

Thermal and Power Characterization of Field-Programmable Gate Arrays

Abdullah Nazma Nowroz
School of Engineering
Brown University
Providence, RI 02912
Email: abdullah_nowroz@brown.edu

Sherief Reda
School of Engineering
Brown University
Providence, RI 02912
Email: sherief_reda@brown.edu

ABSTRACT

In this paper we propose new techniques for thermal and power characterization of Field Programmable Gate Arrays (FPGAs) using infrared imaging techniques. For thermal characterization, we capture the thermal emissions from the backside of an FPGA chip during operation. We analyze the captured emissions and quantify the extent of thermal gradients and hot spots in FPGAs. Given that FPGAs are fabricated with no knowledge of the potential field designs, we propose soft sensing techniques that can combine the measurements of hard sensors to accurately estimate the temperatures where no sensors are embedded. For power characterization, we propose algorithmic techniques to invert the thermal emissions from FPGAs into spatial power estimates. We demonstrate how this technique can be used to produce spatial power maps of soft processors during operation.

ACM Categories & Subject Descriptors

B.7.1 [Integrated Circuits]: Types and Design Styles.

General Terms: Design, Performance, Algorithms.

Keywords: Thermal characterization, sensors, power mapping.

1. INTRODUCTION

Given that FPGA-based implementations are less area and power efficient than their ASIC counterparts, power characterization for FPGAs is an active topic of research recent literature [11, 4]. The challenges in power characterization of heterogeneous FPGAs pose an interesting problem as these chips host programmable soft or hard processors. The power consumption of these processors is not entirely determined during design time as runtime workloads can impact the exact power consumption. Power consumption leads to heat dissipation which increases junction temperatures. Elevated temperatures and hot spots directly impact all key circuit metrics, including: lifetime and reliability, speed, power, and costs. Thus, it is necessary to conduct comprehensive thermal characterization of FPGAs to reduce the risk of increased temperature operation.

Previous works in FPGA power characterization focused on modeling techniques to estimate the power consumption during design time [11, 4, 10]. Compared to ASICs, the interconnect structure

of FPGAs (especially long routing tracks) consumes a significant amount of dynamic power [10]. A recent work describes a method to track percentage changes in dynamic power by tracking percentage changes in the frequencies of nearby ring oscillators [12]. With respect to thermal characterization, a number of papers attempt to quantify the magnitude of thermal gradients in FPGA chips [5, 12]. Due to the lack of knowledge of hot spot locations in FPGAs, many thermal sensor placement methods spatially allocate the sensors in a uniform way across the die [5, 12]. If the knowledge of a target set of benchmarks is first available, then benchmark-specific sensor insertion techniques can be used [8, 3]. Lopez-Buedo *et al.* propose reconfigurable thermal sensors using ring oscillators [6].

In recent years infrared imaging has received increased attention as a powerful tool for thermal and power characterization of real chips [2, 7, 9]. Despite the research efforts on the application of thermal imagery to processors, there are no reported results in the literature for applying infrared imaging to FPGAs. The objective of this paper is to provide thermal and power characterization techniques for FPGAs from the infrared emissions emitted from the backside of the silicon during operation. The main contributions of this paper are as follows.

- Previously published results on FPGA power and thermal characterization and modeling used either simulations or measurements from a limited number of on-die sensors. This paper complements the research of previous papers by providing thermal and power characterization methods based on infrared emissions emitted from FPGAs. Infrared imaging has not been applied before to FPGAs (at least in the literature), and our paper is the first to make such contribution.
- We quantify the extent of thermal gradients during runtime, and we propose a new technique, *soft sensing*, that can utilize the measurements of hard sensors to accurately estimate the temperatures where no sensors are embedded.
- We propose a power characterization methodology that inverts the spatial thermal emissions from FPGAs into power maps.
- Using a real FPGA chip, we demonstrate the effectiveness of applying our proposed thermal and power characterization methodologies. We estimate the spatial power maps of a soft processor during runtime.

