Accurate Power Density Map Estimation for Commercial Multi-Core Microprocessors

Jinwei Zhang, Sheriff Sadiqbatcha, Wentian Jin, Sheldon X.-D. Tan
Electrical and Computer Engineering, University of California, Riverside, CA 92521

***Abstract*—In this work, we propose an accurate full chip steady-state power density map estimation method for the commercial multi-core microprocessors. The new approach is based on the measured steady-state thermal maps (images) from an advanced infrared (IR) thermal imaging system to ensure its accuracy. The new method consists of a few steps. First, based on the first principle of heat transfer, 2D spatial Laplace operation is performed on the given thermal map to obtain the so-called *raw* power density map, which consists of both positive and negative values due to the steady-state nature and boundary conditions of the microprocessors. Then based on the total power of the microprocessor from an online CPU monitoring tool, we develop a novel scheme to generate the actual real positive-only power density map from the raw power density map. At the same time, we develop a novel approach to estimating the effective thermal conductivity of the microprocessors. To further validate the power density map and the estimated actual thermal conductivity of the microprocessors, we construct a thermal model with COMSOL, which mimics the real experimental set up of measurement used in the IR imaging system. Then we compute the thermal maps from the estimated power density maps to ensure the computed thermal maps match the measured thermal maps using FEM method. Experimental results on intel i7-8650U 4-core processor show 1.8°C root-mean-square-error (RMSE) and 96% similarity (2D correlation) between the computed thermal maps and the measured thermal maps.**

I. Introduction

Power, thermal and related reliability issues are the major limiting factors on today's high performance multi-core processors, especially after the breakdown of the so-called Dennard scaling, since power density starts to increase as IC technology advances [1], [2]. To enhance reliability, researchers have proposed many power/thermal regulation or dynamic thermal management methods, including clock gating, power gating, Dynamic Voltage and Frequency Scaling (DVFS), and task migration techniques [3]–[6].

However, one important aspect of those works depends on how to correctly estimate the full-chip temperature. The on-chip temperature is mainly obtained by performing the thermal analysis based on the run-time functional unit (or component-wise) power estimation of the processor. Estimating component power inputs still remains challenging for commercial off-the-shelf microprocessors.

A few existing works have proposed to estimate the component and the total power of a *real* microprocessor [7]–[10]. One idea is to tune each component unit power until the summation matches with the total power that can be measured experimentally [7], [8]. The main difficulty of those approaches, however, is that the searching for component unit power values still remains an ad-hoc approach, which almost always involves manual tunings. Wu *et al.* tried to mitigate this problem by performing linear regression with K-means-based method to identify the unique power track patterns from the running programs [9]. In [10], the authors frame the problem as constrained optimization problem once the thermal models were obtained from finite element simulation and measurement. Recently an RNN-based approach has been proposed to fast estimate the thermal and power hotspot based on the system performance metrics such as Intel's Performance Counter Monitor (IPCM) using machine learning based approaches [11].

In this paper, instead of focusing on the functional units or component power,we try to obtain the full-chip power density map from the measured thermal maps/images of commercial multi-core microprocessor. Once power density maps are obtained, component power can be easily obtained by area integration over the chip layout. The estimated power map also provides many insights into power consumption of different modules and cores and uncores in a microprocessor and can be instrumental to many other power/thermal management applications. Our main contributions are as followings:

1. First, based on the first principle of heat transfer, 2D spatial Laplace operation is performed on the given thermal map to obtain the so-called *raw* power density map, which consists of both positive and negative values due to the steady-state nature and boundary conditions of the microprocessors. We study two simple cases to provide many insights into relationship between the raw power density map and real power density maps. This work is enabled by an advanced thermal measuring platform with a high-precision thermal camera and a cooling system installed on the back side of the CPU. It allows us to take explicit temperature images (thermal maps) of CPU die while the CPU is under load.
2. Then based on the total power of the microprocessor obtained using an online CPU monitoring tool, we develop a novel scheme to generate the actual real positive-only power density map from the raw power density map. At the same time, we develop a novel method to compute effective thermal conductivity of microprocessor dies.
3. To further validate the power density map and the estimated actual effective thermal conductivity of the microprocessors, we construct a thermal model with COMSOL, which mimics the real experimental setup: with the same boundary conditions used in the IR imaging system. We use the thermal measurements when CPU is in idle status to determine the boundary conditions of the thermal simulation model. Then we compute the thermal map based on the estimated power density map to ensure the computed thermal maps match the measured thermal maps using FEM method.

