out heat constraint is satisfied, MSA runs the learning-based temperature-dependent HTC simulation framework to get the maximum chip temperature T_{Max} (line 7). The algorithm then randomly adds a perturbation to one of the geometry parameters and makes a new geometry G_{Nbr} (line 9). Next, MSA checks if G_{Nbr} is in the valid parameter range and and also satisfies the dry-out constraint (lines 11-13). If the constraints are met, MSA runs the learning-based simulation framework for G_{Nbr} to get T_{Max_Nbr} (line 14). If the G_{Nbr} results in a lower peak temperature than G, the algorithm sets G to G_{Nbr} and T_{Max} to T_{Max_Nbr} (lines 15-17). Otherwise, the algorithm sets G to G_{Nbr} based on the probability function (lines 18-20). The algorithm terminates based on the $numstart$, energy E, and decay factor δ (lines 1, 8-9, and 21). It also saves the best G and T_{Max} into an array Opt (lines 22-23). For each type of coolant, we execute this MSA algorithm and select the best geometry and coolant that result in the minimum T_{Max}.

Algorithm 1: Multi-Start Simulated Annealing

```
   Initialize: E, iter, δ, num_start, E_Min
1  while num_start > 0 do
2      randomly select hybrid wick geometry G
3      calculate dry-out heat flux Q_Dry_G [8]
4      if Q_Dry < PD_Max then
5          continue
6      else
7          run thermal simulation using G and get T_Max
8          while E > E_Min and iter > 0 do
9              randomly select a neighbor geometry G_Nbr
10             iter -= 1
11             if G_Nbr in valid parameter range then
12                 calculate dry-out heat flux Q_Dry_Nbr [8]
13                 if Q_Dry < PD_Max then
14                     run thermal simulation using G_Nbr and get
                       T_Max_Nbr
15                     if T_Max_Nbr < T_Max then
16                         G = G_Nbr
17                         T_Max = T_Max_Nbr
18                     else if Random(0,1)< (T_Max_Nbr − T_Max)/(T_Max * E) then
19                         G = G_Nbr
20                         T_Max = T_Max_Nbr
21                     E = E * δ
22         Save G and T_Max into an array Opt
23         num_start -= 1
24 Pick the best G from Opt based on T_Max
```

TABLE I: Hybrid wick geometry parameters and valid range.

Symbol	Parameters	Valid range
t	Nanoporous membrane thickness	250-1000 nm
dp	Membrane pore diameter	50-200 nm
ϕ	Membrane porosity	0.2-0.8
AR	Microchannel aspect ratio	0.5-2
SF	Microchannel wall solid fraction	0.1-0.4
w	Microchannel width	2-8 μm
T_{Sat}	Coolant saturation temperature	$50^{\circ}C$
T_{Cur}	Current temperature of the grid	NA

V. Experimental Results

In this section, we first show the validation results of our learning-based temperature-dependent HTC simulation framework. We then demonstrate our MSA results using various chip floorplans and power profiles against grid search results. We select water, R245fa, and R141b as our coolants and the saturation temperature is set to $50^{\circ}C$. We use the CTM for two-phase VCs with hybrid wick evaporators as our thermal model and show comparisons to COMSOL.

A. Validation of the Machine Learning Model

Before validating the learning-based temperature-dependent HTC simulation framework, we first perform cross-validation of the machine learning regression model to show that our regression model can accurately predict HTC for various hybrid wick geometries and coolants. For each hybrid wick geometry parameter range shown in Table I, we select the minimum value, 25 percentile value, 75 percentile value, and the maximum value as our training and testing geometry

parameters. There is a total number of 4096 geometries. We use COMSOL to generate the temperature-dependent HTC correlations for these 4096 geometries for three different coolants. For each temperature-dependent HTC correlation, we range $T_{Cur} - T_{Sat}$ from 0 to 40°C with a step of 1°C to generate the golden HTC data. The total HTC data size for all the three coolants is 163840. For each coolant, we do k-fold cross-validation to show that our machine-learning regression model is capable of predicting HTC for arbitrary hybrid wick geometries. We use the training data to train a random forest model with a number of trees equal to 100 for each coolant. We also test with various machine learning models including support vector regression (SVR), neural network regression (NNR), decision tree, gradient boosting regression (GDR), etc. Since random forest results in the best accuracy, we only report the regression accuracy results of random forest regression. We show the worst-case scenario results from the k-fold cross-validation for each coolant in Table II. As we can see from the table, our machine learning model can successfully predict HTC for an extensive selection of valid hybrid wick geometries.

