Work-in-Progress: Determining MPSoC Layout from Thermal Camera Images

Michal Sojka
Czech Technical University in Prague
Prague, Czech republic

Ondřej Benedikt
Czech Technical University in Prague
Prague, Czech republic

Zdeněk Hanzálek
Czech Technical University in Prague
Prague, Czech republic

ABSTRACT

In many safety-critical applications, Multi-Processor Systems-on-Chip (MPSoC) must operate within a given thermal envelope under harsh environmental conditions. Meeting the thermal requirements often requires using advanced task allocation and scheduling techniques that are guided by detailed power models. This paper introduces a method that has the potential to simplify the creation of such models. It constructs so-called heat maps from thermal camera images. By comparing the heat maps of different workloads, we identify the locations of on-chip components and the amount of heat produced by them. We demonstrate our method on the i.MX8QuadMax chip from NXP, where we identify the locations of CPU clusters, bigger CPU cores, GPUs, and DRAM controllers.

KEYWORDS

thermal camera, MPSoC, NXP i.MX8, heat map, thermal-aware

ACM Reference Format:
Michal Sojka, Ondřej Benedikt, and Zdeněk Hanzálek. 2021. Work-in-Progress: Determining MPSoC Layout from Thermal Camera Images. In *2021 International Conference on Embedded Software Companion (EMSOFT'21 Companion), October 8–15, 2021, Virtual Event, USA.* ACM, New York, NY, USA, 2 pages. https://doi.org/10.1145/3477244.3477619

1 INTRODUCTION

Thermal-aware design methods have become increasingly important in the development of embedded systems. In many applications, e.g., in future avionics systems, Multi-Processor Systems-on-Chip are needed for their high computing performance. However, at the same time, these MPSoCs must operate within a given thermal envelope under harsh environmental conditions. Meeting the stringent thermal requirements is often in contradiction with the high performance demands. For this reason, system designers look for methods that decrease the operating temperature while maintaining computational performance. The effectiveness of such methods often depends on the availability of accurate power models.

In this paper, we introduce an easily applicable method that has the potential to improve the quality of existing power models. It reveals interesting details about the MPSoC chip layout as well as the location and amount of heat generated by the given workload at different parts of the chip. The method is based on images from a thermal camera captured while running the workload. Compared to previously published works on a similar topic [2, 5], our method does not require mounting a special cooling device on the analyzed board. To demonstrate our method, we analyze the chip layout of NXP i.MX8QuadMax [3], which is an ARMv8 chip with six CPUs in big.LITTLE configuration (4× Cortex A53 and 2× Cortex A72) and 2 GPUs with a total of 16 shader processors with 64 execution units. This chip is not yet officially released by NXP, and little information about it is publicly available.

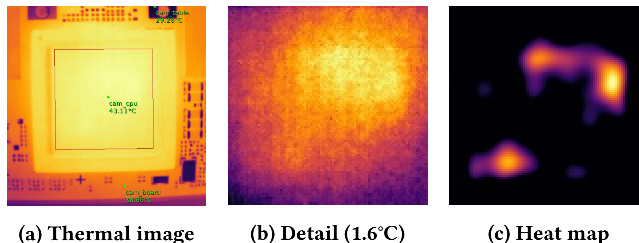

(a) Thermal image **(b) Detail (1.6°C)** **(c) Heat map**

Figure 1: Thermal camera image processing (idle system)

We publish the software for online image processing of the thermal images as open-source [4] to allow other researchers to analyze different chips and using the results in their work.

We believe that the obtained results will allow us to design better power models to guide various thermal-aware workload allocation and scheduling methods [1].

2 METHOD

The method to obtain the chip layout and the corresponding power density map is based on images from a thermal camera. We use a Workswell WIC camera with a spatial resolution of 336x256px, frame rate 9 Hz and thermal resolution 40 mK. The camera takes images of the board (see Fig. 1a). The location of the chip is marked with the red square whose edges are 20 mm long. We transform this area to a regular square of 100×100 pixels (see Fig. 1b). Then we use a method inspired by Zhang et al. [5]. By using a heat diffusion equation and considering only steady state, the spatial heat energy g_T generated at point (x, y) is

$$g_T(x, y) = -\kappa \nabla^2 T, \qquad (1)$$

where κ is thermal conductivity, ∇^2 is the Laplace operator, and T is the spatial temperature profile, i.e., Fig. 1b. Equation (1) allows determining the location of heat sources directly from the thermal camera image. As the image is noisy, we apply a Gaussian blur filter before calculating the Laplacian. Thermal conductivity κ can be determined as described in [5], but in this paper, we use $\kappa = 100$ to avoid too small numbers in our graphs (Fig. 2). We are interested in determining the chip layout and not the exact amount of generated

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
EMSOFT'21 Companion, October 8–15, 2021, Virtual Event, USA
© 2021 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-1-4503-8712-5/21/10…$15.00
https://doi.org/10.1145/3477244.3477619