Experimental results on intel i7-8650U 4-core processor show that the root-mean-square-error (RMSE) of the computed thermal maps and the measured thermal maps are around 1.8°C. Furthermore, computed thermal maps and the measured thermal maps have 96% similarity (2D correlation).

This work is supported in part by NSF grant under No. CCF-1527324, and in part by NSF grant under No. CCF-1816361 and in part by NSF grant under OISE-1854276.

978-3-9819263-4-7/DATE20/©2020 EDAA

II. The power density modeling setup

In this section we briefly outline the framework of the proposed approach, thermal imaging system, and necessary data collection from the commercial multi-core processor.

A. The power density modeling framework

Power map (surface power density distribution) has tight relationship with the temperature distribution, the Laplace transform of temperature and the thermal conductivity. Our proposed approach involves two kinds of data. First one is the thermal maps of CPU measured through a high-precision thermal camera, which senses the infrared emissions from CPU surface and transforms them into images of temperature distribution. The second is the total CPU power consumption along time, which can be obtained through the processor's Performance Counter Monitor. Fig. 1(a) shows the *validation flow* of the power map estimation model and evaluate how accurate the estimations are. Fig. 1(b) illustrates the *estimation flow* from data resources to estimated results for real processors.

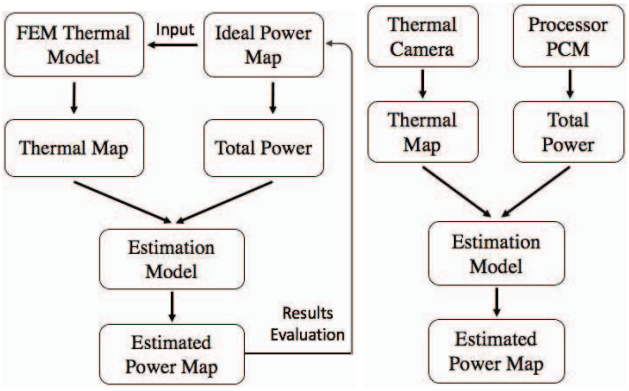

Fig. 1. Frame work overview: (a) validation flow (b) estimation flow.

B. Thermal imaging system

High precision and resolution of measuring the thermal maps are critical to the estimation results. In order to achieve these, we take advantage of a high-precision thermal camera installed closely over the CPU die along with a liquid cooling system attached underneath the motherboard, as shown in Fig. 2. This system adapts the thermal imaging method proposed in [12]. It maximizes the explicitness of thermal maps by directly exposing the top surface of CPU die to the camera, while ensuring the CPU's normal thermal condition by cooling it from the back side of the motherboard. Massive heat generated from CPU flows downwards through the motherboard into the cooling system, and is finally dissipated by the quickly circulated coolant. The model of our thermal camera is FLIR A325SC, 240×320px, 16bit and 60Hz capturing rate. Thanks to a close-up lens, the camera makes temperature difference 50mK clearly visible within as small as 50μm/px.

III. New power density map estimation method

In this section, we present our new method to estimate the full-chip power density from the real multi-core processors. We start with a simple example, which leads to an important observation for the proposed method. Then we will present how to compute the thermal conductivity of the real microprocessors, which is important for later thermal analysis and validation.

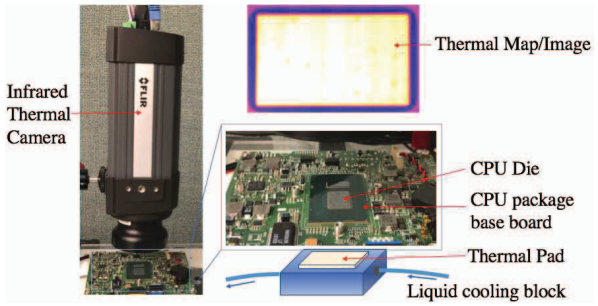

Fig. 2. Thermal imaging system

A. Proposed power density map based on Laplace operation

We first start with the fundamental heat diffusion equation (1), which gives the relationship between temperature and heat generation:

$$\rho C_P \frac{\partial T}{\partial t} - \nabla(\kappa \nabla T) = g \tag{1}$$

where T is temperature (K), ρ is the mass density of the material (kg·m^{-3}), C_P is the mass heat capacity (J·kg^{-1}·K^{-1}), κ is the thermal conductivity (W·m^{-1}·K^{-1}) and g is the spatial heat energy generation (W · m^{-3}).

When CPU runs into steady state the transient term can thus be ignored and the equation (1) can be simplified as:

$$-\kappa \nabla^2 T = g_T(x, y) \tag{2}$$

where ∇^2 is the Laplace operator. From the simplified heat equation (2), we can see that *the negative spatial Laplacian* of the temperature distribution across the die is equal to the spatial heat generation. Therefore, we can perform the 2D spatial Laplacian on a given thermal map to locate the underlining heat-sources $g_T(x, y)$.

