

Fast Thermal Modeling of Liquid, Thermoelectric, and Hybrid Cooling

Fulya Kaplan*, Sherief Reda†, and Ayse K. Coskun*

* Boston University, Boston, MA – {fkaplan3, acoskun}@bu.edu

† Brown University, Providence, RI – sherief_reda@brown.edu

ABSTRACT

Localized hot spots result in elevated on-chip temperatures, significantly limiting the performance, energy efficiency and reliability of today’s processors. For efficient removal of hot spots, using superlattice-based thermoelectric coolers (TECs) in cooperation with microchannel liquid cooling has recently been explored. For the design and evaluation of such hybrid cooling solutions, fast and accurate modeling is essential. In this paper, we present a modeling methodology to account for the complex thermal behavior of TECs and liquid microchannels using compact thermal modeling (CTM). CTM provides a desirable tradeoff between accuracy and speed; thus, it is usually preferred over computationally heavy multi-physics simulations when designing and evaluating thermal management techniques. In this paper, we first describe how to jointly model liquid microchannels and TECs for a hybrid cooling design. We then validate the accuracy of our models by comparing the temperatures obtained from them against the temperatures from COMSOL and 3D-ICE. In comparison to COMSOL, the proposed model provides an average error of $2.07^{\circ}C$ for TECs and $0.36^{\circ}C$ for liquid microchannels, while providing four orders of magnitude faster simulation time. Compared to 3D-ICE, the proposed liquid cooling model achieves $0.02^{\circ}C$ average error. We point out challenges related to integrating a new cooling model into an existing compact thermal simulator and show that modeling decisions can affect the reported temperature by up to $20^{\circ}C$. We conclude our paper by demonstrating an example use case of our proposed model, where we compare the cooling performance of a hybrid cooling design involving microchannels and TECs against a design that adopts liquid cooling only.

INTRODUCTION

Handling high power densities has been a continuing challenge in modern processors. Maintaining safe on-chip temperatures while providing performance improvements will become even harder in next generation processors, as hot spots with power densities reaching $1\text{-}2\text{ kW/cm}^2$ are anticipated [14]. Meeting thermal constraints plays a key factor in processor energy efficiency, since elevated on-chip temperatures degrade performance, increase leakage power, and reduce processor lifetime significantly.

A number of novel cooling solutions have been proposed to overcome this challenge, including microchannel-based liquid cooling [16], [7], [5], TEC cooling [3], [20], [22], fan-cooled heat sinks, and two-phase cooling [14], [21], each serving different target platforms. Recently, hybrid cooling schemes that combine microchannel liquid cooling and TECs on the same platform have been explored for mitigating high density hot spots more efficiently [22], [12], [13]. Figure 1 illustrates an example of such a hybrid cooling design, where a TEC layer is placed above the processing layer and a microchannel liquid cooling layer is placed on top. TECs are placed right on top of high heat flux locations, as they are suitable for removing

localized high density hot spots, but become highly inefficient when cooling down large areas. On the other hand, microchannel liquid cooling can effectively remove background heat on larger chips, but incurs large pumping power when dealing with high heat fluxes. Energy savings of this hybrid scheme has also been experimentally demonstrated [13].

Having fast and accurate thermal modeling tools is essential for the design and evaluation of the future hybrid cooling systems. These models will enable researchers to explore the design space at a broader scale, develop their own optimization techniques easily, and provide a means of fair comparison against the state-of-the-art. There are a number of compact thermal models developed to simulate the behavior of microchannel-based liquid cooling [4], [18], [19], [6], [9] and superlattice-based TECs [11], [20], [22], [3]. However, these models focus on modeling the two cooling methods separately, thus, they do not provide the ability to model a hybrid-cooled system. Shakouri *et al.* demonstrate hybrid cooling benefits in simulation environment [22], where they use compact models for TECs, but represent the effect of microchannel-based liquid cooling by defining a high value of effective heat transfer coefficient at the boundary. Using this simplified way of modeling microchannels, one cannot capture the increase in liquid temperature as it flows through the channel from the inlet to the outlet, leading to loss of accuracy.

Commercial multi-physics simulators such as COMSOL Multiphysics [2] and ANSYS [1] are able to model hybrid cooling with very high accuracy. However, such tools are prohibitively expensive as it takes substantially long time to construct system-specific models, and at runtime, they incur long solution times as well as large memory requirements. Such factors limit the use of multi-physics simulators for modeling hybrid cooling.

To this end, we propose a thermal modeling methodology to simulate the steady-state behavior of hybrid cooling systems

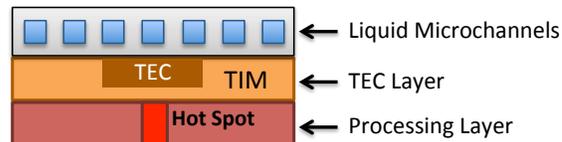


Fig. 1: Front view of an example hybrid design combining microchannel liquid cooling and TECs. TECs are placed on top of high heat flux areas to remove hot spots, while microchannels are used to remove the heat pumped by the TECs and the background heat.

with microchannel liquid cooling and TECs. We adopt a compact thermal modeling approach and integrate our model into a commonly-used simulator, HotSpot [17]. Our model provides a fast and modular way of steady-state thermal evaluation with sufficient accuracy. We validate the accuracy of our TEC model by comparing it against COMSOL Multiphysics software and our liquid cooling model by comparing against both COMSOL and 3D-ICE [18] (which has been previously validated using ANSYS). Finally, we discuss how modeling decisions such as the construction of the thermal resistance network affect the accuracy, or how assumptions about the thermal interface material (TIM) impact fair comparison of cooling designs.

The rest of the paper starts with a review of the existing thermal modeling work in *Related Work* section. We present our modeling method in *Thermal Modeling Methodology* section and provide implementation details. *Validation of the Hybrid Model* section describes the steps for model validation and presents comparison results. *Demonstration of Hybrid Cooling Benefits Using the Proposed Model* section provides a selection of results showing the thermal benefits of hybrid cooling. Finally, we present our conclusions and discuss future directions in *Conclusion* section.

RELATED WORK

This section provides a review of the existing thermal models on TEC cooling, microchannel liquid cooling, as well as hybrid cooling. It then discusses the distinguishing features of our proposed model from prior models.

TEC Cooling: TECs have gained attraction due to their ability to remove hot spots with high power densities. Modeling of TEC thermal behavior is widely studied in the research community [11], [20], [22], [3]. Compact thermal models represent the heat absorbed and rejected on either side of the TEC elements using current sources entering and leaving the thermal nodes [20], [22], [3]. Chowdhury *et al.* compare their numerical compact model against measurements on a test device and show the impact of non-idealities on the cooling potential [3]. Others perform comparison of their 1D analytic TEC model against 3D numerical simulations in ANSYS [22].

