Characterizing non-Ideal Impacts of Reconfigurable Hardware Workloads on Ring Oscillator-based Thermometers

Moinuddin A. Sayed
Department of Electrical and Computer Engineering
Iowa State University
Ames, Iowa 50011
Email: moin@iastate.edu

Phillip H. Jones
Department of Electrical and Computer Engineering
Iowa State University
Ames, Iowa 50011
Email: phjones@iastate.edu

***Abstract*—Thermal issues have resulted in growing concerns among industries fabricating various types of devices, such as Chip Multiprocessors (CMP) and reconfigurable hardware devices. Since passive cooling costs have risen considerably and packaging for worst-case is no longer practical, dynamic thermal management techniques are being devised to combat thermal effects. For such techniques to be applied effectively, it is necessary to accurately measure device temperatures at run time. Although several techniques have been proposed to measure the on-chip temperatures of reconfigurable devices, ring-oscillators in many ways are a preferred choice due to their strong linear temperature-dependence and compact design using available spare reconfigurable resources.**

A major problem in using ring-oscillators to measure temperature, however, is their strong dependence on the core voltage of, and current distribution throughout the device under test. One of the reasons for variations in these properties is changes in the workload running on the device. Researchers have seen large shifts in the output frequencies of ring-oscillators due to core voltage swings on reconfigurable devices, and have tried to find alternate ways of measuring temperature that attempt to mitigate these effects. The need, however, is to have a workload-compensated ring oscillator-based thermometer for reconfigurable devices. To obtain this, it is first necessary to characterize the non-ideal effects of workload variations on ring oscillator response. Where non-ideal refers to impacts on ring oscillator oscillation frequency due to phenomena other than the workload's impact on device temperature.

This paper performs such a characterization, in which the effects of workload variation on ring oscillator output frequency is quantified. A complete hardware-software setup is designed to collect temperature and power related data along with ring oscillator response to varying workload configurations. In addition, a potential issue with using the Xilinx System Monitor to measure die temperature at high ranges is also briefly discussed.

***Index Terms*—reconfigurable hardware; FPGA; thermal characterization; ring oscillator; temperature measurement; voltage variation;**

I. Introduction

Power dissipation and rising temperatures are core issues in today's microprocessor and FPGA-based circuits. One of the major reasons for increased power densities in microprocessors is up-scaling in terms of frequency and die-sizes, while supply voltages have not seen proportional downscaling [1]. The number of transistors per chip has been rising according to Moore's law, worsening the power density problem every year. To help keep power dissipation from continuing to be a major bottleneck in the functioning of future processors, die-size growth needs to be controlled to reduce active capacitance and supply voltage downscaling needs to be pushed further [1].

Since heat is a manifestation of power consumption, thermal issues are very closely related to power issues in today's multiprocessor and FPGA circuits. It has become imperative for circuit designers to take thermal issues into consideration while proposing a new design [2], [3]. To help manage rising costs associated with cooling a device, replacing external mechanisms such as thermocouples, thermistors and associated circuitry, with compact and easy-to-use thermal monitoring solutions that can accurately estimate the temperature of the die, and give estimates of intra-die variations in temperature would be ideal. To this end for reconfigurable hardware architectures, ring-oscillators have often been used for estimating the temperature of FPGAs.

Ring-oscillators are compact design elements that can be built using minimal reconfigurable logic resources, and the linear dependence of their oscillation frequency on temperature allows their potential use as thermal sensors [4], [5], [6], [7], [8], [9], [10], [11]. One major issue associated with the use of ring-oscillators, however, is that their oscillation frequency changes drastically with small changes in the core voltage of the FPGA, rendering their application as thermometers cumbersome [5], [6], [7], [9], [11].