However, this view only works well for identifying locations for heat source peaking, analysis for quantitative relationship between the Laplace transform and the power distribution for entire chip was not enough. To identify the major power and thermal hotspots, it is sufficient to merely find the locations of local minimum values across the Laplace transform. This is because in steady state a local power density peak will always cause a local minimum Laplacian value on that same spot. But when looking at the whole die area, negative-Laplacian of spatial temperature is not exactly equivalent to the CPU power distribution. One important observation is that the negative-Laplacian of the temperature can have *negative values*. This cannot be explained by CPU power distribution since CPU power will never be negative. Fig. 3 shows one of such negative-Laplacian maps based on experimental measurements, where negative parts have remarkable amplitude comparable to the positive parts.

In order to closely study the relationship between the Laplace transform of temperature and the CPU power distribution, we build a simple ideal case in COMSOL Multiphysics heat transfer tool [13]. This case structure contains a rectangular base which has geometry 10×15×0.5mm, and a 4×4×0.5mm heat source block embedded in the base, whose total power is set at 3W (0.1875W · mm^{-2}) and homogeneous in space (Fig. 4(a)). The geometries can be flexible, we set it to approximately match the usual size of CPU die and core. Material of the structure in this case has κ of 400W · m^{-1} · K^{-1}. For the boundary conditions, a convective heat flux set at 1000W · m^{-2}K^{-1} is applied on the bottom surface. This convective heat flux mimics the heat dissipation

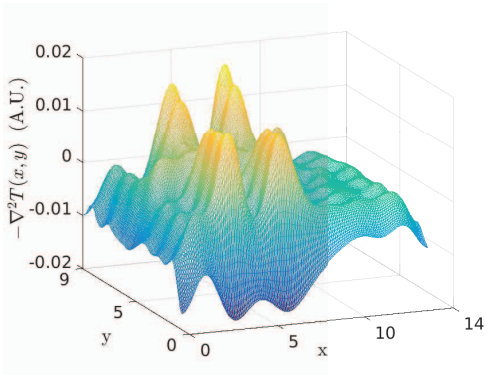

Fig. 3. A negative-Laplacian map example of experiment thermal measurements in 3D view.

through bottom surface. Ambient temperature is set at room temperature 297K.

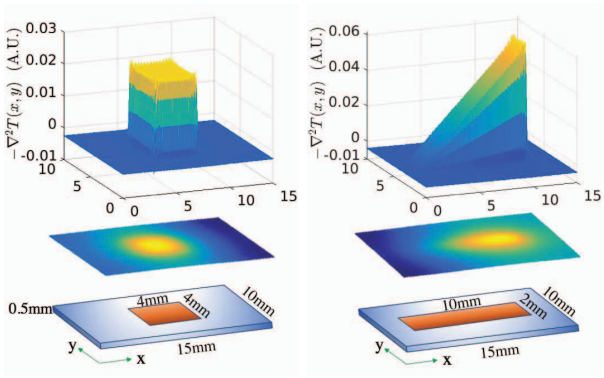

(a) (b)

Fig. 4. Simple ideal cases: (a) homogeneous heat source in orange region with total power 3W. (b) linear heat source in orange region (10×2mm), with areal power density $0.05(x-2)$ W $\cdot$ mm^{-2}, $x \in [2, 12]$ and total 5W.

As shown in Fig. 4(a), it is obvious that the high-raising portion of negative-Laplacian map reflects the area of active power density. Furthermore, we observe that *integration of negative-Laplacian map over all the area (pixels) is always zero*, no matter how the power setting and geometry changes. The reason is that the thermal map we obtained comes from steady state of the CPU with specific thermal boundary conditions. This means that power generation and power dissipation are balanced in such equilibrium state. The negative power density value actually represents more power dissipation than generation at the specific location due to the thermal transfer and convection process at the boundaries. Where positive value means the opposite. For the very positive high-raising portion, which means the heat generation is significantly larger than the dissipation, typically indicates the hotspots of the chip.

In another example, we have an ideal linear heat source, whose power density increase linearly along the x-axis with total power 5W. Fig. 4(b) illustrates the location of power source, power setting and its corresponding negative-Laplacian map. We can observe the negative-Laplacian in such rectangular power region shows an important linear trend as well, while the surrounding region is negative.