The organization of this paper is as follows. We describe our thermal characterization and sensor allocation techniques in Section 2. Our power characterization technique is described in Section 3. In Section 4 we present our experimental results for both thermal and power characterization. Finally, Section 5 summarizes the main conclusions of this work.

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. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

FPGA'11, February 27–March 1, 2011, Monterey, California, USA.

Copyright 2011 ACM 978-1-4503-0554-9/11/02 ...\$10.00.

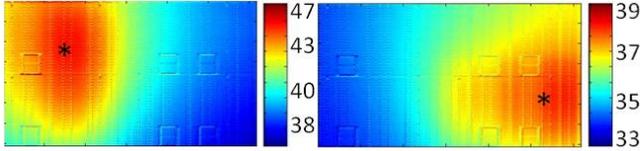


Figure 1: Example of possible thermal gradients in an FPGA.

2. THERMAL CHARACTERIZATION

The fact that FPGA functionality is programmable implies that a diverse range of possible thermal gradients could exist during operation. For thermal characterization, we consider a section of the FPGA die which is approximately 20 mm in width and 10 mm in length. We embed two different designs and capture the thermal emissions as given in Figure 1. The figure demonstrates that the magnitude and location of the maximum temperature, or hot spot, vary significantly among the two traces. For instance, the thermal trace of the first figure shows a maximum temperature of 47°C with a thermal gradient of about 9°C between the hottest and the coldest part of the die’s section. The second trace shows a maximum temperature of 39.5°C and a difference of 7°C between the hottest and coldest spots. Furthermore, the hot spot in the right trace is about 14 mm away from the location of the hot spot of the left trace. These variations make the task of allocating thermal sensors particularly challenging.

To enable effective thermal management of FPGAs, it is necessary to insert thermal sensors at potential hot spot locations. All current FPGAs (e.g., Altera’s Stratix family or Xilinx’s Virtex family) use only one thermal sensor. The possibility of large thermal gradients on the FPGA die as demonstrated in Figure 1 implies that the measurements of one sensor will not necessarily capture the maximum temperature of the FPGA chip.

Because FPGAs are manufactured to accommodate a large number of possible designs, it is not possible to know the locations of hot spots *a priori* during manufacturing time as hot spot locations depend on the design and the design’s layout configured into the FPGA. Thus, hot spot locations can only be known after a particular design is programmed into the FPGA. As the locations of hard sensors are fixed into the manufactured FPGA, the locations of these sensors could be far from the true hot spot locations. Increasing the number of sensors is not a feasible idea as digital thermal sensors consume die area because their need to accommodate analog to digital converters, and thus designers limit their numbers. To circumvent this limitation, we propose the idea of *soft sensors* to augment the placed hard sensors measurements and improve thermal tracking by estimating the temperature at locations where no sensor are embedded.

In our proposed technique, a soft sensor measurement is equal to a weighted linear combination of the measurements of the hard sensors. Thus, if $\hat{t}(l)$ denotes the estimated soft sensor temperature at location l then this estimated temperature can be expressed by

$$\hat{t}(l) = \sum_{i=1}^k w(l, s_i) \mathbf{t}_m(s_i), \quad (1)$$

where $\mathbf{t}_m(s_i)$ ’s are the measurements reported by the k sensors and $w(l, s_1), \dots, w(l, s_k)$ are the weights corresponding to the k sensors that change depending on the location, l , where the temperature needs to be estimated. To determine the best set of weights for some location l , thermal characterization traces collected from the infrared imaging system or from simulation can be used to learn the

optimal way to combine the measurements of the hard sensors. Let the collected thermal traces be denoted by $\mathbf{t}_m^1, \dots, \mathbf{t}_m^n$. Given the n traces, we can construct the following set of equations:

$$\begin{aligned} \mathbf{C}\mathbf{w}(l) &= \begin{pmatrix} \mathbf{t}_m^1(s_1) & \cdots & \mathbf{t}_m^1(s_k) \\ \vdots & & \vdots \\ \mathbf{t}_m^n(s_1) & \cdots & \mathbf{t}_m^n(s_k) \end{pmatrix} \begin{pmatrix} w(l, s_1) \\ \vdots \\ w(l, s_k) \end{pmatrix} \\ &= \begin{pmatrix} \mathbf{t}_m^1(l) \\ \vdots \\ \mathbf{t}_m^n(l) \end{pmatrix} = \mathbf{T}(l), \end{aligned} \quad (2)$$