B. Learning-based Temperature-Dependent HTC Simulation Framework Validation

Next, we perform an accuracy and speedup comparison among the learning-based temperature-dependent HTC simulation framework, temperature-dependent HTC simulation framework, and COMSOL model using various chip power profiles and hybrid wick geometries. Both the temperature-dependent HTC simulation framework and learning-based temperature-dependent HTC simulation framework are integrated into our proposed thermal modeling tool. We model a 2 mm × 2 mm chip with a thickness of 100 μm in COMSOL, the learning-based framework, and the original temperature-dependent framework. We run two sets of simulations for each floorplan as shown in Fig. 5: (i) processing layer with a uniform power density, and (ii) processing layer with a non-uniform power density with 500 × 500 μm^2 hot spots. Each simulation set uses three different hybrid wick geometries as shown in Table III and three different coolants. We train the machine learning regression model with 4096 geometries described in the previous section. The three validation geometries are excluded from the training set. For uniform power density tests, we use 100 W/cm^2 and 200 W/cm^2. For non-uniform power density tests, we set the background power density to 50 W/cm^2 and hot spots power density to 100, 500, and 1000 W/cm^2 [8]. The COMSOL model uses 428 nodes to compute the temperature distribution, while we use 16×16×3 nodes in our modeling tool to simulate the temperature. We compare the average simulation runtime of COMSOL, temperature-dependent HTC simulation framework, and learning-based temperature-dependent HTC simulation framework. For uniform power density validation tests, the maximum and average errors of the learning-based temperature-dependent HTC simulation framework are less than 0.46 °C and 0.24 °C, respectively. The accuracy results for non-uniform validation tests are shown in Fig. 6. Compared to the COMSOL model, the temperature-dependent HTC simulation framework has a maximum error of 1.34°C with an average speedup of 22×. The maximum error and average error of our proposed learning-based temperature-dependent HTC simulation framework are 2.59°C and 0.98°C, respectively and the average speedup is 21×. Note that the accuracy results include all of the validation floorplans and geometries. Since we replace temperature-dependent HTC correlation lookup table with machine learning regression models, the simulation speedup and accuracy are expected to decrease. However, our proposed model still achieves good accuracy and speedup when compared to the COMSOL model. Most importantly, our proposed learning-based temperature-dependent HTC simulation framework enables accurate thermal simulations with valid and comprehensive geometries.

TABLE II: Worst-case results from the k-fold cross-validation tests. MAE stands for mean absolute error, RMSE stands for root mean square error. The errors are normalized with respect to the golden HTC data.

5-fold				3-fold			
Coolant	MAE	RMSE	R2	Coolant	MAE	RMSE	R2
Water	0.11%	0.16%	99.95%	Water	0.13%	0.19%	99.93%
R245fa	0.41%	0.76%	99.92%	R245fa	1.1%	2.1%	99.91%
R141b	0.17%	0.31%	99.94%	R141b	0.8%	1.2%	99.92%

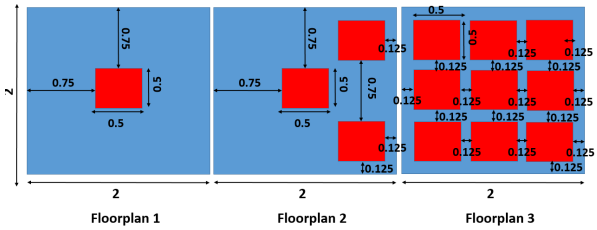

Fig. 5: Floorplans used in validation. Dimensions are in mm.

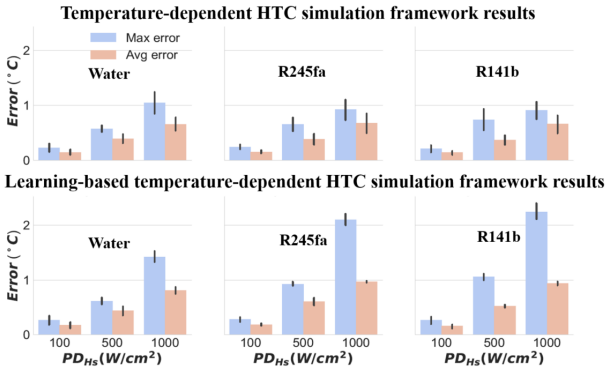

Fig. 6: Non-uniform power profile maximum and average temperature error validation results. PD_{Hs} stands for hot spot power density.

C. MSA vs. A (Coarse) Grid Search

To evaluate the efficiency of our proposed optimization flow, we compare the optimal geometries, coolants, the corresponding peak temperatures, and the searching time of MSA and grid search. For grid search, we select the solution space to be 4096 hybrid wick geometries defined in Section V-A. For each geometry, we first compare the dry-out heat flux of the geometry to the maximum power density PD_{Max}. We then collect all the geometries that satisfy the dry-out constraint and select the one that has the minimum peak temperature. For each coolant, we perform this grid search and pick the optimal coolant and the best hybrid wick geometry. This grid

TABLE III: Hybrid wick geometries for validation tests.