Based on the observations from these ideal examples, we can see that the positive part of the negative-Laplacian map are the region where most of the real power density is located. In this two simple cases, they contain the 100% real power density distribution. As a result, we can just use the *positive part* of the the negative-Laplacian map as the estimated power map. But the actual values of power map in those region are yet to be determined, which will be answered in the following section by calculating the accurate thermal conductivity κ.

B. Estimation of real thermal conductivity

Modern microprocessor die is usually as thin as within 0.5mm thickness, thus thermal characteristics along z-axis can be viewed as homogeneous. Power density distribution is only important on the surface x-y plane.

In reality, heat density is a combination of CPU power and heat dissipation by heat sink. Assume the thickness of CPU die is Δz, $p(x,y)$ stands for surface power density at location (x,y) and $p_d(x,y)$ denotes heat dissipated locally. Heat density can be expressed as:

$$g_T(x,y) = \frac{p(x,y) - p_d(x,y)}{\Delta z} \tag{3}$$

Then for location (x,y), (2) can be rewritten as:

$$-\kappa \nabla^2 T(x,y) = \frac{p(x,y) - p_d(x,y)}{\Delta z} \tag{4}$$

Considering the entire chip, integrate both sides on the whole die area,

$$-\kappa \int \nabla^2 T(x,y) dxdy = \int \frac{p(x,y) - p_d(x,y)}{\Delta z} dxdy \tag{5}$$

Suppose P is total CPU power, and P_d is total heat dissipation (mainly through convective heat flux), (5) can be further written as:

$$-\kappa \int \nabla^2 T(x,y) dxdy = \frac{P - P_d}{\Delta z} \tag{6}$$

At steady state P should be equal to P_d. This infers the integration on the right hand side of (6) would give zero total heat, as CPU power is balanced with heat removal. It also implies the integrated Laplacian should be zero. In fact, this zero result has been observed both in our experiments and aforementioned simulation. Based on the discussion in the previous sub-section, (6) can be approximated as:

$$-\kappa \int_{S_P} \nabla^2 T(x,y) dxdy \approx \frac{P}{\Delta z} \tag{7}$$

$$\kappa \approx \frac{P/\Delta z}{-\int_{S_P} \nabla^2 T(x,y) dxdy} \tag{8}$$

where S_P indicates area where negative-Laplacian of temperature is positive. Since die thickness Δz is constant, once negative-Laplacian map is obtained from temperature image, the equivalent thermal conductivity κ can be obtained. It basically means that the proportional factor κ can be acquired from dividing total power by thickness and by areal integration of the positive parts of negative-Laplacian. Having this κ, CPU power density map becomes straightforward, which is expressed as:

$$p(x,y) = \begin{cases} \kappa \Delta z[-\nabla^2 T(x,y)], & -\nabla^2 T(x,y) > 0 \\ 0, & -\nabla^2 T(x,y) \leq 0 \end{cases} \tag{9}$$

Using the above equations to estimate the power map for the homogeneous heat source example and the linear heat source example, the results are showed in Fig. 5. Fig. 5(c) and Fig. 5(d) are the estimated power densities for the two

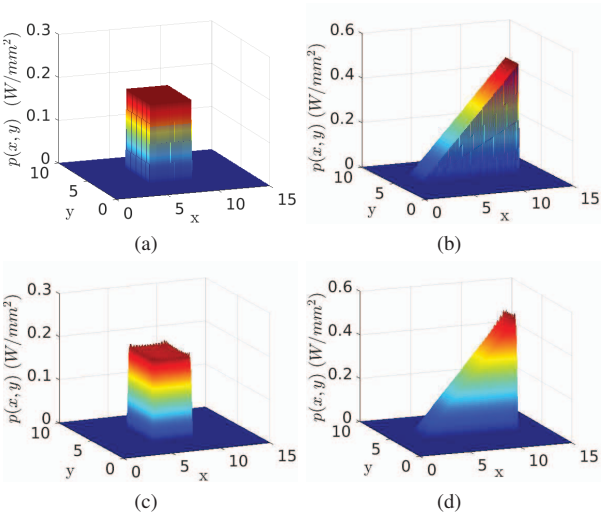

Fig. 5. Comparison between estimated power density maps and exact ones for simple ideal examples. (a) and (b) are original power density maps for homogeneous heat source and linear heat source, respectively. (c) and (d) are the corresponding estimated power density maps

cases, while Fig. 5(a) and Fig. 5(b) are the corresponding original power density maps. As we can see, some spikes exist at corners in the estimation results due to numerical noise.