which can be written succinctly in matrix notation as $\mathbf{C}\mathbf{w}(l) = \mathbf{T}(l)$. The best set of weights $\mathbf{w}(l)$ that minimizes the total least square error can be computed as follows:

$$\mathbf{w}(l) = (\mathbf{C}^T \mathbf{C})^{-1} \mathbf{C}^T \mathbf{T}(l), \quad (3)$$

where T indicates the matrix transpose operation. Computing the optimal weights should be carried only once off-line either during (1) design time using the results from thermal modeling and simulation tools, or (2) after fabrication where infrared imaging techniques are typically used to provide the necessary characterization and calibration of the embedded thermal sensors. The computed weights can be then distributed by the FPGA vendor as part of their tool chain. In current FPGA tools, the user must select the exact target device of the tool before synthesizing any design. Thus, storing a set of soft sensor weights for each device fits naturally with existing FPGA tool chain flows. During design time, the FPGA tool can create the necessary soft sensors based on the placement results of the programmed design or at any desired input locations supplied by the user. The physical implementation of a soft sensor requires only one adder-accumulator multiplier to compute the weighted sum as described by Equation (1).

The proposed soft sensing technique has advantages over previously proposed reconfigurable sensors based on ring oscillators [6, 12]. Ring oscillator based sensors are affected by process variability and operating voltages and thus thermal calibration needs to be conducted on every manufactured device independently. On contrast, the weights of the proposed soft sensing technique need to be estimated only once for an entire device family. These weights reflect the thermal conductances between different parts of the die which are not affected by underlying nano-scale electrical variations. Furthermore, these weights can be determined either using infrared measurements (as in our approach) or from the final design layout using thermal modeling tools. Once these weights are determined, they are distributed with the vendor’s tools as part of the characteristics of the device family.

3. POWER CHARACTERIZATION

The physical relationship between power and temperature is governed by the heat diffusion equation. In steady state, the heat equation is given by

$$\nabla \cdot (k(x, y) \nabla t(x, y)) - ht(x, y) = -p(x, y), \quad (4)$$

where $k(x, y)$ is the thermal conductivity at location (x, y) , $t(x, y)$ is temperature at location (x, y) relative to the ambient temperature, $p(x, y)$ is the power density at location (x, y) , and h is the heat transfer coefficient. The explicit use of location in the thermal conductance $k(x, y)$ is important as different parts of the chip have different material composition which impacts lateral heat diffusion.

In any practical implementation, the heat equation has to be discretized. This discretization comes from the finite memory size of any computer and more importantly from the use of thermal imaging equipment with limited spatial resolution. In a discretized form, the continuous temperature signal $t(x,y)$ is represented by a vector \mathbf{t} that gives the temperatures at a discrete set of die locations, and the continuous power signal $p(x,y)$ is represented by a vector \mathbf{p} that gives the power at the same set of discrete die locations. In such case, Equation (4) is approximated by the following linear matrix formulation

$$\mathbf{R}\mathbf{p} + \mathbf{e} = \mathbf{t} + \mathbf{e} = \mathbf{t}_m, \quad (5)$$

where the matrix \mathbf{R} gives the thermal resistivity between different discrete locations of the die, the vector \mathbf{e} denotes the errors in measurements created by the infrared imaging system, and \mathbf{t}_m denotes the vector of measured temperatures.