One of the reasons for core voltage variations in FPGAs is due to changes in the amount of computation being performed (i.e. changes in the workload). The ***main goal of this work*** is to characterize the non-ideal effects of workload variations on ring oscillator response. Where non-ideal refers to impacts on ring oscillator oscillation frequency due to phenomena other than the workload's impact on device temperature. This characterization has been performed on two Xilinx Virtex-5 FPGAs (the XCV5LX110T and the XCV5LX330). Transient and steady-state temperature, supply current and ring-oscillator frequency have been collected to show how workload variation

impacts ring-oscillator frequency. This characterization serves as a first step towards compensating for the effects of core voltage variations, due to changing workloads, on ring-oscillator frequencies, and perhaps towards the design of auto-calibrated workload compensated ring oscillator-based thermometers for FPGAs.

The remainder of this paper is organized as follows. Section II discusses related work. Section III provides an overview of the problem being addressed and the approach used to characterize this problem. A description of the tests that were conducted is given in Section IV. This is followed by an analysis of the results in Section V. This section also gives a brief discussion of odd behavior observed while using the Xilinx System Monitor [12] to measure die temperature. Section VI presents conclusions and future directions for this work.

II. Related Work

A. Thermal and Power Issues

Since power issues have been a cause of concern in recent years, a number of techniques have been adopted to combat them, and perhaps the simplest is to design a chip's package for the worst-case. The aim is to remove heat from the chip at a faster rate than it is generated. However, packaging for worst case has not only become difficult, but also prohibitively expensive, which forces designers to look at other design-level techniques to control power density. Also, quite often the difference between the average and maximum power consumption of a microprocessor is large, thus the use of dynamic thermal management (DTM) mechanisms can be leveraged. Designing passive cooling mechanisms for maximum power and temperature scenarios is not practical, since cooling costs are rising fast and have already reached $1-$3 per watt [1] with existing processors that consume 100 plus watts.

In addition to expense, there are several other reasons why thermal issues have become a major design challenge. For example, leakage current in transistors is one of the major contributors to on-chip power. Since this leakage current is directly proportional to the temperature of the chip, a positive feedback loop is formed between the leakage and temperature of the chip [4]. Creation of thermal hotspots, due to differences in activity rates within a die, poses challenges by accelerating the failure mechanisms in semiconductor devices [13]. An excellent overview of temperature-related failures due to the failure mechanisms such as Electromigration and Time Dependent Dielectric Breakdown in a 65nm FPGA is given in [8]. [8] also discusses performance degradation due to Negative Bias Temperature Instability, which gives additional motivation for devising DTM techniques that keep temperature and power levels from crossing critical thresholds [14].

B. Measuring Temperature

Implementing DTM often necessitates measuring on-chip temperature, which itself has several challenges. Use of devices like thermocouples and thermistors requires that the associated wiring and hardware be immune to high-frequency signals due to cross-talk on the board, and also require special care from the designer for sensor positioning, coupling and instrumentation [5]. To avoid using such devices, techniques at the design level have been proposed. An embedded diode is fabricated on to the die of modern FPGAs to measure the junction temperature; the latest series of Xilinx FPGAs (Virtex 5 onwards) also provide a way to read the output of this diode in a digitized form [12].

C. Ring Oscillators as Thermal Sensors and Related Issues

Several researchers have made use of ring-oscillators as compact thermal sensors on reconfigurable devices [4], [5], [6], [7], [15]. The problem of ring-oscillator frequency variation due to changes in a device's core voltage has been discussed in [6], [7], [9], [11], which also develop workarounds to the problem by applying different approaches to temperature-measurement. [7] uses a sample mode in which the application is momentarily paused in order to obtain stable measurements; [6] suggests making measurements over a range of temperatures and supply voltages, building an empirical model, then pausing the application to make measurements of the frequency and voltage, and finally plugging these values into a model that estimates temperature; in [9], experiments are conducted that show increased non-linearity of ring-oscillator frequencies at low-voltages and shows that each frequency point has a unique voltage-temperature pair; and [11] suggests using the CMOS delay coefficient and applying it after measuring the oscillation frequency at a known room temperature. Our work differs from other works to date in that it provides a detailed characterization of non-ideal impacts of workload variation on ring oscillator response. This is a first step toward developing a workload compensation scheme in the future.

