

HotSpot 6.0: Validation, Acceleration and Extension

Runjie Zhang^{1,2}, Mircea R. Stan¹, and Kevin Skadron¹

¹University of Virginia, Charlottesville, VA

²IBM T. J. Watson Research Center,
{runjie,mircea,skadron}@virginia.edu

May 11, 2015

1 Introduction

The challenge of removing the heat generated by microprocessors with reasonable cost has long been identified as a critical issue [9]. As a matter of fact, the cooling constraint (a.k.a., thermal wall) has become a major hurdle that limits the scaling of operating frequency and transistor density. As a fast and accurate thermal model, HotSpot [1] enables early-stage evaluation of chip temperature and therefore supports architectural study of dynamic thermal management strategies, chip floorplanning, and novel cooling solutions.

The recent industry trend of moving toward tighter in-package integration (e.g., stacked DRAM) brings both opportunities and challenges. From a thermal design point of view, such integration along the third dimension significantly increases the density of heat emission and therefore requires more efficient cooling solutions as well as thermal management techniques. From a thermal modeling perspective, 3D integration raises the complexity of the problem in terms of both problem size and system heterogeneity. In order to further facilitate thermal modeling for 3D integrated silicon chips, we developed a new version of HotSpot (version 6.0) with the following new features:

- Calibration/validation against representative reference data created under mentorship from an experienced IBM POWER-family power/thermal modeling team.
- Improved steady-state solver.
- Support layers with non-uniform thermal resistivity and heat capacity. This feature is originally contributed by Prof. Ayse Coskun’s group in Boston University.
- Improved support for secondary heat transfer path.

HotSpot version 6.0 is available at <http://lava.cs.virginia.edu/HotSpot>. This report provides a detailed description for all the major new features.

2 Calibration and Validation Reference

In order to further validate HotSpot’s accuracy as well as to provide users a set of modeling parameters that represents the cooling system of an up-to-date, commercial and high-performance processor, one of us (the first author) spent time at IBM T. J. Watson Research Center, and worked with domain experts (associated with IBM’s POWER-family processors) to create a suitably representative HotSpot modeling framework. In particular, we started with the published chip dimensions of IBM’s POWER7+ chip [10], since that serves as the baseline architectural reference in the DARPA PERFECT project that provided majority funding for this validation work. POWER7+ is a 32nm, 8-core processor with 80MB of L3 cache; the particular system product offering that provided our package-level modeling context is the model 730 [2]; the chip in that product can operate at clock frequencies at or above 4.2 GHz.

In working with the IBM team, we got access to representative power and temperature maps. These did not necessarily correspond to any real multi-core workload, but were artificially created power-thermal correspondences, based on synthetic stressmarks. The chip floorplan was uniformly divided. The reference thermal map was created from the artificially created power map

	Calibrated Parameters
Convection Resistance (K/W)	0.17
Heat Sink Thickness (mm)	6
Heat Sink Side (mm)	66
Heat Sink Thermal Conductivity ($W/m.K$)	400
Heat Spreader Thickness (mm)	2
Heat Spreader Side (mm)	48
Heat Spreader Thermal Conductivity ($W/m.K$)	400
Interface Material Thickness (mm)	0.1
Interface Material Thermal Conductivity ($W/m.K$)	5
Silicon Chip Thickness (mm)	0.78
Silicon Chip Thermal Conductivity ($W/m.K$)	140

Table 1: Calibrated HotSpot Parameters

using IBM’s internal, low-level physical thermal model. Both the power and temperature data was provided in steady-state form only.

Table 1 shows the HotSpot package model parameters assumed for this particular validation exercise. We did not have access to the real product parameters, but we were guided by the IBM mentors to make reasonable assumptions based on non-IBM external references [3, 4]. The only empirical parameters are those corresponding to the thickness of the thermal interface material (TIM) and heat spreader, which we tuned within ranges suggested by [5], and are not representative of any IBM product packages. We note that the thermal interface material between the heat spreader and heat sink (TIM2) is not directly supported by HotSpot; hence, we had to ignore it in this work. Overall, our objective was to start with industry-validated power-thermal maps and make reasonable adjustments to our package model (based on externally published data and guidance from IBM mentors) in order to make sure that HotSpot produces thermal results that build confidence in the overall modeling infrastructure.




Figure 1: Validation Results

To validate HotSpot, we created a synthetic floorplan using ArchFP [6] and converted IBM’s

power consumption data into HotSpot-compatible format accordingly. Fig. 1 shows the resulting error map. Across the floorplan, the errors of HotSpot results range from -3.41% to +2.15% with an average absolute error ratio of 0.90% (Fig. 1a). If we compare the absolute temperature values (Fig. 1b), HotSpot gives a error range of -1.40 to 1.15 degree Celsius with an average absolute error of 0.43 degree Celsius. These validation results indicate that HotSpot has high accuracy.