The objective of the *thermal to power inversion problem* is to find the vector \mathbf{p} that leads to the temperatures \mathbf{t} . Given the thermal measurements \mathbf{t}_m from the infrared camera, the objective is to find the best power pattern \mathbf{p} that minimizes the total squared error between the temperatures as computed using Equation (5) and the measured temperatures; that is,

$$\arg_{\mathbf{p}} \min \|\mathbf{R}\mathbf{p} - \mathbf{t}_m\|_2^2, \quad (6)$$

where $\|\cdot\|_2$ indicate the L_2 norm, and under the constraints that the sum of the elements of the power vector is equal to the total power consumption of the chip p_{total} and that the power estimates must be greater than zero; i.e.,

$$\|\mathbf{p}\|_1 = \sum_i p_i = p_{total} \text{ and } \mathbf{p} \geq \mathbf{0}, \quad (7)$$

where $\|\cdot\|_1$ indicate the L_1 norm and p_{total} is the total power consumption of the chip. We use MATLAB's quadric optimization solver (`lsqlin`) to minimize (6) under the constraints of Equations (7). The solver uses the active-set strategy (also known as a projection method) which relies on two step solution. The first step calculates a feasible solution point, and the second phase generates an iterative sequence of feasible solution points that converge to the final solution.

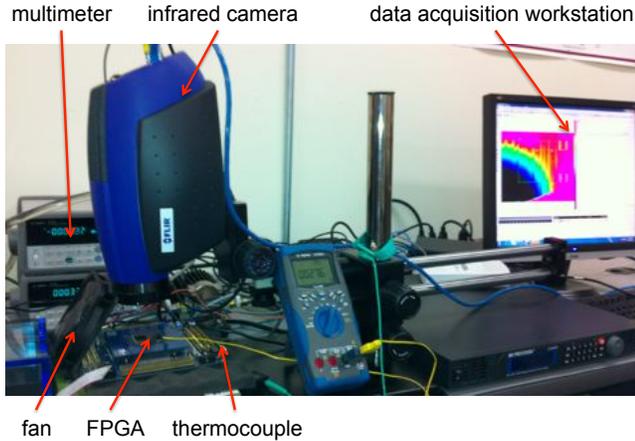


Figure 2: Experimental setup.

4. EXPERIMENTAL RESULTS

We use a 90 nm Altera Stratix II EP2S180 FPGA with 180,000 logic elements. For our experiments, we only consider a portion the die of dimensions 20 mm \times 10 mm. To capture the thermal emissions from the backside of the FPGA, we use a FLIR SC5600 infrared camera with a mid-wave spectral range of 2.5 μm – 5.1 μm . The camera is capable of operating at 100 Hz with a spatial resolution of 30 μm with a 0.5 \times microscopy kit. The camera's cryogenic cooled InSb detectors have a sensitivity of about 20 mK noise equivalent temperature difference. We devise a pixel-by-pixel calibration process that translates the captured emissions (as measured by the digital levels of the camera's analog to digital converters) to temperatures to get an accurate thermal imagery. For total power measurement, we externally intercept the current to the FPGA by a 1 m Ω shunt resistor. The voltage (1 mV for every 1 A) across the shunt resistor is measured through an Agilent 34410 digital multimeter. Our setup is illustrated in Figure 2.

The objective of the first experiment is to demonstrate the advantages of using soft sensors to augment hard sensors. For the hard sensors, we consider their locations to be allocated in a regular and uniform way across the die. We first collect tens of thermal traces that arise from different designs embedded in the FPGA. These designs have different spatial power allocations which vary the thermal characteristics. To quantify the error in hot spot tracking, we define the *thermal error* for a thermal trace as the difference between the maximum of the measurements at the locations of the thermal sensors and the maximum temperature in the trace. We report the *worst thermal error* across all traces as a function of the number of sensors in Figure 3 for both hard and soft sensors. The results show that soft sensing can cut the tracking error by 40% bringing the error to within less than 1 $^\circ\text{C}$ with just three sensors.

In the second experiment we evaluate our spatial power mapping technique by inverting the thermal emissions of the Nios II soft processor, while running the standard Dhrystone 2.1 application, into spatial power estimates. We consider two different configurations of the Nios II processor: the standard model Nios II/s with multipliers implemented in the FPGA's DSP blocks and the full-performance model Nios II/f. We first constrain the logic blocks of the Nios II processor to fit into a 30 \times 30 array of logic blocks (about 6.4 mm \times 7.0 mm) of the layout as shown in Figure 4.a and then capture the steady-state thermal emissions \mathbf{t}_m from this area as shown in Figure 4.b. The total power consumption of the two configurations are 315 mW (Nios II/s) and 477 mW (Nios II/f).