Microchannel Liquid Cooling: Liquid cooling with microchannels is an attractive solution for especially 3D-stacked architectures, where the temperature problem is escalated due to vertical layer stacking. Various researchers focus on fast and accurate modeling of the liquid-cooled ICs [4], [18], [19], [6], [9]. Coskun *et al.* incorporate a liquid cooling model into HotSpot-4.01 simulator, where a grid level thermal RC network is constructed and thermal properties of different interlayer materials (i.e., TSVs, microchannels) are specified [4]. Sridhar *et al.* propose 3D-ICE [18], which has the ability to model the thermal gradient between the inlet and outlet ports introduced by the flow of the liquid. They validate 3D-ICE against ANSYS CFX computational fluid dynamics (CFD) tool. The follow-up of this work [19] adds the support for modeling enhanced heat transfer geometries such as pin-fin structures. This model also simplifies the computation in

the microchannel layers by homogenizing them into porous medium. Another body of work focuses on speeding up the long simulation time observed in liquid-cooled ICs [6], [9]. ICTherm [6] is a recently introduced simulator that implements an efficient algorithm to compute the transient temperature in linear-time complexity in liquid-cooled ICs. Other researchers [9] tackle the long simulation time problem by using GPU-accelerated generalized minimum residual (GMRES) method and provide one or two orders of magnitude speedup compared to single-threaded CPU-GMRES method.

Hybrid Cooling with TECs and Liquid Microchannels:

Hybrid cooling with TECs and liquid microchannels has been proposed as an energy-efficient solution for mitigating high density hot spots [22], [12], [13]. Sahu *et al.* show the thermal benefits and characterize the behavior of such hybrid cooling scheme on an experimental setup incorporating on-chip TEC units and a microchannel heat sink [13]. Other work rely on compact models to demonstrate the cooling energy savings of a hybrid solid-state and microfluidic cooling system over solely using microfluidic cooling [22], [12]. They use the aforementioned compact models for TEC modeling, and represent the effect of microchannel-based liquid cooling using a high effective heat transfer coefficient at the boundary. This is a simplified way of modeling hybrid cooling and it does not consider important aspects of liquid cooling, such as the rise of coolant temperature as it flows from the inlet to the outlet. Such aspects become critical when, for example, exploring the impact of hot spot locations on the resulting cooling power. Hot spots that are located closer to the outlet of the microchannels get hotter than the ones that are closer to the inlets, and failing to model this effect results in optimistic evaluation of systems.

Distinguishing Aspects of Our Work: Our work is the first to develop a compact hybrid thermal model for the design and evaluation of systems using TECs and liquid microchannels with sufficient detail and modularity. We integrate our model into an open source thermal simulator (i.e., HotSpot [17]) that is commonly used in the research community. Our work contributes the following improvements over the existing models:

- 1) When modeling liquid microchannels in a hybrid cooling environment, our model avoids assumptions such as representing the liquid cooling layer with a heat transfer coefficient.
- 2) Our model is applicable to a wide range of platforms and applications.
- 3) It is modular in the sense that the users can plug-in the cooling elements (TECs, microchannels, or both) with desired size, properties, and granularity. In this way, the proposed model enables researches to explore a wider design space in a fast and accurate manner compared to existing models.
- 4) Compared to using COMSOL, our compact model provides sufficient accuracy while saving considerable amount of time and processing resources.

THERMAL MODELING METHODOLOGY

Our goal is to develop a thermal model that provides a fast, accurate, and modular way of simulating hybrid cooling systems. We adopt a compact thermal modeling approach for implementing our hybrid cooling model. In this approach, temperature is modeled based on an equivalent RC network of thermal nodes, where R and C correspond to thermal resistance and thermal capacitance, respectively. The individual node temperatures are found by solving this network using differential equation solvers. We implement our proposed hybrid thermal model in HotSpot-6.0 [17] temperature simulator. We begin this section by providing background on the HotSpot simulator. We continue with a detailed description of our compact thermal model and explain how we integrate it into HotSpot.

Background on the HotSpot Simulator

HotSpot is a thermal simulator that has been developed to model processor temperature using compact thermal modeling approach. HotSpot models lateral and vertical heat flows, and it also includes a processor package model to represent the heat spreader and a heat sink with fan. The 3D stacking feature of HotSpot allows the user to define multiple layers with desired properties, such as silicon and TIM layers. The inputs to the simulator are (i) the physical geometry of the chip stack, which includes the width and length of the processor and packaging layers as well as the thickness of each stacked layer, (ii) the floorplan of each layer, which specifies the size and location of rectangular blocks constituting the layer, (iii) the thermal properties of the materials used in the layers, and (iv) the power dissipation of the blocks. Based on the input parameters, the simulator constructs a 3D RC network representation of the chip stack. This representation is then converted into a set of matrix equations in the form of $\mathbf{GT}(t) + \mathbf{CT}\dot{\mathbf{T}}(t) = \mathbf{U}(t)$, where \mathbf{G} and \mathbf{C} are the thermal conductance and capacitance matrices, \mathbf{U} is the power dissipation matrix, and \mathbf{T} is the matrix of node temperatures to solve for. Solving this system of matrix equations using a differential equation solver gives the temperature of each node.

There are two different simulation modes in HotSpot: steady-state and transient simulation. The steady-state simulation gives the stable temperature that the system will reach after running under the same condition for a long time. As there is no time dependence in steady-state (i.e., $\dot{\mathbf{T}}(t) = \mathbf{0}$), the system is represented as a thermal R (resistance) network and the set of matrix equations take the form of $\mathbf{GT}(t) = \mathbf{U}(t)$. For steady-state simulation, HotSpot includes both an iterative solver and a SuperLU-based matrix solver, which provides higher speed at the cost of using more memory resources.

Recently, Meng *et al.* [10] have added the functionality to model *heterogeneity within each layer* in HotSpot such that the user can assign different thermal properties to individual blocks residing on the same layer (e.g., copper TSVs going through a TIM layer). This feature is included in the most recent release (i.e., HotSpot-6.0), and we make use of this feature when implementing our proposed hybrid model.

In this work, we design a steady-state thermal model to quantify the behavior of microchannel liquid cooling and TEC cooling. We implement our model using the *grid model* and the SuperLU-based matrix solver in HotSpot-6.0. We propose a hybrid cooling model allowing the user to define microchannel-based liquid cooling together with TECs on the same platform and provide example chip stacks throughout the paper for demonstration. Our models are modular as they enable the user to select a single or a combination of the cooling methods, configure the number and ordering of the cooling layers in the stack, as well as the location of the microchannels and TECs on the floorplan. For example, one can model a hybrid system with TECs and fan-cooled heat sink, a system with liquid cooling and TECs, a system that adopts liquid cooling only, or any other combination of these cooling methods.

Proposed TEC Model

A TEC operates based on the Peltier effect such that when current passes through the device, heat is absorbed from one side (cold side) and rejected to the other side (hot side), creating a thermal gradient across the two sides [22], [3]. The amount of heat removed by the TEC depends on the Seebeck coefficient (S), applied current (I), electrical resistivity (ρ_{tec}), thermal conductivity (k_{tec}) of the TEC device, and the temperatures of the hot (T_h) and cold (T_c) sides. Superlattice-based thin film TECs made of Bi_2Te_3 have high figure-of-merit (ZT). They are silicon micro-fabrication compatible and can be directly integrated or fabricated on the back of a silicon chip [3], [13]. Existing on-chip TEC devices are composed of ultra-thin (5-10um) Bi_2Te_3 -based p-n thermocouples sandwiched between copper mini-headers and are covered with ceramic insulator plates at the outmost surfaces [3].