To compare the similarity of the two power maps, we introduce *2D correlation coefficient*, or simply *correlation* to evaluate the similarity between the real power map and the estimated power map, which is defined as

$$r = \frac{\sum_m \sum_n \left(A_{mn} - \overline{A}\right)\left(B_{mn} - \overline{B}\right)}{\sqrt{\left(\sum_m \sum_n \left(A_{mn} - \overline{A}\right)^2\right)\left(\sum_m \sum_n \left(B_{mn} - \overline{B}\right)^2\right)}} \quad (10)$$

where $\overline{A}$ and $\overline{B}$ are mean of all entries in A and B, respectively. r is a scalar between 0 and 1, the more it approaches 1 the more they look alike. For the above two examples, the correlations of the first and second example are 0.977 and 0.973 respectively. In addition, RMSE of estimated power map on the active powered region is $0.005\text{W} \cdot \text{mm}^{-2}$ and $0.015\text{W} \cdot \text{mm}^{-2}$ respectively for the two cases as well. As a result, we can see that the proposed power map estimation method is quite accurate.

The thermal conductivity κ of silicon is about $130\text{W} \cdot \text{m}^{-1} \cdot \text{K}^{-1}$, copper is about $400\text{W} \cdot \text{m}^{-1} \cdot \text{K}^{-1}$. Due to the mixture of silicon, copper and some other materials in real die, the overall κ could be somewhere around $130 \sim 400\text{W} \cdot \text{m}^{-1} \cdot \text{K}^{-1}$. The material in the motivation examples in simulation has κ of $400\text{W} \cdot \text{m}^{-1} \cdot \text{K}^{-1}$. In our case, the estimated κ by the proposed method is about $417\text{W} \cdot \text{m}^{-1} \cdot \text{K}^{-1}$, about 5% error for the estimation.

Therefore, we have verified the approach of estimating the power map in simulations. Moreover, the estimated power density maps sufficiently match the original power setting.

IV. Power Map Estimation for Real Microprocessors

As processor power in different locations changes over time, thermal map of CPU die also changes. Observed from thousands of thermal maps from multiple different workloads, thermal state can run into steady state in a few seconds after each power level change. Furthermore, during majority of running time CPU is running in steady state. As a result, we can estimate the CPU power map from the measured thermal map in the steady state. Based on the power map model derived in section III, we can also find the equivalent thermal conductivity κ of the die from the measurement of thermal maps.

For our work, the total power of CPU is also needed. Intel Performance Counter Monitor (PCM) provides users a software interface that computes the internal resource utilization of the latest Intel core processors. One metric of the PCM dataset is CPU energy consumption between two accesses. To ensure precision, power data has to be synchronized with the thermal maps. As mentioned in the system setup, if the capturing frequency of infrared camera is f, PCM data should be recorded in this same frequency. Suppose the CPU energy along discretized time points is series E, then total power $P = E/\Delta t$ is also a time series, and $\Delta t = 1/f$.

One thermal map is related to one negative-Laplacian map, and it will result in one κ value. Different thermal maps may result in different κ values. To obtain a reasonable κ, sufficient data samples from different workloads and time points are needed. A reasonable κ should be a constant despite different workloads and time. In this work, we execute eight workloads of different kinds and capture over 14000 thermal maps for each workload. Using the approach described in section III, κ computed with respect to different workloads along time line are plotted in Fig. 6. As we can see, κ comes out quite constant and even for different workloads although there is a large overlap, which is expected. Some glitches exist due to CPU changing power levels thus not running in steady state. During those times, temperature transient term $(\rho C_P \frac{\partial T}{\partial t})$ cannot be ignored. Furthermore, the arithmetic mean of κ is $174\text{W} \cdot \text{m}^{-1} \cdot \text{K}^{-1}$. We define this κ as the estimated equivalent thermal conductivity of CPU die.

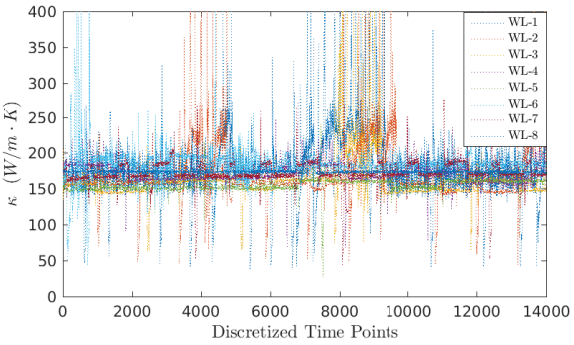

Fig. 6. Estimated thermal conductivity of die with respect to multiple workloads in time domain.