It is worth mentioning that according to Fig. 1, HotSpot gives higher temperature estimations (positive error values) for areas with higher power consumption and lower temperate estimations (negative error values) for the less active blocks. This phenomenon indicates that HotSpot is underestimating the amount of heat transferred laterally. If users want to further improve modeling accuracy, reducing lateral thermal resistance would reduce error.

3 Supporting Layers with Non-uniform Thermal Properties

Compared with modeling 2D chips, simulating heat transfer in 3D ICs involves a much higher level of complexity due to the increase of problem size and system heterogeneity. To be more specific, stacking multiple silicon dies together not only proportionally increases the number of nodes to simulate, but also introduces layers with non-uniform thermal properties. For example, through-silicon-vias (TSVs) travel vertically to connect different silicon layers together. Although the TSVs consists of uniform material themselves, their thermal property usually differs from the layers they go through (e.g., TIM between two silicon layers). Consequently, those layers can no longer be treated as uniform thermal conducting sheets.

The previous versions of HotSpot assumed that the thermal resistance and specific heat of each layer (e.g., silicon, TIM, package) is uniform. Therefore, components like TSV could not be precisely modeled. To solve this problem, Prof. Ayse Coskun’s research group in Boston University developed an extension to HotSpot (based on version 5.02) that allows blocks within the same layer to have different thermal properties. The basic idea is to: 1) add a data structure that records per grid-node thermal RC instead of per-layer RC; 2) take per-block thermal resistance and specific heat as input (along with other per-block geometric information in the .flp input file); 3) calculate per-grid-node thermal RC based on the provided block-level thermal properties; and 4) solve the heterogeneous RC network. More details and experiments can be found in their recent paper [8].

In HotSpot-6.0, we incorporated this extension and made it compatible with other new features like the new steady-state solver (Sec.4) and secondary heat transfer path (Sec.5.1). To enable this new feature in simulation, command line option ‘detailed_3D’ should be set to ‘on’. There is a slight change in the format of the floorplan file (.flp) that allows units in the floorplan to have different heat capacitance and resistivity. Regardless of these changes, the command line interface and file I/O format are backward compatible. We note that currently, heterogeneous layers could only be modeled in grid-mode with .lcf file specified. Also, all the major contributions to the source code by BU are commented with “BU_3D”.

4 Fast Steady-State Solver Based on SuperLU

For steady-state simulation, the previous versions of HotSpot solves the resistor network by applying Kirchhoff’s current law (KCL) iteratively until the result differences between two consecutive iterations are lower than a threshold. To speed up this iterative approach, multi-grid method was also adopted to accelerate the convergence rate. Nevertheless, the grid-mode, steady state simulation time with fine-grained on-chip resolution (i.e., large grid size) is still relatively long. To make matters worse, the new feature of supporting layers with heterogeneous thermal properties reduces the efficiency of multi-grid method and therefore slows down 3D simulations. To further speedup grid-mode’s steady-state simulation, we implemented a new solver that directly solves the resistor network with LU decomposition. Since LU decomposition is a common algorithm, there exist many highly optimized libraries. We choose SuperLU [7], an open source library for the direct solution of large, sparse systems of linear equations as the engine of our new steady-state solver. Documentations, installation instructions and source code of SuperLU can be found here: <http://crd-legacy.lbl.gov/~xiaoye/SuperLU/>. The HOWTO file included in HotSpot-6.0 package gives detailed instructions on how to equip HotSpot with SuperLU and simulate steady-state with the new solver.

Problem Size	2D_Old	2D_New	2D_Speedup	3D_Old	3D_New	3D_Speedup
64x64	5.46	0.26	21.3x	10.51	0.63	16.7x
128x128	11.50	2.06	5.6x	32.33	5.57	5.8x
256x256	35.86	16.88	2.1x	96.56	49.18	2.0x
512x512	137.37	153.75	0.9x	408.91	456.54	0.9x

Table 2: Simulation time comparison between previous iterative solver and new direct solver. Except speedup results, all values are in seconds.