III. Workload-Variation Impact on Ring Oscillator-Based Thermometer Frequency

A. Overview

Figure 1 depicts the problem that has motivated the characterization performed in this work. While running an application on an FPGA that employs ring oscillators as thermometers, sudden changes in the workload results in significant abrupt shifts in the frequency of the ring oscillator's oscillation. Two postulated reasons for this frequency shift are: 1) a sudden increase in workload causes dips in the core operating voltages, which in turn impacts the ring-oscillator frequency, and 2) a sudden increase in workload stresses the power distribution network of the FPGA, thus providing less current to logic elements, which in turn decreases the ring oscillator's oscillation frequency. Without compensating for workload dependencies, using ring oscillators as thermal sensors is challenging.

To clarify the gravity of the voltage and current variation problem, it should be noted that even very small variations in voltage (on the order of 1-5mV) can cause significant shifts in the output frequency of a ring-oscillator [6], [7],

[9] [11]. Taking a specific example from this work, an instantaneous change in workload from 0% to 80% utilization of a Virtex-5 LX110T FPGA running at 100 MHz, causes the output count value obtained from the ring-oscillator to instantaneously decrease by 295, resulting in a 73°C error in estimated temperature. Thus making it imperative to account for the impact of workload variation.

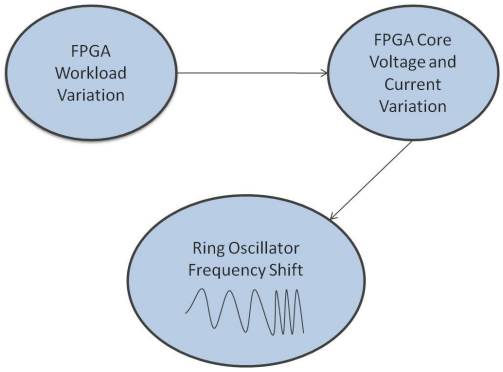

Fig. 1. Ring-oscillator frequency dependence on workload.

B. Architecture for Characterizing Ring Oscillator Dependency on Workload Variation

This work was implemented on two different Virtex-5 FPGAs, an LX110T mounted on a Xilinx XUP-V5 board [16] and an LX330 mounted on a Hitech Global board [17]. For both the chips, the design consists of a ring-oscillator placed in the middle of the chip and a flexible thermal benchmark circuit that occupies the rest of the resources on the chip. This benchmark circuit, made up of Core Blocks, is based on a thermal benchmark architecture described in [18]. Each Core Block is a chain of D-type flip-flops connected to each other through logic gates to form an array.

Figure 2 shows the implementation of a Core Block using FPGA resources. Figure 3 illustrates the formation of a Thermal Workload Unit using these core blocks. A Thermal

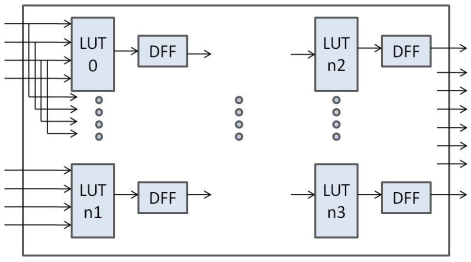

Fig. 2. Core Block.

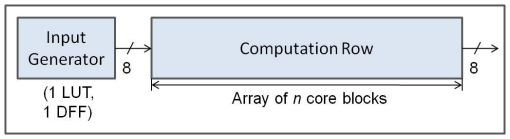

Fig. 3. Thermal Workload Unit.

Resources	LX110T	LX330
LUTs	1152	3312
FFs	1152	3312
% of chip	1.667	1.597

Fig. 4. Size of one workload unit on the LX110T and LX330.