Geometry	t	dp	ϕ	AR	SF	w
1	450	120	0.4	2	0.25	4
2	300	100	0.2	1	0.2	5
3	700	150	0.45	0.6	0.1	8

TABLE IV: Optimization results for grid search and MSA. The optimal hybrid wick geometry is written in a format of $\{t\ (\mu m), dp\ (\mu m), \phi, AR, SF, w\ (\mu m)\}$. Table on the top shows the grid search results.

Floorplan	Hot Spot Power Density (W/cm^2)			
	100	500	1500	2000
1	R245fa$\{0.75, 0.2, 0.4, 1, 0.2, 4\}$	R245fa$\{1, 0.2, 0.6, 1.5, 0.3, 4\}$	R245fa$\{0.75, 0.2, 0.8, 2, 0.3, 4\}$	R245fa$\{1, 0.2, 0.4, 1.5, 0.3, 6\}$
2	R245fa$\{1, 0.05, 0.2, 2, 0.2, 8\}$	R141b$\{1, 0.05, 0.2, 2, 0.2, 8\}$	R141b$\{0.25, 0.05, 0.2, 2, 0.2, 8\}$	R141b$\{0.25, 0.05, 0.2, 2, 0.2, 8\}$
3	R245fa$\{1, 0.05, 0.2, 2, 0.2, 8\}$	R141b$\{1, 0.05, 0.2, 2, 0.2, 8\}$	Water$\{1, 0.2, 0.6, 1.5, 0.2, 6\}$	Water$\{1, 0.2, 0.6, 1.5, 0.2, 6\}$

Floorplan	Hot Spot Power Density (W/cm^2)			
	100	500	1500	2000
1	R245fa$\{0.76, 0.17, 0.44, 1.15, 0.18, 4.57\}$	R245fa$\{0.96, 0.19, 0.42, 1.69, 0.32, 3.53\}$	R245fa$\{0.86, 0.19, 0.47, 1.23, 0.3, 6.11\}$	R245fa$\{0.98, 0.19, 0.5, 1.8, 0.36, 4.96\}$
2	R245fa$\{0.98, 0.05, 0.22, 1.99, 0.19, 7.95\}$	R141b$\{0.99, 0.05, 0.23, 1.88, 0.2, 7.98\}$	R141b$\{0.25, 0.05, 0.23, 1.74, 0.15, 7.94\}$	R141b$\{0.27, 0.08, 0.36, 1.94, 0.12, 7.94\}$
3	R245fa$\{0.99, 0.05, 0.33, 1.99, 0.22, 7.99\}$	R141b$\{0.99, 0.05, 0.2, 1.94, 0.3, 8\}$	Water$\{0.98, 0.19, 0.48, 1.69, 0.19, 6.4\}$	Water$\{0.99, 0.19, 0.49, 1.37, 0.2, 6.8\}$

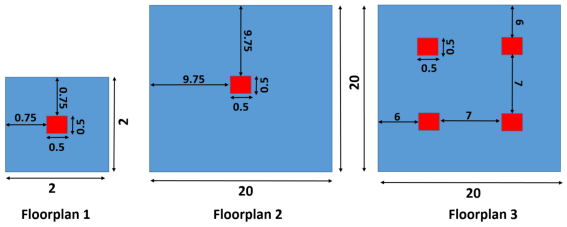

Fig. 7: Experimental floorplans. Dimensions are in mm.

search is coarse-grained because the solution space does not contain all possible valid hybrid wick geometries. As for MSA, we set the initial energy E to 1, the decay factor δ to 0.9, and minimum energy E_{Min} to 0.01. num_start and $iter$ are set to 10 and 100, respectively. We use three different floorplans (see Fig. 7) with a background power density PD_{Bg} of 50 W/cm^2 and hot spot power density PD_{Hs} of $\{100, 500, 1500, 2000\}$ W/cm^2. The results of the grid search and MSA are shown in Table IV. In all of the experiments, our proposed optimization flow selects combinations of coolant and geometry that result in lower temperatures (average 0.67 $°C$ and maximum 1.78 $°C$) compared to the selections made by grid search. Most importantly, the maximum and average searching and simulation time for 4096 solutions grid search are 19 and 6.67 hours, respectively. However, the maximum and average simulation times of our proposed MSA are 2.05 and 1.57 hours, respectively. The maximum and average speedup of our proposed MSA are 9.4× and 4×, respectively, which means our proposed optimization flow is more efficient than the grid search. We also observe that the dry-out limit is highly correlated with the chip size and number of hot spots. If the chip size is larger and there are a larger number of hot spots, only water meets the dry-out constraint. R245fa and R141b generally have better HTC than water, but they suffer from low critical dry-out heat flux. R245fa and R141b can be used as coolants for small-size chips and fewer hot spots.