Table 2 shows our internal test results on simulation speed. The test platform is a linux machine with Intel Xeon X5550 CPU and 64GB main memory. Runtime results were measured with command “time”. We evaluated both the default 2D case (one layer of silicon, one layer of TIM) and a 3D case with two layers of silicon and two layers of TIM. The on-chip grid size (controlled by command line option `-grid_rows` and `-grid_cols`) was swept from 64-by-64 to 512-by-512. Note that larger on-chip grid provides fine-grained estimation for on-chip temperature at the cost of larger problem size. Our results show that the new solver achieves up to 20x speed up against the old solver. However, due to its large memory footprint (e.g., 8GB virtual memory usage with 512x512 grid size), direct solver will significantly slowdown when the problem size is large. For this reason, we recommend using direct solver when problem size is not extremely large (e.g. no larger than 256x256 for a 4-layer 3D chip). If the host machine’s memory is limited (less than 8GB), smaller grid-size is preferred. Also, we recommend using direct solver when users enable the modeling for heterogeneous layers (through command line option `-detailed_3D`). We note that the solver could be easily selected through makefile option at compile time.

5 Other New Features and Updates

5.1 Improved support for secondary heat transfer path

The secondary heat transfer path refers to the heat escape path from the silicon to on-chip interconnect layers, C4 pads, packaging substrate, solder balls and printed-circuit board. Previous versions of HotSpot do not support secondary heat transfer path when 3D IC is modeled. This version relaxes this constraint and now secondary heat transfer path can be modeled along with 3D stacks.

5.2 Improved support for plotting 3D floorplan

In the past, we released a script (tofig.pl) along with HotSpot source code to help users visualize their 2D floorplans. HotSpot-6.0 includes an updated script (3Dfig.pl) that automatically plots all layers of a 3D stack. It has similar functionality to tofig.pl but instead reads in an .lcf files and creates a .FIG file for each floorplan listed in the .lcf file. Additionally, it will show a unit's resistivity and capacitance if it specified in the floorplan file. More details and instructions are available in the README-6.0 file. We note that this script is also contributed by the BU team.

We also note that starting from version 6.0, we no longer provide maintenance for the Excel interface.

6 Conclusions

To better facilitate thermal modeling for contemporary and near-future 2D/3D silicon chips, we developed a new version (6.0) of HotSpot that provides faster and more flexible modeling capabilities. By pursuing a validation/calibration exercise that uses industrial design team approved power-thermal maps as reference, and package parameters that are assumed from credible published references (as approved by our IBM mentors), we find that HotSpot could estimate steady-state on-chip temperature with less than 3.5% error. With the help of a direct solver, we speed up steady-state state simulation by up to 20x. By incorporating an extension developed by a BU research team, we are able to model layers with heterogeneous thermal properties. HotSpot version 6.0 is available at <http://lava.cs.virginia.edu/HotSpot>. For further questions, please e-mail the HotSpot user group at "hotspot@cs.virginia.edu".

Acknowledgments

The majority of this work is supported by DARPA MTO under contract no.HR0011-13-C-0022. The views expressed are those of the authors and do not reflect the official policy or position of the Department of Defense or the U.S. Government. This document is Approved for Public Release. Distribution Unlimited.

For the POWER7+ calibration/validation work, we would like to thank Pradip Bose and Prithish R. Parida from IBM.

For the heterogeneous 3D modeling work, we would like to thank Ayse Coskun, Katsutoshi Kawakami, Daniel Rossell, Samuel Howes, Tiansheng Zhang, and Fulya Kaplan at Boston Uni-

versity, David Atienza at EPFL, Mohamed Sabry at EPFL, Yusuf Leblebici at EPFL, and Tajana Rosing at UCSD.

References

- [1] <http://www.cs.virginia.edu/HotSpot>.
- [2] <http://www-03.ibm.com/systems/power/hardware/730/perfdata.html>.
- [3] http://www.aavid.com/product-group/microprocessors/intel_processors.
- [4] <http://www.chipsetc.com/silicon-wafers.html>.
- [5] <http://www.electronics-cooling.com/2006/08/packaging-challenges-for-high-heat-flux-devices/>.
- [6] Gregory G Faust, Runjie Zhang, Kevin Skadron, Mircea R Stan, and Brett H Meyer. ArchFP: Rapid prototyping of pre-rtl floorplans. In *VSLI-SoC*, 2012.
- [7] Xiaoye S. Li. An overview of SuperLU: Algorithms, implementation, and user interface. *ACM TOMS*, 31(3), 2005.
- [8] Jie Meng, Katsutoshi Kawakami, and Ayse K Coskun. Optimizing energy efficiency of 3-D multicore systems with stacked dram under power and thermal constraints. In *Design Automation Conference (DAC)*, 2012.
- [9] K. Skadron, M. R. Stan, W. Huang, S. Velusamy, D. Tarjan, and K. Sankaranarayanan. Temperature-Aware Microarchitecture. In *International Symposium on Computer Architecture (ISCA)*, June 2003.
- [10] S Taylor et al. POWER7+™: IBM's next generation POWER microprocessor. In *Hot Chips*, volume 24, 2012.