There are three main contributors to heat flow within a TEC unit: (i) the Peltier term which accounts for the heat absorbed/rejected on the cold/hot sides, (ii) the conductive heat flow term, and (iii) Joule heating term that represents the resistive heat generated by passing current through the TEC. Mathematical representation of these terms are:

$$Q_c = N(SIT_c - \frac{T_h - T_c}{R_t} - \frac{1}{2}I^2R_e) \quad (1)$$

$$Q_h = N(SIT_h - \frac{T_h - T_c}{R_t} + \frac{1}{2}I^2R_e) \quad (2)$$

where Q_c and Q_h stand for the heat absorbed and rejected on the cold and hot sides, respectively. T_c and T_h are the cold and hot side temperatures. N is the number of p-n couples placed in area A . $R_t = h_{tec}/k_{tec}A$ is the thermal resistance and $R_e = \rho_{tec}h_{tec}/A$ is the electrical resistance of a TEC unit of thickness h_{tec} and area A .

We implement this model in HotSpot in the following way. We use the *grid model* in HotSpot, in which, each layer on the processor stack is divided into smaller grid cells representing a thermal node in the thermal R network. We add functionality to define a block on the floorplan as a TEC unit. We then assign TEC thermal properties only to the grid cells corresponding to these TEC units. For this purpose, we use the heterogeneous

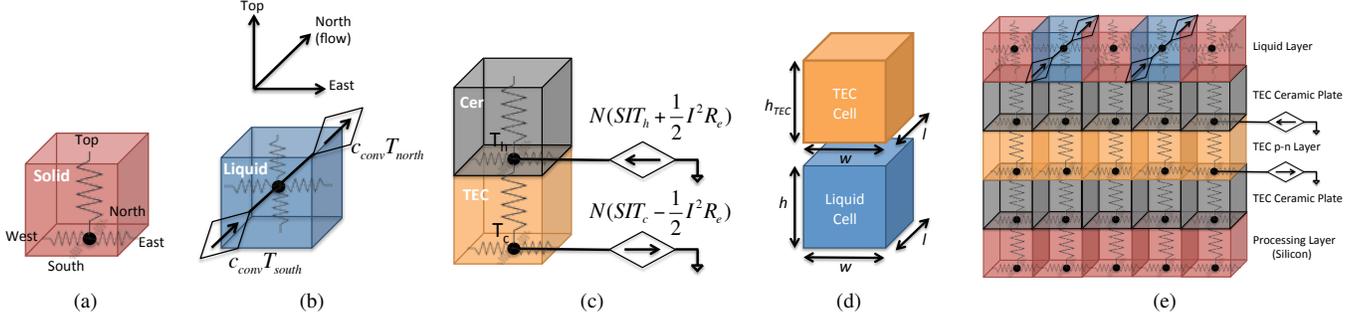


Fig. 2: (a) Solid grid cell, (b) Liquid grid cell, (c) TEC grid cell, (d) Dimensions of the grid cells, (e) Connectivity of the grid cells building a chip stack. Current sources are shown only for the rightmost TEC and ceramic cells for clarity.

3D modeling feature of HotSpot as mentioned earlier. HotSpot by default accounts for the conductive heat flow (term (ii)) for solid cells as shown in Figure 2(a). In order to represent the Peltier term and Joule heating term on the cold and hot side of the TEC units described in Equations (1) and (2), we define current sources entering and leaving the TEC cells as illustrated in Figure 2(c). In the Figure, bottom surface of the TEC cell corresponds to the cold side temperature, while the bottom surface of the cell in the upper adjacent layer (i.e., the ceramic plate) corresponds to the hot side temperature.

Proposed Microchannel Liquid Cooling Model

We adopt the 4 resistor model-based (4RM) liquid cooling model presented in 3D-ICE [18]. In the 4RM-based model, the discretization of the thermal grids is done such that the entire cross-section of a microchannel forms a liquid grid cell. There are two main contributors to heat flow regarding a liquid grid cell: (i) convective heat transfer from the walls of the channel to the liquid and (ii) convective heat transfer in the direction of the liquid flow into and out of the current liquid cell. Figure 2(b) illustrates a liquid grid, where the term (i) is represented by resistive elements in four directions and the term (ii) is represented by using current sources in the direction of the flow (from South to North). The numerical values of the resistances are given as follows [18]:

$$R_{top,bottom} = \frac{1}{h_{f,vertical} \cdot w \cdot l} \quad (3)$$

$$R_{east,west} = \frac{1}{h_{f,side} \cdot h \cdot l} \quad (4)$$

where $h_{f,vertical}$ and $h_{f,side}$ are the heat transfer coefficients for microchannel forced convection; w , l , and h are the width, length, and height of the microchannel cell, respectively (See Figure 2(d) for the cell dimensions.). As also stated in 3D-ICE work [18], $h_{f,vertical}$ and $h_{f,side}$ (i.e., the vertical and side heat transfer coefficients) can be obtained from empirical correlations or numerical presimulation for a given system. For computing the heat transfer coefficients, prior work provides the following formulas assuming imposed axial heat flux and radial isothermal conditions:

$$h_{f,vertical} = h_{f,side} = \frac{k_{coolant} \cdot Nu}{d_h} \quad (5)$$

$$Nu = 8.235 \cdot (1 - 2.0421AR + 3.0853AR^2 - 2.4765AR^3 + 1.0578AR^4 - 0.1861AR^5) \quad (6)$$

In these formulas, $k_{coolant}$ is the thermal conductivity of the coolant and $d_h = \frac{2h \cdot w}{h+w}$ is the hydraulic diameter of the channel. Nusselt number (Nu) was derived in prior work [15] as a function of channel aspect ratio ($AR = \min\{h/w, w/h\}$). As Equations (5) and (6) may differ under different system assumptions, the original 3D-ICE simulator defines $h_{f,vertical}$ and $h_{f,side}$ as input parameters specified by the user.

Next, the values of the convective terms in the flow direction (i.e., the current sources) are computed as follows:

$$I_{in} = c_{conv} \cdot T_{south} \quad (7)$$

$$I_{out} = c_{conv} \cdot T_{north} \quad (8)$$

$$c_{conv} = C_v \cdot u_{avg,y} \cdot \Delta A_y \quad (9)$$

where I_{in} and I_{out} denote the convective heat flow into and out of the cell, respectively. T_{south} and T_{north} are the interface temperatures at the south and north surfaces of the cell. C_v is the specific heat capacity of the coolant, $u_{avg,y}$ is the average coolant velocity, and $\Delta A_y = w \cdot h$. The surface temperatures are approximated as the average of the cell temperatures which share that interface. We assume that for the southmost cell, $T_{south} = T_{inlet}$ (i.e., temperature of the coolant at the microchannel inlet) and for the northmost cell $T_{north} = T_{cell}$.

Note that by default, HotSpot places the virtual temperature node at the bottom surface of the grid cell in the vertical direction as illustrated in Figure 2(a). This convention is useful for modeling the TEC cells as the thermal effect is observed at the bottom and top surface of the TEC device. However, for liquid cells, we need to place the virtual node in the middle of the cell to be able to include the heat flow from the top/bottom walls in an accurate manner. Doing otherwise results in underestimation of the chip temperature by up to $20^\circ C$ for liquid-cooled systems, according to our analysis (Refer to *Placement of the Virtual Thermal Node* section for more detail). Thus, we construct the thermal resistance network in our model such that for liquid cells, the node is placed in the middle; while for all other cells including TECs, the node

is placed at the bottom surface. This way of constructing the thermal resistance network is one of our novel contributions. In Figure 2(e), we demonstrate how the grid cells of each type are connected in the chip stack building a thermal R network, for a single row of cells.

