

An EDA Tool for Co-designing High-Performance Processors and Emerging Cooling Technologies

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

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

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

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

Abstract—Future high-performance systems face elevated thermal challenges as power densities are expected to reach (and surpass) several thousand watts per centimeter square. High-density hot spots resulting from these power densities cause serious performance degradation or reliability issues, if not handled efficiently. Co-optimizing performance, power consumption, and the cooling capability can lead to enhanced performance without violating thermal constraints. This paper introduces a new EDA tool that facilitates co-design of high-performance processors with highly-efficient novel cooling techniques. The proposed tool is under-development and contains accurate thermal models of emerging cooling technologies such as liquid cooling, thermo-electric coolers, two-phase cooling, phase change material, or a hybrid cooling design. The planned tool is also capable of reducing simulation times by decoupling power/performance from thermal simulations. Further, we also plan to include several optimization modules that enable design-time/runtime policy decisions to achieve high performance while maintaining safe chip temperatures.

I. INTRODUCTION

Power densities in high-performance processors is expected to reach 1-5 kW/cm^2 [1], [2]. These high power densities will lead to thermal hot spots, and significantly degrade chip performance and reliability. As a result, temperature will be a key challenge in future high-performance processors because existing cooling techniques are not sufficient to cool down such high-density hot spots. Therefore, there is a need for novel cooling technologies.

Several emerging efficient cooling technologies (such as two-phase cooling or liquid cooling) are being developed by thermo-mechanical engineers to target high-density hot spots. However, there is no infrastructure to investigate and integrate these cooling technologies into a target high-performance system. To solve this problem, we plan to build an EDA tool that has the ability to model these cutting-edge cooling mechanisms, and co-design high-performance computing systems together with these highly-efficient cooling systems. Such a co-design helps in achieving improved performance while avoiding high chip temperatures. The main features of our planned EDA tool are:

- The tool enables co-optimization of high-performance processors with novel cooling techniques (Section II).
- The EDA tool can decouple the power/performance simulations from thermal simulations and save substantial amount of simulation time (Section II).

- We plan to make the tool modular so that different database and optimization modules of interest can be plugged in (Section II).
- Thermal models of different cooling techniques can be developed and plugged into the tool (Section III).

II. OVERVIEW OF PROPOSED EDA TOOL

Our proposed EDA tool provides an interface to create an optimized computing system along with an efficient cooling subsystem. The thermal simulation infrastructure in our proposed EDA tool contains thermal models for various emerging cooling techniques as well as more traditional cooling methods. We plan to design the proposed tool such that it has the ability to optimize computing system design together with cooling optimization. The computing systems can be optimized in terms of performance and/or energy-efficiency, while the cooling subsystem selects the most-efficient cooling technique for the target system. To build such a tool, we have to develop a system-level simulation framework that combines performance, power, and temperature modeling with several design-time/runtime optimization and management control knobs. These control knobs include dynamic voltage and frequency scaling, thread migration, thermal-aware floorplanning, etc., and are used in different policy decisions. The policies and control knobs form part of an optimization module, which is implemented in the thermal simulation infrastructure.

Figure 1 shows a flow diagram of our proposed system-level simulation framework for the EDA tool. “Option A” shows a sequential flow of performance, power, and thermal simulations. While such a flow may work well when power traces are not changing with time, it is practically impossible to run simulations with design-time/runtime policy decisions. This is because of two reasons: (i) performance and power simulations are already highly time-consuming, and (ii) for runtime optimization, there must be a feedback loop between temperature and power/performance simulations which results in even longer simulation times. Therefore, it is necessary to decouple thermal simulations from power/performance simulations (implemented as “Option B”). The decoupling enables faster simulation times, and provides support for optimized policy decisions at design-time/runtime.

In “Option B”, we generate offline power/performance traces corresponding to each application and runtime control knob of interest (eg., different thread counts for multi-threaded applications or different V/f setting for the cores), and store these traces in a database. This database can be generated using simulations or from measurements collected on real processors. The thermal simulator polls this database based on the executing application and policy decisions.

