RETROSPECTIVE: Power model validation through thermal measurements

Jose Renau
Dept of Computer Science and Engineering
UC Santa Cruz
renau@ucsc.edu

I. Context

Research on power and temperature in computer architecture became a primary focus for the MASC group from 2005 to 2006. It quickly became clear that establishing a ground truth was critical to developing more detailed power and thermal models. The primary challenge was the need for detailed breakdowns, but access to RTL/layout data was not feasible with cutting-edge commercial processors. Since we lacked RTL access and could not measure power consumption directly for each block, we decided to examine the temperature impact that arose from power usage.

We began exploring the use of infrared cameras, yet faced several challenges: cooling the chip after removing the heatsink, and managing heat diffusion caused by the packaging and silicon substrate.

Our early prototypes were lower power chips, in which we could lower the voltage to maintain a functional design without a heatsink. We soon started delidding the chips. It became apparent that some chips, such as AMD, were simpler to handle as Intel chips utilized a solder-based thermal interface material. Although not stated in the ISCA paper, we did send some Intel chips for polishing to remove the solder residue. However, we did not include these results in the paper due to the challenges posed by heat diffusion.

II. Cooling down

Once we had some cleanly delidded packages, we commenced chip measurements. For this ISCA work, the AMD Athlon 64 with a 130nm SOI became our primary focus, although we also assessed Xilinx FPGAs, as they were easy to operate at low power levels.

To produce representative results, it was necessary to devise an infrared-transparent cooling setup. However, budgetary constraints at the time did not permit the purchase of an infrared camera, so we opted to rent one. We allocated funds for two rental days, a window within which we had to decide whether to proceed with or abort the project. On the first day, we prepared numerous chips for measurement, but the main aim was to explore effective methods for chip cooling. Since our lab was not rated for handling hazardous chemicals, we sought potential IR transparent liquids, testing commonplace kitchen liquids like pure alcohol, vinegar, and olive oil. To our surprise, olive oil was quite effective, laying the foundation for refining our cooling medium in future experiments. This led us to opt for purer oils, such as medical-grade diuretics or Fluka Mineral Oil 69808.

In this initial setup, we demonstrated that we could cool the chip using a liquid flow, though we had concerns about the turbulence potentially complicating the modeling process. We soon discovered that sapphire is infrared transparent, prompting us to construct a heatsink with two sapphire windows, as depicted in Figure 1-(a). We aimed to circulate oil between the two sapphire windows. However, this did not work as anticipated. First, it could not sufficiently cool the chip at nominal voltage. Second, maintaining good contact between the silicon substrate and the die proved difficult. Consequently, we abandoned this approach for the ISCA paper but utilized a simplified version in a subsequent ASPLOS paper [2].

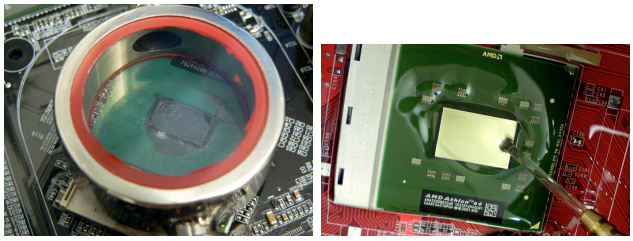

Fig. 1. (a) Failed prototype of saphire chamber not used in the paper. (b) An early oil-flow picture with a non-aligned flow.

III. Measuring

The final measurement setup is depicted in Figure 2. It utilizes an oil flow over the delidded processor, monitored by a FLIR SC-4000 camera. Additionally, we used a multimeter to trace the overall chip power consumption and a pump and radiator to maintain a constant oil temperature. In later studies, we expanded the reservoir size to circumvent the need for the radiator, which occasionally experienced ”leakage”.

The camera had to be a Midwave Infrared (MWIR) model, operating in the 3-5um range because the AMD and Xilinx silicon substrates were transparent at these frequencies. Interestingly, Intel chips were not transparent within this range, a factor that we could not determine if due to polishing or doping. Operating in this frequency range allowed us to measure the junction temperature rather than the top of the silicon, which is blurred due to heat diffusion.