VALIDATION OF THE HYBRID MODEL

Validation of the TEC Model

In order to validate our TEC model, we compare its temperatures against the ones obtained from COMSOL simulations. For this purpose, we first select a prototype TEC device that has been fabricated on the back of a silicon chip and has been characterized in prior work [3]. We then create a model of this TEC device in COMSOL using the *heat transfer module*.

Validation Using COMSOL. The TEC device and the chip layers we model in COMSOL are illustrated in Figure 3. It is a superlattice-based thin film TEC made of Bi_2Te_3 as the bulk material and has high intrinsic figure-of-merit (ZT) [3]. TEC is composed of an array of 7×7 p-n thermocouples and has a total size of $3.5\text{mm} \times 3.5\text{mm}$. Thermocouples are sandwiched between copper mini-headers and the top and bottom surface of the device is covered by ceramic plates to provide electrical insulation. Legs of the p-n thermocouples are ultra-thin ($8\mu\text{m}$) and the total thickness of the TEC device including the ceramic plates is $100\mu\text{m}$. Since the length and width of the thermocouple legs were not specified in prior work, we estimated them such that the 7×7 array fits in the $3.5\text{mm} \times 3.5\text{mm}$ area. Based on this estimation, the leg width and leg length are $400\mu\text{m}$ and $150\mu\text{m}$, respectively. This corresponds to 0.833 p-n thermocouples per mm^2 area and it is used when calculating the parameter N (i.e., the number of p-n thermocouples per TEC area) in the proposed model. Detailed parameters of the TEC are given in Table I. Note that for the temperature dependent parameters such as S , ρ_{tec} and k_{tec} , we assume constant values at steady-state temperature as reported in prior work [3]. We use the reported thermal conductivity, k_{tec} , for calculating the vertical thermal resistance. Since there is air between the p-n thermocouples, lateral heat transfer within the TEC unit is minimal. Thus, we assign a very large number to the thermal resistance in the horizontal direction for the TEC device. For all other layers, we include the lateral heat flow based on the corresponding material properties.

Next, we model the processing layer using a $100\mu\text{m}$ -thick silicon layer at the cold side of the TEC, and assign a heat flux value (i.e., power dissipated per unit area) to it to represent the generated heat. As TECs pump heat from the cold side to the hot side, an additional cooling mechanism is usually needed on the hot side of the TEC to avoid overheating and provide proper operation. Thus, at the hot side of the TEC, we define another layer, which represents the chip package and an additional cooling mechanism (e.g., heat sink with fans, cold plates) that removes the heat pumped by the TEC. We assume silicon properties for this layer, set its thickness as $40\mu\text{m}$, and assign a heat transfer coefficient (htc) at the surface to the ambient to

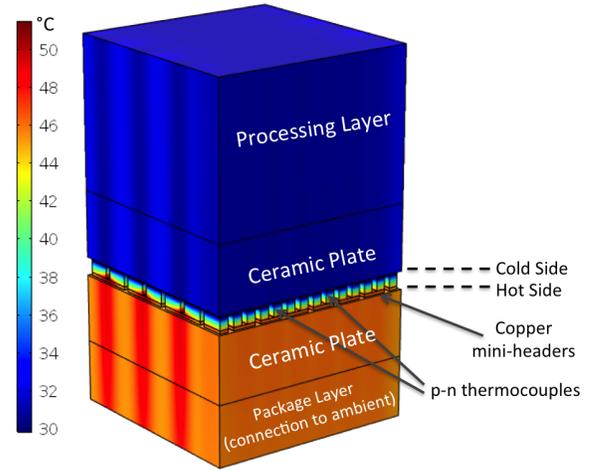


Fig. 3: TEC device that we modeled in COMSOL. Example temperature distribution corresponds to a bias current of 4A.

TABLE I: The parameters we used for the liquid microchannel and TEC models.

Microchannel height	h	$100\mu\text{m}$
Microchannel width	w	$50\mu\text{m}$
Grid cell width & length	$w = l$	$50\mu\text{m}$
Microchannel length	L	10mm
Coolant thermal conductivity	$k_{coolant}$	0.6069W/mK
Coolant specific heat	C_v	4181J/kgK
Coolant inlet temperature	T_{inlet}	27°C
Coolant density	$\rho_{coolant}$	998kg/m^3
Coolant viscosity	μ	$8.89 \times 10^{-4}\text{Pa.s}$
Average coolant velocity	u_{avg}	$< 3\text{m/s}$
TEC width & length	$w_{tec} = l_{tec}$	3.5mm
Seebeck coefficient	S	$301\mu\text{V/K}$
Thermocouple thickness	h_{tec}	$8\mu\text{m}$
Copper mini-header thickness	h_{Cu}	$2\mu\text{m}$
Ceramic plate thickness	h_{Cer}	$44\mu\text{m}$
TEC electrical resistivity	ρ_{tec}	$1.08 \times 10^{-5}\text{Ohm.m}$
TEC thermal conductivity	k_{tec}	1.2W/mK
Copper thermal conductivity	k_{Cu}	400W/mK
Ceramic thermal conductivity	k_{Cer}	175W/mK
Silicon thermal conductivity	k_{Si}	130W/mK

represent the additional cooling mechanism. Htc corresponds to the cooling capability of the additional cooling method with a higher number representing more effective cooling. We modify HotSpot's package model so as to define a similar layer with connection to ambient using the htc parameter.

In COMSOL, we model the TECs in detail by defining the individual p-n legs, the copper mini-headers connecting the thermocouples in series, the VDD and ground nodes one by one. In the proposed model, we define the TEC device as a block, where the details of individual p-n legs and the empty space between them are omitted for the sake of simplicity and speed. In order to account for the differences introduced by the simplifications, we calibrate our proposed model empirically based on COMSOL. Based on our experiments, we observe that such effects demonstrate themselves as a scaling factor on the electrical resistivity of the TEC, which we experimentally determine as 14.

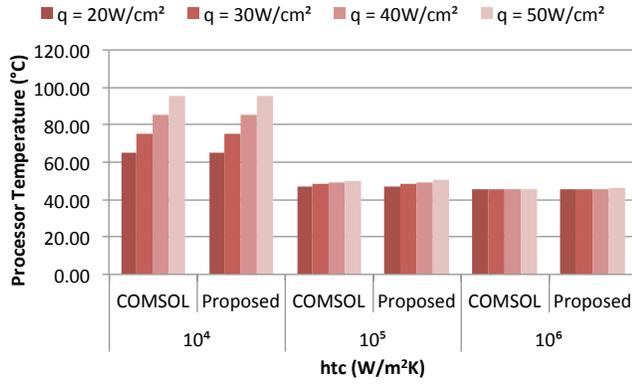


Fig. 4: Comparison of processor layer temperature for the case without TECs and with varying heat transfer coefficient (htc) and heat flux (q) values.

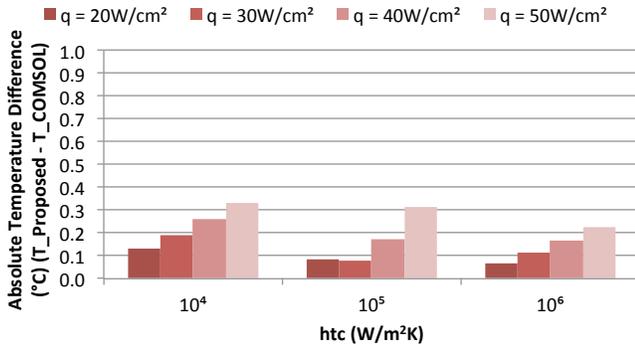


Fig. 5: Absolute temperature error for the case without TECs and with varying heat transfer coefficient (htc) and heat flux (q) values.