III. COOLING TECHNIQUES

Compact thermal models (CTM) use the duality between thermal and electrical properties for fast thermal modeling and simulations. In CTM, the chip is divided into a network of grid cells, and each cell has an equivalent thermal resistor and a thermal capacitor (RC) circuit. The processor temperature is then calculated by solving a lumped RC network. In the proposed EDA tool, we incorporate compact thermal models for four different cutting-edge cooling technologies: phase change material, microchannel liquid cooling, thermo-electric cooler, and two-phase cooling. Currently, we have integrated these models inside HotSpot 6.0 simulator [3]. We plan to integrate the above cooling models into a new modular, high-performance tool, with easy-to-use user interfaces. We also intend to investigate other fast and accurate thermal modeling techniques and incorporate those in our tool.

A. Phase Change Material

Phase change material (PCM) removes heat through solid-liquid phase change while maintaining a near constant temperature. In our previous work, we build a compact thermal model for phase change material cooling, and assign a temperature-dependent specific heat capacity to each PCM cell [4]. We integrate the PCM model inside the transient analysis solver, where specific heat capacity for each PCM cell is updated dynamically based on temperature change. This model gives a maximum error of 3.35°C on validation against COMSOL.

B. Microchannel Liquid Cooling

Microchannel liquid cooling is an attractive solution to remove high heat fluxes in 3D-stacked architectures. In this technique, liquid is pumped through microchannels etched at the back of silicon, and heat is dissipated in the form of convection. We adopt the 4-Resistor model [5], and assign a heat transfer coefficient (htc) for conduction from the sidewalls [6]. In addition, we use voltage-controlled current source to represent the heat convection along the liquid flow. The microchannel liquid cooling model shows an average error of 0.36°C on validation against COMSOL.




Fig. 1. System-level simulation framework of our EDA tool

C. Thermo-electric Coolers

When a bias current is applied to a thermo-electric cooler (TEC), heat is absorbed on one side and rejected on the other side. In our recent work, we incorporate the Peltier term (heat absorption and rejection), and Joule heating term (resistive heat generated by the bias current as it passes through TEC) as a current source [6]. HotSpot simulator accounts for the heat conduction term in thermal RC network by default. The TEC compact thermal model shows an average error of 2.07°C when validated against COMSOL.

D. Two-phase Cooling

Two-phase cooling is an emerging technique that removes heat through evaporation process of a coolant, and has the potential to remove high heat fluxes. Our two-phase cooling structure employs a micropillar array as a thin-film evaporator on the heated bottom surface of the microchannel. These micropillar structures suppress dry-out and increase the heat transfer coefficient. We plan to derive htc correlations for sidewalls and bottom surface based on the prototype in [7]. Furthermore, the latent heat of evaporation can be represented as a voltage-controlled current source in our compact thermal model.

IV. DISCUSSION AND FUTURE WORK

In our previous works, we have designed and implemented CTMs for the following cooling techniques: PCM, liquid cooling, TEC, and a hybrid cooling design of liquid cooling and TEC [4], [6]. We also have built the framework to decouple power/performance simulations from thermal simulations [4]. Currently, we are working on implementing a compact thermal model for two-phase cooling with micropillar evaporators as well as finalizing a version of the decoupled simulation framework. Eventually, we will release the EDA tool to the research community.

ACKNOWLEDGEMENT

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

REFERENCES

- [1] Schultz *et al.*, “Embedded two-phase cooling of large three-dimensional compatible chips with radial channels,” *Journal of Electronic Packaging*, vol. 138, no. 2, p. 021005, 2016.
- [2] 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*, vol. 6, no. 7, pp. 1056–1065, 2016.
- [3] Skadron *et al.*, “Temperature-aware microarchitecture,” in *Computer Architecture, 2003. Proceedings. 30th Annual International Symposium on*. IEEE, 2003, pp. 2–13.
- [4] Kaplan *et al.*, “Modeling and analysis of phase change materials for efficient thermal management,” in *Computer Design (ICCD), 2014 32nd IEEE International Conference on*. IEEE, 2014, pp. 256–263.
- [5] Sridhar *et al.*, “3d-ice: Fast compact transient thermal modeling for 3d ics with inter-tier liquid cooling,” in *Proceedings of the International Conference on Computer-Aided Design*. IEEE Press, 2010, pp. 463–470.
- [6] F. Kaplan, S. Reda, and A. K. Coskun, “Fast thermal modeling of liquid, thermoelectric, and hybrid cooling,” in *Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm), 2017 16th IEEE Intersociety Conference on*. IEEE, 2017, pp. 726–735.
- [7] Zhu *et al.*, “Surface structure enhanced microchannel flow boiling,” *Journal of Heat Transfer*, vol. 138, no. 9, p. 091501, 2016.