For real CPUs, as we do not know exact power density distribution, to validate the estimated power density map, we have compared their corresponding thermal maps. The idea is to build a thermal model and thermal simulation framework, which mimics the experimental set up of the chip die in the thermal imaging system with similar thermal boundary conditions and thermal structures. Thus, the validation flow can be summarized as following:

1) Obtain sufficient number of estimated power maps based on the proposed method. The experimental measurements should include an *idle* status, meaning CPU has

extremely low power, which will be used to set boundary conditions.

2) Build a finite element method (FEM) thermal simulation model that mimics the real structure of the processor die in the thermal imaging system.
3) Substitute the estimated thermal conductivity κ as well as the estimated power map into thermal simulation model as parameters and inputs.
4) Examine similarities between the simulated thermal maps and the experiment measured thermal maps. Higher degree of identity indicates higher precision of power map estimation, vice versa.

We will expand the details of implementing this validation flow in the result section.

V. Experimental Results and Discussions

In this section, we will present estimated power map results for the Intel i7-8650U, which has 4 CPU cores, and an integrated GPU. As all the existing works focus on component-wise power estimations, it is difficult to do the direct comparison. Instead, we validate the estimated power density map by finite element based thermal simulation using COMSOL Multiphysics as discussed earlier and show the results in this section [13].

A. Thermal model to imitate experiment setup

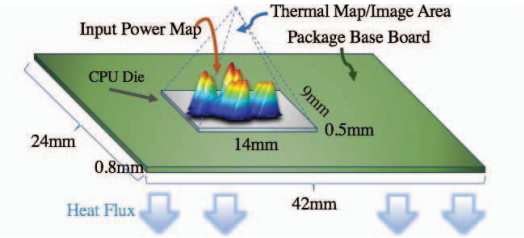

Fig. 7. Thermal model created to imitate the real experiment setup.

Once we obtain the estimated power density maps and equivalent thermal conductivity κ of the CPU die, we then start to build a heat transfer structure that mimics the real experiment setup. Fig. 7 illustrates the structure created using COMSOL Multiphysics which matches the real CPU package geometry. Geometries of Intel i7-8650U are acquired from the open resource from WikiChip organization [14]. The CPU package dimension is 42×24×1.3mm, thickness of base circuit board is 0.8mm. There are two dies soldered on the base board, the CPU die, which is the object we will study, has dimension 14×9×0.5mm. In our model, material of the CPU die and package base board are initialized with silicon and FR4 (Circuit Board), respectively. However, thermal conductivity of die part (silicon) is set to the computed κ, i.e. $\kappa = 174\text{W} \cdot \text{m}^{-1} \cdot \text{K}^{-1}$ in our case. A convective heat flux is applied to the bottom surface of the package base board, which simulates the heat flow from CPU package through motherboard to the cooling system, as indicated by the arrows at bottom in Fig. 7. The convective heat flux rate and thermal conductivity of base board will be determined as boundary conditions later.

For thermal simulation, we also need to know the correct thermal boundary conditions of the die. One idea is to explore the idle status of CPU (its boundary conditions are same of the CPU in other workload modes) as it is easy to extract the power map in this status. Specifically, since the idle status has extremely low power, power map is pre-known as almost zero,

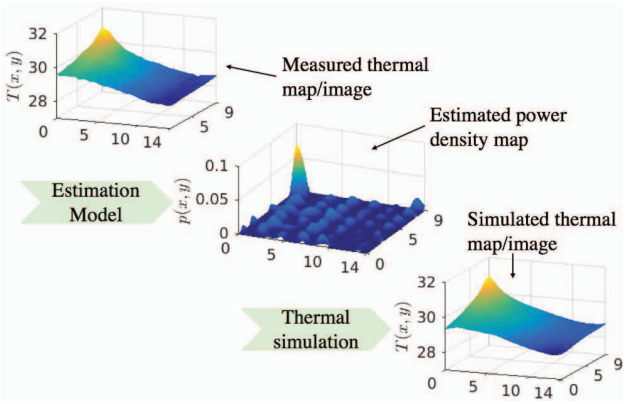

Fig. 8. Setting proper boundary conditions for thermal simulation model using measurements of CPU's idle status

except for very few places that have slight power. Majority of CPU appears to be approaching ambient temperature, spacial temperature appears relatively flat. At the beginning, simulated thermal map according to the estimated idle power map has the same trend with the measured thermal map, whereas the amplitude and the range have a little discrepancy. This will guide us how to adjust the thermal conductivity of package base board and bottom convective heat flux rate such that the simulated thermal map matches the measured thermal map as much as possible for *idle* status (Fig. 8). After some efforts, we found that convective heat flux rate is about $600\text{W} \cdot \text{m}^{-2} \cdot \text{K}^{-1}$ and thermal conductivity of base board is about $6\text{W} \cdot \text{m}^{-1} \cdot \text{K}^{-1}$