We run two sets of experiments in COMSOL: (i) without the TEC device for varying htc and heat flux (q) levels, and (ii) with the TEC device using a bias current changing from 0 to 7A with varying q levels. For the case with TECs, we define a multi-physics problem, which combines *heat transfer in solids* with *thermoelectric effect*, *electromagnetic heat source* and *thermal coupling* elements. We use the segregated solver in COMSOL to solve the multi-physics problem iteratively using GEMRES method for both sub-parts of the problem. The resulting mesh consists of 164088 domain elements, 125204 boundary elements and 16422 edge elements. Number of degrees of freedom solved for is 1810396.

For the rest of the paper, we will refer to the results corresponding to our proposed hybrid model as *proposed*. Figure 4 compares the temperature of the processor layer for the case without TECs, and Figure 5 reports the absolute temperature difference between the proposed model and COMSOL for the processor layer. As seen from the figure, there is a good match between the two simulators with an absolute error of less than $0.5^{\circ}C$ across all htc and q combinations.

Next, we present the comparison results for the case with TECs. Figure 6 compares the average temperature of the

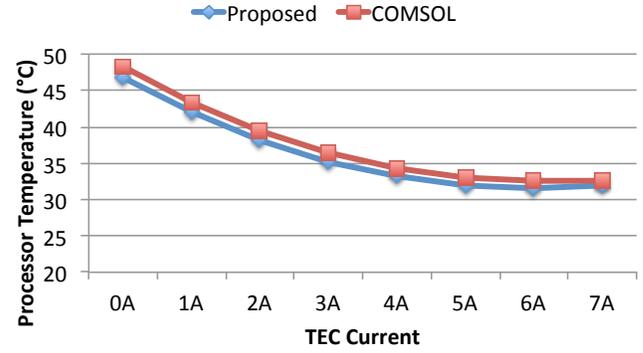


Fig. 6: Comparison of processor temperature over TEC current for COMSOL and the proposed model. $htc = 10^6 W/m^2 K$ and $q = 20 W/cm^2$.

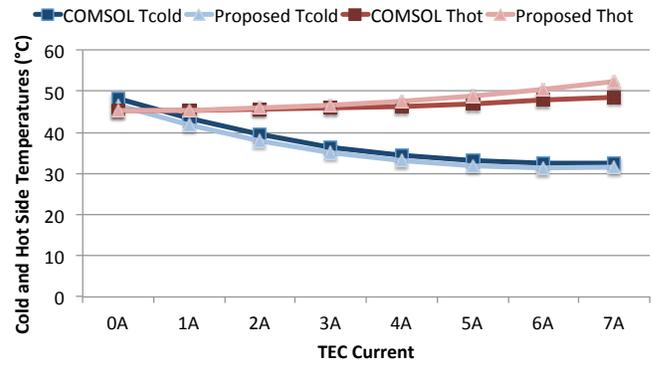


Fig. 7: Comparison of the cold and hot side temperatures over TEC current for COMSOL and the proposed model. $htc = 10^6 W/m^2 K$ and $q = 20 W/cm^2$.

processor layer over a range of TEC bias currents. For this experiment, $htc = 10^6 W/m^2 K$ and $q = 20 W/cm^2$. Our proposed TEC model closely follows the temperature results obtained from COMSOL with an error less than $1.5^{\circ}C$. As expected, the processor temperature starts to reduce as the TEC bias current increases. At some point (i.e., around 6A), impact of Joule heating becomes dominant, resulting in a slight increase in the processor temperature. In Figure 7, we report the cold and hot side temperatures of the TEC for the same simulation. At 0A bias current, $T_{cold} > T_{hot}$ due to the additional resistance presented by the TEC device. At around 0.5A, amount of heat that is pumped by the TEC overcomes its own resistance and $\Delta T = (T_{hot} - T_{cold})$ becomes positive and starts to increase.

In Figure 8, we compare the thermal maps obtained from the two simulations for $q = 20 W/cm^2$ and TEC current of 4A. The plots on the left correspond to the temperatures of the processing layer, while the plots on the right show the temperatures on the hot side. We carry out similar analysis for other q values ranging from 20 to $50 W/cm^2$ and observe that the absolute maximum error is $3.57^{\circ}C$. We also report $2.07^{\circ}C$ of average and $2.25^{\circ}C$ of RMS error.

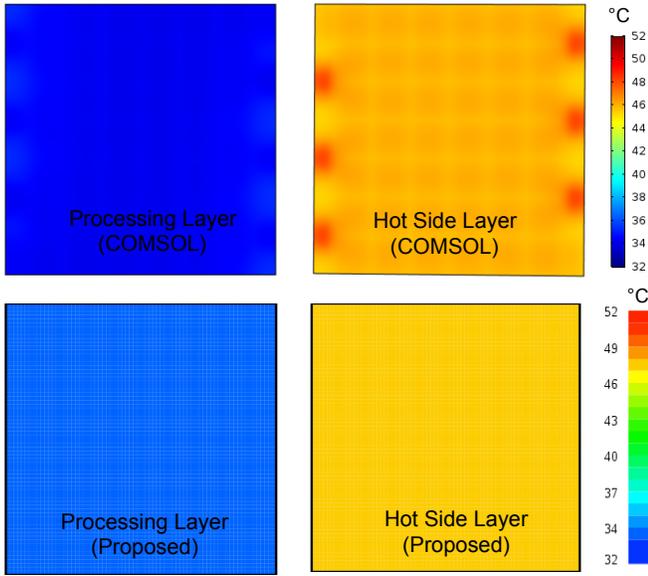


Fig. 8: Comparison of thermal maps corresponding to the processing layer and the TEC hot side, for COMSOL and the proposed model. $h_{tc} = 10^6 W/m^2K$, $q = 20 W/cm^2$ and $I = 4A$.

Effects of TIM Assumptions on the Fairness of Comparison Among Cooling Designs. Selection of the TIM properties is highly important for fair comparison of different cooling designs. Next, we will describe an example case from prior work and show how the assumptions on the TIM thickness and properties may lead to overestimation of TEC benefits.

Prior work demonstrates the benefits of TECs regarding the removal of high density hot spots [3]. There is a very small hot spot area placed at the center of the processing layer. The size of the hot spot is $400\mu m \times 400\mu m$ and hot spot heat flux is $q = 1.25 kW/cm^2$, while the background heat flux is $q_{bgnd} = 42.7 W/cm^2$. On top of the processing layer, there is TIM followed by a top packaging layer representing the heat sink. In order to demonstrate the benefits of TECs, two cases are compared: (i) chip stack with processor, $125\mu m$ TIM, and the package layers, and (ii) chip stack with processor, $25\mu m$ TIM, $100\mu m$ TEC layer with a TEC unit placed above the hot spot, and the package layers. The TIM conductivity was assumed as $1.75 W/mK$.