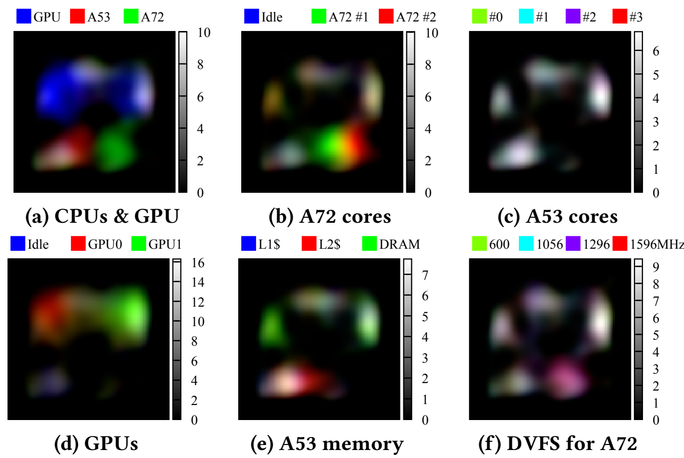

(a) CPUs & GPU (b) A72 cores (c) A53 cores

(d) GPUs (e) A53 memory (f) DVFS for A72

Figure 2: Comparison of heat maps for different workloads

heat. Despite Gaussian filtration, the resulting heat energy map g contains noise. Therefore, in subsequent processing, we work with the thresholded value $g'_T(x, y) = \max(0, g_T(x, y) - g_0)$. The constant g_0 is determined as follows: When no power is supplied to the board, the value of g represents only the noise. We set g_0 to be the maximum value of the observed noise so that for the powered-off board, we get $g'_T(x, y) = 0,\ \forall x, y$. The values of $g'_T(x, y)$, called *heat map* in the following, are shown in Fig. 1c.

The heat maps calculated from subsequent frames of the same workload are not identical. To increase the reproducibility of our results, we average the results from multiple frames. The heatmaps in Fig. 2 represent averages over 2 minutes of execution, however, even a single-frame heat map (Fig. 1c) provides acceptable precision.

3 RESULTS

We applied our method to the Toradex Apalis i.MX8 board running Yocto Linux distribution. In the idle state (Fig. 1) only the necessary services were running (systemd, dbus, agetty, dropbear). The idle power consumption of the whole board was 5.2 W.

We run different workloads our Themobench suite, collect the heat map g'_T for each workload, and compare the heatmaps of different workloads by combining them to a single image – see Fig. 2. Each heatmap is drawn with a different hue; the intensity of the color is proportional to the amount of generated heat. The hues of different components are selected so that the combination of all used hues with the same intensity produces a shade of gray.

The location of main MPSoC components can be seen in Fig. 2a. The GPU heatmap was produced by running a compute-bound OpenCL program on the GPU. CPU heatmaps result from running a workload utilizing all CPU cores of the same type and the associated L2 cache memory.

Fig. 2b shows the difference between the two A72 cores running the tinyrenderer application. The individual cores can be clearly distinguished. On the other hand, when running the same workload on A53 cores (Fig. 2c), we observe no significant difference between the cores. By comparing Figs. 2b and 2c, we see that A72s produce more heat than A53s, which matches the typical characteristics of big.LITTLE cores.

To reveal the layout of the GPU, we constructed an OpenCL kernel that performs compute-bound work in the selected GPU workgroups and keeps the rest idle. Figure 2d compares the execution in workgroups 0–7 and 8–15, which appears to correspond to the two available GPUs. When we tried to show the differences between individual workgroups or their couples and quadruples, no significant differences were observed, and the resulting figures were mostly gray like Fig. 2c.

Fig. 2e shows the results of a benchmark stressing different parts of the memory hierarchy on an A53 core. We can observe that the L2 cache is located at the right of CPU cores and that the two DRAM memory controllers are located on the left and right side approximately in the middle. The DRAM controller on the right is co-located with other parts of the chip, which are "always-on".

The effect of dynamic voltage and frequency scaling (DVFS) for A72 cores can be seen in Fig. 2f. The CPU frequency influences the heat generation only at the location of the A72 cores and not at the other parts of the chip.

The chip layout presented in Fig. 2 matches pretty well the locations of the power supply pins in the SoC datasheet [3].

We performed the experiments both with and without the fan-induced airflow. The airflow prevents the chip from overheating during some experiments. The results (g'_T) do not depend on the presence of the airflow; the locations of chip components were the same in both cases.

4 CONCLUSION AND FUTURE WORK

We have described a method to obtain a so-called chip heat map from thermal camera images. This method does not require mounting any cooling devices used in similar works. Furthermore, software performing all necessary calculations in real-time is available as open-source [4]. We showed the applicability of our method on the NXP i.MX8QuadMax chip, whose complete documentation is not yet publicly released, and demonstrated that the location of some chip elements (CPU clusters, bigger CPU cores, GPUs, and DRAM controllers) could be clearly identified.

We plan to use the heat maps of various workloads to construct accurate per-workload power models. Using such models in thermal-aware task allocation and scheduling techniques should reduce the operating temperature, which is a vital property of many safety-critical embedded systems.

ACKNOWLEDGMENTS

This research has received funding from the Clean Sky 2 Joint Undertakingnder the European Union's H2020 research and innovation programme under grant agreement No 832011 (THERMAC).

REFERENCES

[1] O. Benedikt, M. Sojka, P. Zaykov, et al. 2021. Thermal-Aware Scheduling for MPSoC in the Avionics Domain: Tooling and Initial Results. In *The 27th IEEE International Conference on Embedded and Real-Time Computing Systems and Applications*. IEEE.
[2] Dev, Paul, Huang, et al. 2018. Implications of Integrated CPU-GPU Processors on Thermal and Power Management Techniques. (Aug. 2018). arXiv:1808.09651 [cs]
[3] NXP Semiconductors. 2020. *i.MX 8QuadMax Automotive and Infotainment Applications Processors*. Data Sheet IMX8QMAEC.
[4] M. Sojka. 2021. Thermocam-PCB. https://github.com/CTU-IIG/thermocam-pcb
[5] J. Zhang, S. Sadiqbatcha, W. Jin, and S. X.-D. Tan. 2020. Accurate Power Density Map Estimation for Commercial Multi-Core Microprocessors. In *DATE*. 1085–1090. https://doi.org/10.23919/DATE48585.2020.9116545