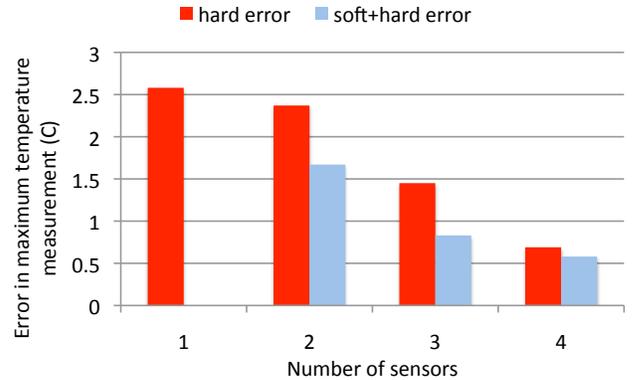


Figure 3: Error in temperature tracking in $^\circ\text{C}$ as a function of the number of sensors using soft sensors.

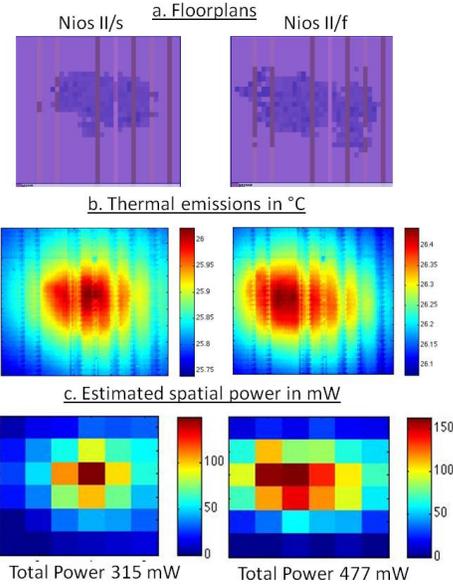


Figure 4: Spatial power estimates of Nios II processors running Dhrystone 2.1 application.

For the purpose of spatial power mapping, we discretize the layout area into 6×6 regions, and thus the \mathbf{p} vector is comprised of 36 elements that need to be estimated.

To compute the power map \mathbf{p} from the thermal emissions \mathbf{t}_m , we need to estimate the matrix \mathbf{R} as given by Equation (5). We estimate the matrix \mathbf{R} in a column-by-column basis. We note that the k^{th} column of matrix \mathbf{R} can be obtained by setting the vector \mathbf{p} to be equal to $[0 \ 0 \ \dots \ 1 \ \dots \ 0]^T$, where the "1" is at the k^{th} location of the \mathbf{p} vector, and then use the resultant emissions \mathbf{t}_k directly as the k^{th} column of matrix \mathbf{R} . To realize this setting, we utilize the fact that FPGAs are programmable. For each power region k , we embed, and turn on, ring oscillators precisely into the logic array blocks that are available in the region, while the rest of the blocks in the design are inactive. Such precise embedding is possible with Altera's Quartus II tool. The resultant thermal emissions \mathbf{t}_k from such embedding are then normalized by the total power p_k that is measured externally through the digital multimeter. Thus, column k of matrix \mathbf{R} is equal to \mathbf{t}_k/p_k . We automate the whole process in order to measure the 36 columns of \mathbf{R} with fast turn-around time. Our earlier work describes the automated embedding procedure in more details [1].

With the estimated matrix \mathbf{R} and thermal emissions from the Nios II processor (Figure 4.b), we compute the spatial power maps using the procedure given in Section 3. The estimated spatial power maps in mW are plotted in Figure 4.c. Our estimated spatial maps augment the floorplan with valuable spatial power density estimates. The revealed detailed power estimation maps can be used to calibrate the estimates from high-level dynamic power modeling tools. Given that the design of the Nios II processor is proprietary, it is not possible for us to match the spatial power consumption estimates to the various functional blocks of the processor.