Prior work [3] claims that by simply adding the TEC layer, even at $0A$ bias current, there would be a passive cooling effect introduced by higher thermal conductivity of the TEC material. This claim is true, if we assume a $125\mu m$ -thick TIM as the baseline without TECs (let us call this baseline #1). However, in a system without TECs, a much thinner TIM, i.e., one with $25\mu m$ thickness, can be utilized (let us call this baseline #2). Moreover, there are TIM materials with much higher reported thermal conductivities [8]. Using our simulation framework, we evaluate the results of each assumption. As the heat sink properties were not specified in prior work, we assign $h_{tc} = 10^6 W/m^2K$ without loss of generality to represent

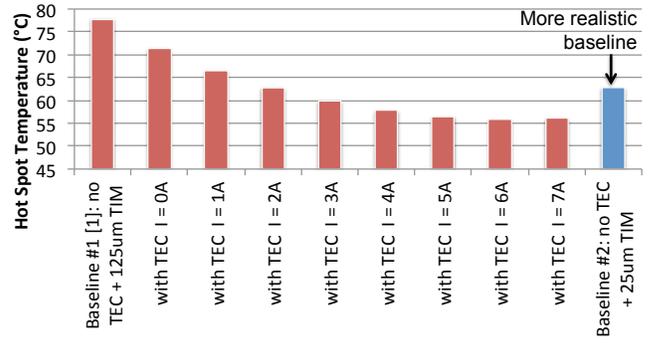


Fig. 9: Comparison of the hot spot temperatures for pessimistic baseline #1 from prior work [3], a more realistic baseline #2, and a system with TECs using different bias currents.

the heat sink. We assume a higher quality TIM material from recent work [8] with $8.5 W/mK$ conductivity. In Figure 9, we compare baselines #1 and #2 against the system with different TEC bias currents and report the maximum temperatures. As seen from the figure, baseline #1 results in about $15^\circ C$ higher temperature compared to a more realistic baseline #2. In fact, when we add TECs and do not apply any bias current, we are introducing additional thermal resistance which increases the temperature by $9^\circ C$ compared to the more realistic baseline #2. However, if one assumes the very thick TIM from baseline #1, it seems like TEC is providing cooling even without being activated, which leads to overestimation of its benefits. For the hot spot heat flux we have in this experiment, TEC starts to provide benefit over the baseline #2 only after $2A$ of bias current.

We think that such assumptions on the TIM thickness can affect conclusions when comparing two different cooling designs, thus, are highly important.

Validation of the Liquid Cooling Model

We validate our microchannel liquid cooling model by comparing it against two different simulators: (i) COMSOL Multiphysics tool, and (ii) 3D-ICE [18] simulator, which has been well validated against ANSYS CFX tool. During validation of the 3D-ICE simulator, two different chip stacks were modeled: (i) two active dies and one microchannel layer in between them and (ii) three active dies and four microchannel layers placed between them. Experiments with various flow rates and heat flux profiles have been carried out and a maximum temperature error of $1.5^\circ C$ was reported.

Validation Using COMSOL and 3D-ICE. For validation of our proposed model in COMSOL, we first create a chip stack with liquid microchannels. Figure 10(a) illustrates the cross-section of the chip stack, where the liquid microchannel layer is placed on top of the processor layer, and an additional bulk silicon layer (with $40\mu m$ thickness) is placed on top to provide closure to the microchannels. We simulate a thin slice of this chip stack as in prior work [18] to reduce the problem size in COMSOL (See Figure 10(b)). The width and length of the slice

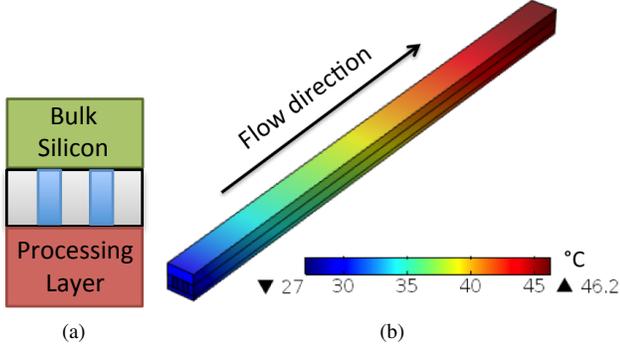


Fig. 10: (a) Front view of the thin slice of chip stack we modeled for liquid cooling, (b) Side view of the chip stack as we modeled in COMSOL.

are $250\mu\text{m}$ and 5mm , respectively. We set the microchannel width as $w=50\mu\text{m}$ (also equal to the channel wall width) and channel thickness as $h=100\mu\text{m}$. With these microchannel parameters, the simulated slice includes two microchannels interleaved between three channel walls made of silicon. At the top surface of the bulk silicon layer, we assign a very small heat transfer coefficient (i.e., $htc = 0.01\text{W}/\text{m}^2\text{K}$) to represent minimal convection to air. We assume water as the coolant and use the coolant properties given in Table I.

Similar to the case with TECs, the problem we define in COMSOL is a multi-physics problem, which combines *heat transfer in solids*, *heat transfer in liquids*, and *laminar flow* elements. We use the segregated solver in COMSOL to solve the multi-physics problem, where the segregated step 1 (corresponding to the laminar flow) is an iterative solver using GEMRES method, and the segregated step 2 (corresponding to heat flow) is a direct solver using PARDISO method. We construct a fine mesh, which consists of 628237 domain elements, 66162 boundary elements and 4332 edge elements. Number of degrees of freedom solved for is 514554.

We model the same chip stack in 3D-ICE simulator for the second set of comparisons. As the computation of $h_{f,vertical}$ and $h_{f,side}$ coefficients significantly differ in COMSOL and 3D-ICE, we first experimentally estimate the coefficients from COMSOL simulations and then use them as inputs to the proposed model and 3D-ICE simulator. This way, we can carry out a consistent comparison of the three models. We extract the coefficients from COMSOL as follows: to find $h_{f,side}$, we select the surface of a side wall facing a microchannel and record the surface average of the total normal heat flux value ($ht.ntflux$ in COMSOL). We then record the surface average of the side wall temperature (T_{wall}), and the volume average of the liquid temperature (T_{liquid}). Finally, we compute $h_{f,side}$ using the equation below:

$$h_{f,side} = \frac{ht.ntflux}{(T_{wall} - T_{liquid})} \quad (10)$$

We carry out similar computation for $h_{f,vertical}$ using the top

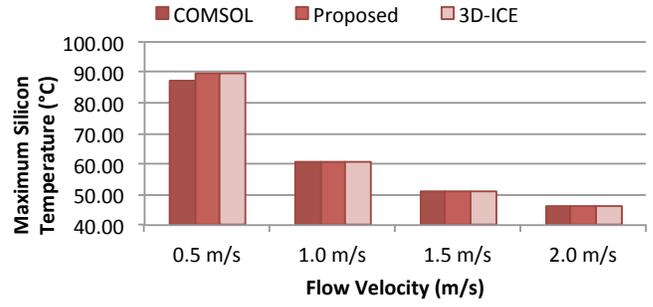


Fig. 11: Maximum processor temperature comparison for COMSOL, 3D-ICE and the proposed model for $q = 100\text{W}/\text{cm}^2$.