B. Accuracy study

We have examined sufficiently large amount of data samples that relate to multiple different steady states. For most data samples, the estimated power maps are able to duplicate the thermal maps that are almost identical to the measured thermal maps, with average similarities over 0.96. Besides, RMSE for thermal maps simulated by estimated power maps is 1.8°C. Fig. 9 lists four typical samples of the power map estimation results. The first row are the experiment measured thermal maps, the second row are the estimated power maps in 3D view, and the last row are thermal maps generated by FEM thermal simulation.

Fig. 10 further illustrates the power density distribution (Fig. 10(b)) on the processor die layout (Fig. 10(a)) at a steady state. Intel i7-8650U has system agent (which contains the Image Processing Unit (IPU), the Display Engine (DE), and the I/O bus), GPU module on the side and 4 core array in the middle. As we can see, 4 cores consume dominant power, whereas system agent and GPU module consumes low power at such steady state (Fig. 10(c)). We can see that power pattern aligns with the location of cores and Ring/Interconnect quite well. We obtain the component power by power density integration for the component, which are presented in Table. I.

TABLE I
Estimated Power for Processor Component (i7-8650U)

Component	Power	Component	Power
System Agent	0.70W	Ring/Interconnect	1.78W
GPU	0.59W	L3 cache	0.53W
Core#1	2.22W	Core#2	2.36W
Core#3	2.27W	Core#4	2.51W

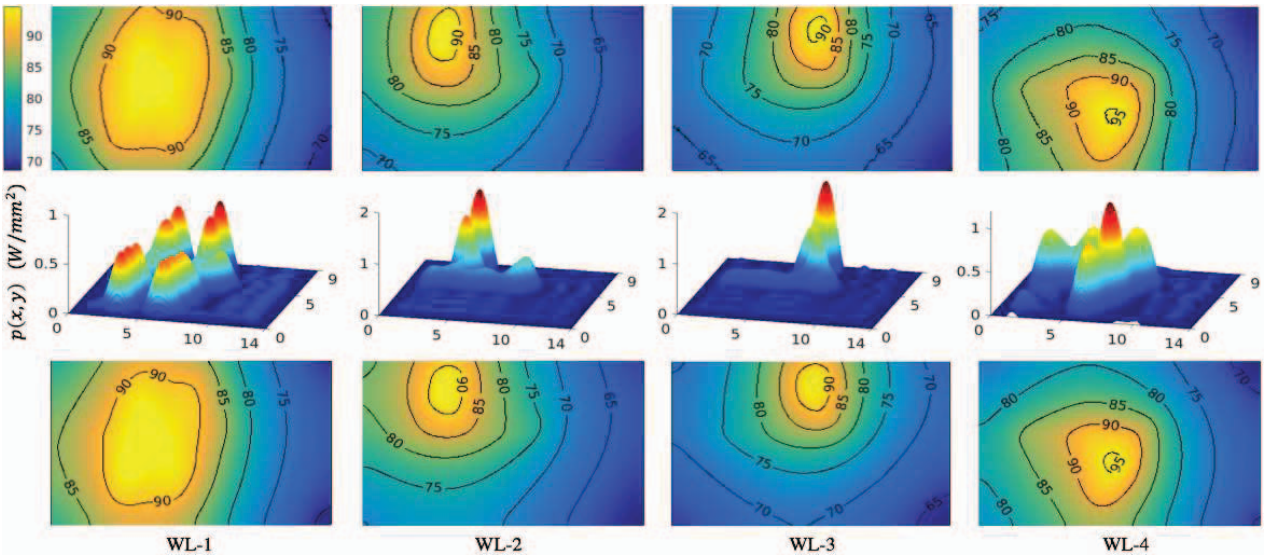

Fig. 9. The power map estimation results, from first to last row are experiment measured thermal maps, estimated power maps in 3D view, and simulation generated thermal maps, respectively. Each column is related to one workload at one steady state.