5. CONCLUSIONS

In this paper we have presented new techniques for thermal and power characterization of FPGAs using infrared emissions. We described experimental techniques to capture the infrared emissions from the backside of FPGAs. Using the thermal images, we elucidated the magnitude of thermal gradients and hot spots in FPGAs.

We have demonstrated that up to 9°C gradients can exist within our tested FPGA. To circumvent the lack of knowledge of hot spots in FPGAs and improve thermal tracking, we proposed a soft sensor technique that combines the measurements of hard sensors to estimate the temperatures at locations where no sensors are embedded. Our technique is capable of reducing the thermal tracking error by 40%. For power characterization, we proposed a new technique for thermal to power inversion for FPGAs using quadratic program optimization. We have proposed experimental techniques to estimate the different parameters required for power mapping. We used our power mapping methodology to estimate the spatial power distribution of an embedded soft processor during operation.

6. REFERENCES

- [1] R. Cochran, A. Nowroz, and S. Reda, "Post-Silicon Power Characterization Using Thermal Infrared Emissions," in *International Symposium on Low Power Electronics and Design*, 2010, pp. 331–336.
- [2] H. Hamann, A. Weger, J. Lacey, Z. Hu, and P. Bose, "Hotspot-Limited Microprocessors: Direct Temperature and Power Distribution Measurements," *IEEE Journal of Solid-State Circuits*, vol. 42, no. 1, pp. 56–65, 2007.
- [3] B. Lee, K. Chung, B. Koo, and N. Eum, "Thermal Sensor Allocation and Placement for Reconfigurable Systems," *Transactions on Design Automation of Electronic Systems*, vol. 4, no. 41, pp. 50:1–23, 2009.
- [4] F. Li, Y. Lin, L. He, D. Chen, and J. Cong, "Power Modeling and Characteristics of Field Programmable Gate Arrays," *IEEE Transactions on Computer Aided Design of Integrated Circuits*, vol. 24, no. 11, pp. 1712–1724, 2005.
- [5] S. Lopez-Buedo and E. I. Boemo, "Making Visible the Thermal Behaviour of Embedded Microprocessors on FPGAs: A Progress Report," in *International Symposium on Field Programmable Gate Arrays*, 2004, pp. 79–86.
- [6] S. Lopez-Buedo, J. Garrido, and E. Boemo, "Dynamically Inserting, Operating, and Eliminating Thermal Sensors of FPGA-Based Systems," *IEEE Transactions on Components and Packaging Technologies*, vol. 25, no. 4, pp. 561–566, 2002.
- [7] F. J. Mesa-Martinez, M. Brown, J. Nayfach-Battilana, and J. Renau, "Measuring Performance, Power, and Temperature from Real Processors," in *International Symposium on Computer Architecture*, 2007, pp. 1–10.
- [8] R. Mukherjee, S. Mondal, and S. O. Memik, "Thermal Sensor Allocation and Placement for Reconfigurable Systems," in *International Conference on Computer Aided Design*, 2006, pp. 437–442.
- [9] A. N. Nowroz, R. Cochran, and S. Reda, "Thermal Monitoring of Real Processors: Techniques for Sensor Allocation and Full Characterization," in *Design Automation Conference*, 2010, pp. 56–61.
- [10] K. Poon, S. Wilton, and A. Yan, "A Detailed Power Model for Field-Programmable Gate Arrays," *ACM Transactions on Design Automation of Electronic Systems*, vol. 10, no. 2, pp. 279–302, 2005.
- [11] L. Shang, A. Kaviani, and K. Bathala, "Dynamic Power Consumption in Virtex-II FPGA Family," in *International Symposium on Field Programmable Gate Arrays*, 2002, pp. 157–164.
- [12] K. Zick and J. Hayes, "On-line Sensing for Healthier FPGA Systems," in *International Symposium on Field Programmable Gate Arrays*, 2010, pp. 239–248.