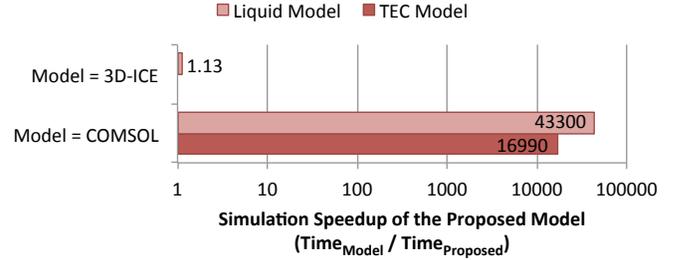


Fig. 12: Comparison of the simulation speed across three simulators. As 3D-ICE does not have a TEC model, the bar is not shown.

and bottom walls. We repeat the same steps for the flow velocities that we experiment with and assign the average computed value to the heat transfer coefficients. For our system, we determine that $h_{f,side} \approx h_{f,vertical} = 1.05 \times 10^5 \text{W}/\text{m}^2\text{K}$. We use these values as inputs to the proposed model and 3D-ICE simulator.

We run steady-state simulations for a range of q values of $12.5, 25, 50,$ and $100\text{W}/\text{cm}^2$ as well as for different flow velocities, $u_{avg} = 0.5, 1.0, 1.5, 2.0\text{m}/\text{s}$, and record the maximum temperature of the processing layer for the proposed model, COMSOL, and 3D-ICE. Figure 11 shows the maximum processor temperatures obtained from COMSOL, 3D-ICE, and our proposed model for all u_{avg} combinations where $q = 100\text{W}/\text{cm}^2$. Among all experiments, compared to COMSOL simulations, our proposed model provides maximum, average and RMS error of 2.46°C (corresponds to 2.8%), 0.36°C , and 0.72°C , respectively. In comparison to 3D-ICE simulator, the error of the proposed model is less than 0.04°C .

Finally, we compare the solution speeds of the simulators against the proposed model for both TEC and liquid cooling. Figure 12 demonstrates the average solution time ratio of the compared simulators over the proposed model. As indicated by the figure, the proposed compact modeling approach can save significant simulation with reasonable tradeoff in accuracy.

Placement of the Virtual Thermal Node. An important aspect of modeling hybrid cooling is related to where to place the virtual thermal nodes on the grid cells while constructing the thermal resistance network. As we briefly discussed in prior

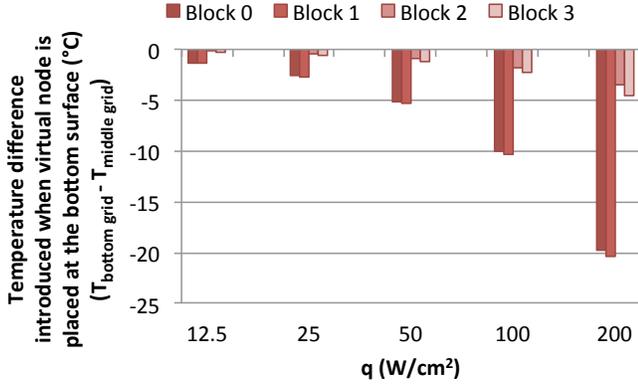


Fig. 13: Temperature difference introduced when the virtual node is placed at the bottom surface of a liquid cell. Placing the virtual node at the bottom surface results in underestimation of the temperature by up to $20^\circ C$.

sections, HotSpot simulator by default places the virtual node at the bottom surface of a grid cell as shown in Figure 2(a). This is very convenient for TEC modeling, where we focus on either side of the TEC cell (i.e., cold and hot sides) when applying Kirchhoff's current law at the nodes and inserting the current terms into the equation (see Figure 2(c)). However, for the liquid cooling model, we have found out that this approach results in significant underestimation of processor temperatures. This is because when modeling the temperature of the liquid cells, one should account for both the heat transferred from the solid cell (conduction) to the walls and from the walls to the liquid (convection). When the virtual node is placed at the bottom surface, the vertical heat transfer from the cell above is fully attributed to convection, whereas the heat transfer from the bottom cell is fully attributed to conduction (instead of a combination of them from each direction). This asymmetric representation of the resistances creates an effect as if liquid absorbs more heat from the processing layer than it actually does. This assumption also affects the convection in the direction of the flow as the values of the convective terms depend on the temperatures of the south and north faces (i.e., T_{south} and T_{north}) of the liquid cells, eventually resulting in underestimation of the processor temperatures.

We demonstrate this effect in Figure 13 for the following system. We assume the same chip stack described in Figure 10(a), but simulate a $10mm \times 10mm$ die composed of four blocks with equal area, representing a conventional chip. We experiment with $q = 12.5, 25, 50, 100, 200 W/cm^2$.

As shown, placing the virtual grid at the bottom surface of a grid may result in up to $20^\circ C$ lower processor temperature in comparison to placing it in the middle (which is the adopted approach in the proposed model and gives matching results compared to 3D-ICE). This is an important factor as it would significantly change the outcome when evaluating different cooling designs.

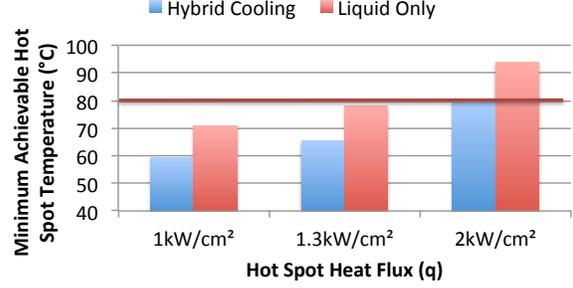


Fig. 14: Minimum achievable hot spot temperature using hybrid cooling and liquid-only cooling for different q values. The red horizontal line represents the maximum temperature constraint.

DEMONSTRATION OF HYBRID COOLING BENEFITS USING THE PROPOSED MODEL

In this section, we demonstrate the benefits of hybrid cooling combining liquid cooling with TECs using our proposed model. For this purpose, we create two different system scenarios: one uses solely liquid cooling and the other adopts hybrid cooling with liquid and TECs. We show that hybrid cooling can mitigate hot spots with very large heat fluxes reaching $2kW/cm^2$, which liquid cooling by itself cannot handle. In a liquid-only scenario, the chip stack looks like the one in Figure 10. For the hybrid cooling scenario, we add a TEC layer between the processing layer and the microchannel cooling layer as illustrated in Figure 1. TEC layer is such that the TEC device described in Figure 3 is placed right above the hot spot area and the rest of the TEC layer is covered with silicon. We assume a $20mm \times 20mm$ large chip as expected in the next generation processors [14]. We also assume a single hot spot with a size of $500\mu m \times 500\mu m$ and $q=1, 1.3, 2kW/cm^2$. We assign a background heat flux of $50W/cm^2$, which is commonly observed on modern processors. We assume that the hot spot is located near the microchannel outlet, representing an extreme case, where the liquid cooling efficiency will be poor due to heating up of the fluid.

The cooling limit of a liquid-cooled system is determined by the allowed pressure drop across the channels, which is determined by the manufacturers. For the TECs, on the other hand, we are limited by the amount of current we can apply without damaging the device and the wires. We use a maximum pressure drop limit of 1.3bars [18] and TEC current limit of 7A [3]. Next, in Figure 14, we compare the minimum achievable hot spot temperature when using the two cooling schemes for each q value. As suggested by the figure, hybrid cooling can provide up to $15^\circ C$ reduction in the hot spot temperature compared to just liquid cooling, at the cost of additional power consumed by the TECs. Moreover, a hybrid-cooled system can maintain maximum temperature under $80^\circ C$ (a temperature constraint widely used in current processors) for all q values. However, when using liquid cooling only, for the case of $q=2kW/cm^2$, we can at most reduce the temperature down to

94°C. As we demonstrated in this section, our proposed model is able capture the benefits of hybrid cooling with minimal modeling effort on the user side and with sufficient accuracy. One can leverage our hybrid model to do similar analysis for various TEC sizes, locations and properties as it does not depend on a specific target system.