VI. Acknowledgement

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

VII. Conclusion

In this paper, we introduce a learning-based temperature-dependent HTC simulation framework for two-phase cooling. This framework has the ability to enable accurate thermal simulations for two-phase cooling technology with a wide range of cooling parameters. We build a compact thermal simulator to integrate the proposed two-phase VCs with hybrid wick evaporators CTM into the simulation framework and validate the temperature results against a COMSOL model as well as a temperature-dependent HTC simulation framework designed for the same cooling technology. Compared to the COMSOL model, our simulation framework with CTM achieves a 21× speedup with an average accuracy loss of less than 0.98$°C$ when simulating with three hybrid wick geometries that are excluded from the training set. We also design an optimization flow for our proposed CTM to finds the best hybrid wick geometry and coolant for a given floorplan and power profile. Our future work includes integrating other two-phase cooling CTMs (e.g., microchannel-based two-phase cooling and two-phase VCs with single wick evaporators) into our simulation framework and investigating the tradeoffs between the training set size and the machine learning model accuracy.

References

[1] M. Bulut, S. G. Kandlikar, and N. Sozbir. A review of vapor chambers. *Heat Transfer Engineering*, pages 1–23, 2018.
[2] I. Chowdhury et al. On-chip cooling by superlattice-based thin-film thermoelectrics. *Nature nanotechnology*, 4(4):235, 2009.
[3] A. K. Coskun, J. L. Ayala, D. Atienza, and T. S. Rosing. Modeling and dynamic management of 3D multicore systems with liquid cooling. In *17th IFIP International Conference on Very Large Scale Integration (VLSI-SoC),*, pages 35–40. IEEE, 2009.
[4] B. Dang et al. Integrated microfluidic cooling and interconnects for 2D and 3D chips. *IEEE Trans. on Advanced Packaging*, 33(1):79–87, 2010.
[5] D. F. Hanks et al. Nanoporous membrane device for ultra high heat flux thermal management. *Microsystems & Nanoengineering*, 4(1):1, 2018.
[6] Y. S. Ju et al. Planar vapor chamber with hybrid evaporator wicks for the thermal management of high-heat-flux and high-power optoelectronic devices. *Intl. Journal of Heat and Mass Transfer*, 60:163–169, 2013.
[7] F. Kaplan, S. Reda, and A. K. Coskun. Fast thermal modeling of liquid, thermoelectric, and hybrid cooling. In *Proc. of Intersociety Conf. on Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm)*, pages 726–735, 2017.
[8] Z. Lu et al. Design and modeling of membrane-based evaporative cooling devices for thermal management of high heat fluxes. *IEEE Transactions on Components, Packaging and Manufacturing Technology*, 6(7):1056–1065, July 2016.
[9] Z. Lu et al. A unified relationship for evaporation kinetics at low mach numbers. *Nature Communications*, (2019):1–8, 2019.
[10] K. Skadron et al. Temperature-aware microarchitecture. In *IEEE Proc. of International Symposium on Computer Architecture (ISCA)*, pages 2–13, 2003.
[11] Y. Sone. Kinetic theoretical studies of the half-space problem of evaporation and condensation. *Transport Theory and Statistical Physics*, 29(3-5):227–260, 2000.
[12] A. Sridhar et al. STEAM: A fast compact thermal model for two-phase cooling of integrated circuits. In *IEEE Proc. of International Conference on Computer-Aided Design (ICCAD)*, pages 256–263, 2013.
[13] A. Sridhar, A. Vincenzi, D. Atienza, and T. Brunschwiler. 3D-ICE: A compact thermal model for early-stage design of liquid-cooled ICs. *IEEE Transactions on Computers*, 63(10):2576–2589, 2014.
[14] J. Thome. Heat transfer engineering data book III, 2010.
[15] G. Vaartstra, Z. Lu, and E. N. Wang. Simultaneous prediction of dryout heat flux and local temperature for thin film evaporation in micropillar wicks. *Intl. Journal of Heat and Mass Transfer*, 136:170–177, 2019.
[16] K. Yazawa et al. 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)*, pages 99–106, 2012.
[17] Z. Yuan et al. Modeling and optimization of chip cooling with two-phase vapor chambers. In *Proc. of IEEE/ACM International Symposium on Low Power Electronics and Design (ISLPED)*, 2019.
[18] Z. Yuan et al. Two-phase vapor chambers with micropillar evaporators: a new approach to remove heat from future high-performance chips. In *Proc. of Intersociety Conf. on Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm)*, 2019.