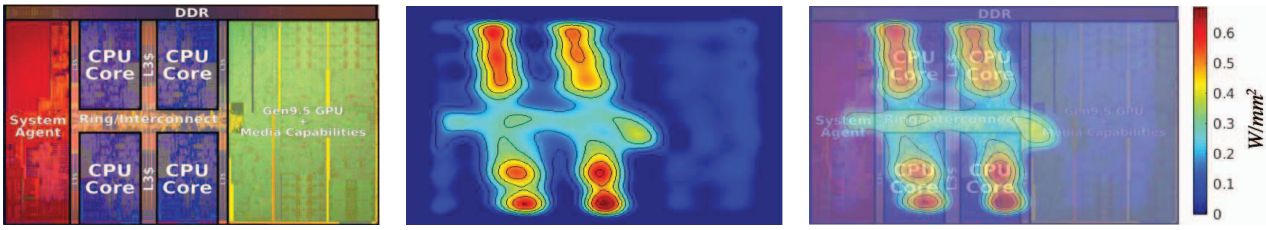

(a) (b) (c)

Fig. 10. (a) Intel i7-8650U processor die layout. (b) an estimated power density map. (c) projection of power density map onto the processor die layout

VI. Conclusion

In this paper, we have proposed an accurate full chip steady-state power density map estimation method for the commercial multi-core microprocessors. The new approach consists of several steps: 2D spatial Laplacian operation of measured thermal maps to generate the raw power density map, then computing the real power density map from the raw density map by re-distributing the total power of the chip. As a byproduct, we can also compute the real effective thermal conductivity of the microprocessor die. To further validate the estimated power map, finite element based thermal simulation was carried out based on the estimated power map and thermal models imitating the experimental set up of measurement used in the IR imaging system. Experimental results on intel i7-8650U 4-core processor shows that RMSE of the computed thermal maps and the measured thermal maps are around 1.8°C. Furthermore, computed thermal maps and the measured thermal maps have 96% similarity (2D correlation).

References

[1] H. Esmaeilzadeh, E. Blem, R. St. Amant, K. Sankaralingam, and D. Burger. Dark silicon and the end of multicore scaling. *Micro, IEEE*, 32(3):122–134, May 2012.

[2] Michael Taylor. A landscape of the new dark silicon design regime. *IEEE/ACM International Symposium on Microarchitecture*, 33(5):8–19, October 2013.

[3] David Brooks and Margaret Martonosi. Dynamic thermal management for high-performance microprocessors. In *Proc. IEEE Int. Symp. on High-Performance Computer Architecture (HPCA)*, pages 171–182, Jan. 2001.

[4] Vinay Hanumaiah and Sarma Vrudhula. Energy-efficient operation of multicore processors by DVFS, task migration, and active cooling. *IEEE Trans. on Computers*, 63(2):349–360, February 2014.

[5] Z. Liu, S. X.-D. Tan, X. Huang, and H. Wang. Task migrations for distributed thermal management considering transient effects. *IEEE Trans. on Very Large Scale Integration (VLSI) Systems*, 23(2):397–401, 2015.

[6] Hai Wang, Jian Ma, Sheldon X.-D. Tan, Chi Zhang, He Tang, Keheng Huang, and Zhenghong Zhang. Hierarchical dynamic thermal management method for high-performance many-core microprocessors. *ACM Trans. on Design Automation of Electronics Systems*, 22(1):1:1–1:21, July 2016.

[7] Russ Joseph and Margaret Martonosi. Run-time power estimation in high-performance microprocessors. In *Proc. Int. Symp. on Low Power Electronics and Design (ISLPED)*, pages 135–140, 2001.

[8] Canturk Isci and Margaret Martonosi. Runtime power monitoring in high-end processors: Methodology and empirical data. In *Proceedings of MICRO*, 2003.

[9] Wei Wu, Lingling Jin, Jun Yang, Pu Liu, and Sheldon X.-D. Tan. Efficient power modeling and software thermal sensing for runtime temperature monitoring. *ACM Trans. on Design Automation of Electronics Systems*, 12(3):1–29, 2007.

[10] K. Dev, A. N. Nowroz, and S. Reda. Power mapping and modeling of multi-core processors. In *International Symposium on Low Power Electronics and Design (ISLPED)*, pages 39–44, Sept 2013.

[11] S. Sadiqbatcha, H. Zhao, H. Amrouch, J. Henkel, and S. X.-D. Tan. Hot spot identification and system parameterized thermal modeling for multi-core processors through infrared thermal imaging. In *2019 Design, Automation Test in Europe Conference Exhibition (DATE)*, March 2019.

[12] H. Amrouch and J. Henkel. Lucid infrared thermography of thermally-constrained processors. In *2015 IEEE/ACM International Symposium on Low Power Electronics and Design (ISLPED)*, pages 347–352, July 2015.

[13] COMSOL Multiphysics. Heat transfer module user's guide. *COMSOL version*, 4, 2014.

[14] WikiChip. Core i7-8650u - intel. https://en.wikichip.org/wiki/intel/core_i7/i7-8650u.