CONCLUSION

In this paper, we have proposed a hybrid cooling model for fast and accurate simulation of systems with liquid cooling and TEC devices. We have integrated our model into a compact thermal simulator and provided comparison against COMSOL Multiphysics tool and 3D-ICE simulator. Our main conclusions are the following:

- Our proposed hybrid model provides a maximum error of less than 3.57°C for TECs and 2.46°C for liquid microchannels compared to COMSOL. Our model achieves this accuracy while providing 16990x and 43300x faster simulation than COMSOL for TECs and liquid cooling, respectively.
- One of the important aspects to be considered for accurate modeling of hybrid cooling with liquid microchannels and TECs is the placement of the virtual thermal node. For the liquid grid cells, placing the virtual node at the bottom surface instead of the middle of the grid can result in underestimation of the processor temperature by 20°C.
- For a fair comparison of different cooling designs with and without TECs, the TIM properties should be realistically selected. Otherwise, temperature benefits of TECs can be overestimated by 15°C.

In this work, we focused on modeling of the steady-state behavior of hybrid cooling systems. For the design and evaluation of runtime management policies, having a transient temperature model is necessary to account for the spatial and temporal changes in the heat flux over time due to the dynamic nature of the application characteristics. Thus, our future work in modeling will include adding transient simulation functionality. We also plan to release the proposed hybrid modeling tool for other researchers to use. The following step in future work will be to develop runtime optimization policies using our model to minimize the cooling power under temperature constraints in a hybrid-cooled system.

ACKNOWLEDGMENT

This project has been partially funded by the NSF CAREER grant #1149703 and the NSF CRI (CI-NEW) grant #1730316/1730003.

REFERENCES

- [1] ANSYS, <http://www.ansys.com>.
- [2] Comsol Multiphysics, <http://www.comsol.com>.
- [3] Ihtesham Chowdhury, Ravi Prasher, Kelly Lofgreen, Gregory Chryslers, Sridhar Narasimhan, Ravi Mahajan, David Koester, Randall Alley, and Rama Venkatasubramanian, *On-chip cooling by superlattice-based thin-film thermoelectrics*, Nature Nanotechnology **4** (2009), no. 4, 235–238.
- [4] Ayse K. Coskun, Tajana S. Rosing, Jose Ayala, and David Atienza, *Modeling and dynamic management of 3D multicore systems with liquid cooling*, IFIP/IEEE International Conference on Very Large Scale Integration (VLSI-SoC), 2009, pp. 60–65.
- [5] 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 **33** (2010), no. 1, 79–87.
- [6] A. Fourmigue, G. Beltrame, and G. Nicolescu, *Efficient transient thermal simulation of 3D ICs with liquid-cooling and through silicon vias*, Design, Automation and Test in Europe Conference and Exhibition (DATE), 2014, March 2014, pp. 1–6.
- [7] Wolfgang Gruener, *IBM cools 3D chips with intergrated water channels*, 2008.
- [8] S. Jayakumar and S. Reda, *Making sense of thermoelectrics for processor thermal management and energy harvesting*, Low Power Electronics and Design (ISLPED), 2015 IEEE/ACM International Symposium on, July 2015, pp. 31–36.
- [9] Xue-Xin Liu, Zao Liu, S.X.-D. Tan, and J. Gordon, *Full-chip thermal analysis of 3D ICs with liquid cooling by GPU-accelerated GMRES method*, Quality Electronic Design (ISQED), 2012 13th International Symposium on, March 2012, pp. 123–128.
- [10] Jie Meng, K. Kawakami, and A.K. Coskun, *Optimizing energy efficiency of 3-d multicore systems with stacked dram under power and thermal constraints*, Design Automation Conference (DAC), 2012, pp. 648–655.
- [11] F. Paterna and S. Reda, *Mitigating dark-silicon problems using superlattice-based thermoelectric coolers*, Design, Automation Test in Europe Conference Exhibition (DATE), 2013, March 2013, pp. 1391–1394.
- [12] V. Sahu, A. G. Fedorov, Y. Joshi, K. Yazawa, A. Ziabari, and A. Shakouri, *Energy efficient liquid-thermoelectric hybrid cooling for hot-spot removal*, Semiconductor Thermal Measurement and Management Symposium (SEMI-THERM), 2012 28th Annual IEEE, March 2012, pp. 130–134.
- [13] V. Sahu, Y. K. Joshi, A. G. Fedorov, J. H. Bahk, X. Wang, and A. Shakouri, *Experimental characterization of hybrid solid-state and fluidic cooling for thermal management of localized hotspots*, IEEE Transactions on Components, Packaging and Manufacturing Technology **5** (2015), no. 1, 57–64.
- [14] M. Schultz, F. Yang, E. Colgan, et al., *Embedded two-phase cooling of large three-dimensional compatible chips with radial channels*, ASME Journal of Electronic Packaging **138** (2016), no. 2, 021005–1–021005–5.
- [15] R. K. Shah and A. L. London, *Laminar flow forced convection in ducts: A source book for compact heat exchanger analytical data*, New York: Academic Press, 1978.
- [16] Chander Shekhar Sharma, Manish K. Tiwari, Severin Zimmermann, Thomas Brunschwiler, Gerd Schlottig, Bruno Michel, and Dimos Poulikakos, *Energy efficient hotspot-targeted embedded liquid cooling of electronics*, Applied Energy **138** (2015), no. C, 414–422.
- [17] K. Skadron, M.R. Stan, W. Huang, Sivakumar Velusamy, Karthik Sankaranarayanan, and D. Tarjan, *Temperature-Aware Microarchitecture*, Proceedings of International Symposium on Computer Architecture, 2003, pp. 2–13.
- [18] A. Sridhar, A. Vincenzi, M. Ruggiero, T. Brunschwiler, and D. Atienza, *3d-ice: Fast compact transient thermal modeling for 3D ICs with inter-tier liquid cooling*, 2010 IEEE/ACM International Conference on Computer-Aided Design (ICCAD), Nov 2010, pp. 463–470.
- [19] A. Sridhar, A. Vincenzi, M. Ruggiero, T. Brunschwiler, and D. Atienza, *Compact transient thermal model for 3D ICs with liquid cooling via enhanced heat transfer cavity geometries*, International Workshop on Thermal Investigations of ICs and Systems (THERMINIC), Oct 2010, pp. 1–6.
- [20] R. A. Taylor and G. L. Solbrekken, *Comprehensive system-level optimization of thermoelectric devices for electronic cooling applications*, IEEE Transactions on Components and Packaging Technologies **31** (2008), no. 1, 23–31.
- [21] John R. Thome, *Boiling in microchannels: a review of experiment and theory*, International Journal of Heat and Fluid Flow **25** (2004), no. 2, 128 – 139, Selected Papers from the 5th {ECI} International Conference on Boiling Heat Transfer.
- [22] K. Yazawa, A. Ziabari, Yee Rui 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*, Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm), 2012 13th IEEE Intersociety Conference on, May 2012, pp. 99–106.