After conducting numerous measurements, we noted certain distortions and realized that the camera required calibration

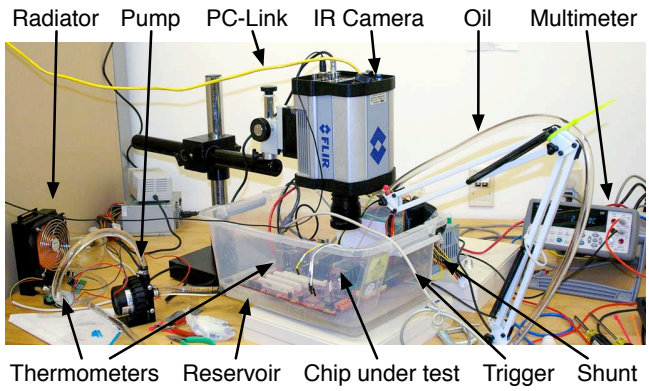

Fig. 2. Measuring setup with an oil-based heatsink with laminar flow.

for each chip due to various factors. The solution involved applying a simple linear correction per pixel after constantly heating the die at several points while the chip was inactive. Figure 3 exhibits the hottest average temperature measured, superimposed on the AMD Athlon floor plan.

The preliminary measurements using oil flow, as seen in Figure 2-(b), was used for the ISCA research. The oil flow had to be carefully aligned and regulated to produce a laminar flow, but it performed quite effectively once this setup was achieved. However, any adjustments to the oil flow necessitated a comprehensive recalibration.

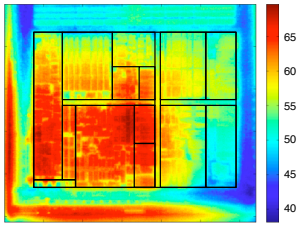

Fig. 3. Highest average temperature frame when running SPEC.

IV. Modeling and Power

Having acquired the ability to measure temperature with a resolution of 100x100um per pixel at a rate of 125 frames per second, our next objective was to ascertain power consumption. To this end, we extended the functionality of HotSpot [3] to model the new oil-based "heatsink". Our setup was carefully adjusted to produce a laminar flow, which is simpler to model.

Leveraging a previous publication that featured the AMD Athlon floorplan, we formulated a power equation (Equation 1) for each block of the floorplan.

$$\texttt{Power} = P_{clock} + P_{op} * AR + P_{leak0} * T^2 * e^{(P_{leak1}/T + P_{leak2})} * (1 - e^{P_{leak3}/T}) \quad (1)$$

We ran numerous benchmarks under varying power temperatures to ascertain the equation parameters. We aimed to capture frames where blocks displayed high and low activity rates at different operating temperatures. The motivation behind this was the significant temperature dependence of leakage, enabling us to determine the dynamic/power breakdown. A genetic algorithm was employed to identify the optimal values best fit the multimeter power and infrared camera temperature readings. The outcome of these measurements allowed us to establish the maximum and minimum dynamic power for each block and to create a breakdown of leakage.

V. Achievements

The most closely related work [1] employed infrared cameras with steady-state power consumption. This paper, however, pioneered a methodology for measuring temperature and extracting dynamic power consumption with fine granularity for an off-the-shelf processor.

There were significant developments following this paper. We assisted NVIDIA and Sun Microsystems in measuring some of their chips. This system was utilized to gain insights and verify power consumption. Sun Microsystems was particularly helpful, providing funding and special packages used in a subsequent paper [2]. This led to the development of a controlled setup that greatly enhanced measurement accuracy. As a result, we developed a sapphire silicon heatsink that closely mimicked the properties of traditional metal heatsinks. This addressed a potential shortcoming of the ISCA paper: it did not operate under typical temperature conditions but instead utilized a new "oil"-based heatsink. While this was sufficient for obtaining power breakdowns, the operating conditions were not typical. Our follow-up work enabled chip measurements under typical operating conditions by developing a sapphire heatsink.

Acknowledgements

The authors of this paper are Francisco J. Mesa-Martínez and Joseph Nayfach-Battilana. We thank Olivier Temam, James Christofferson, Ali Shakouri, Michael Huang, Luigi Capodieci, and Matthew Guthaus for their insightful feedback. Any opinions, findings, conclusions or recommendations expressed in this material are solely those of the authors and do not necessarily reflect the views of the National Science Foundation. The National Science Foundation partly supported this work under grant 0546819, and gifts from SUN Microsystems.

References

[1] H. Hamann, J. Lacey, A. Weger, and J. Wakil, "Spatially-resolved imaging of microprocessor power (simp): hotspots in microprocessors," in *Thermal and Thermomechanical Phenomena in Electronics Systems, 2006. ITHERM'06. The Tenth Intersociety Conference on*. IEEE, 2006, pp. 5–pp.

[2] F. Mesa-Martínez, E. K. Ardestani, and J. Renau, "Characterizing processor thermal behavior," in *ACM SIGARCH Computer Architecture News*, vol. 38, no. 1. ACM, 2010, pp. 193–204.

[3] K. Skadron, M. Stan, W. Huang, S. Velusamy, K. Sankaranarayanan, and D. Tarjan, "Temperature-aware microarchitecture," in *ACM SIGARCH Computer Architecture News*, vol. 31, no. 2. ACM, 2003, pp. 2–13.