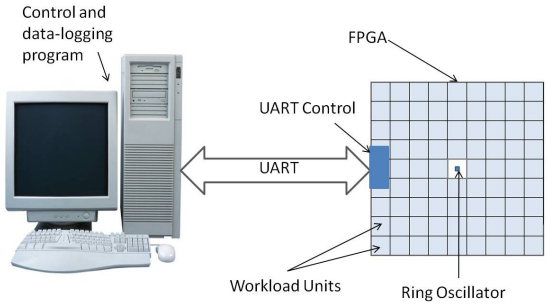

Fig. 5. Hardware-software setup.

Workload Unit consists of a chain of Core Blocks concatenated to form a Computation Row, and an Input Generator that drives the activity rate of the Computation Row. The main purpose of the workload units is to toggle the resources (flip-flops and logic gates) every clock cycle to achieve maximum heating of the chip. Since the entire benchmark circuit is divided into a number of such workload units, they can be selectively enabled to control the amount the FPGA heats. Figure 4 shows the resource utilization of one workload unit on the LX110T and LX330. The LX110T houses 53 and the LX330 houses 54 of these workload units to achieve a maximum utilization of 88% and 86% respectively.

A program running on an external PC was used to selectively enable portions of the chip by activating sets of workload units. In addition, the program continuously logged data during each experiment. The information logged consisted of ring oscillator frequency, current pull from the power supply, and system monitor temperature readings. For activating workload units and logging data, the external program sent simple commands over a UART interface that were acted upon by a small command processing module deployed on the FPGA. This setup is shown in Figure 5, which provides the overall architecture of the measurement system. Tests were run on both chips that enabled 0% though 80% of the available FPGA resources, in steps of 20%, at various frequencies. These workloads are described in more detail in Section IV.

C. Ring-Oscillator (Thermal Monitor)

The architecture of the thermal monitor is organized such that in a fixed period of time the number of ring oscillator oscillations is counted. This is achieved by having the ring oscillator drive the clock, whose frequency is dependent on

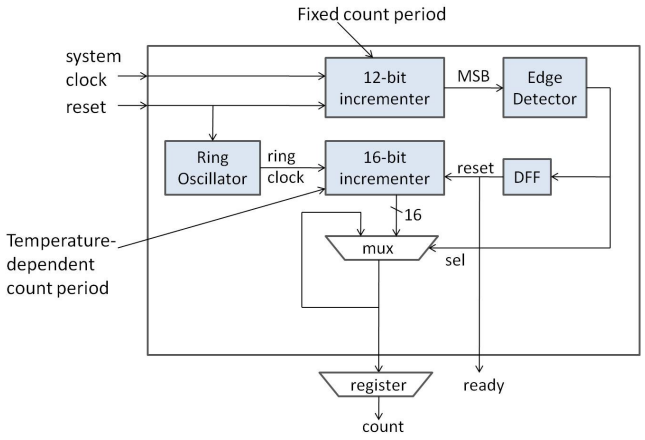

Fig. 6. Thermal Monitor Architecture.

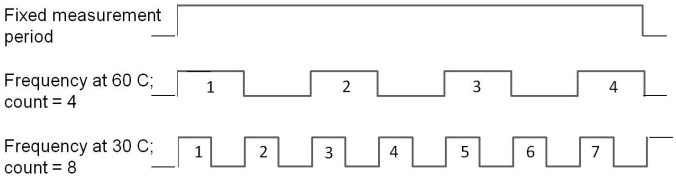

Fig. 7. Ring oscillator frequency temperature dependence.

temperature, of an incrementer circuit. A fixed system clock was then used to measure how many times the temperature dependent incrementer counted over a fixed period of time. Since the ring-oscillator period is a function of temperature, the changes in the count obtained is related to the temperature of the circuit. The thermal monitor architecture is shown in Figure 6. In this implementation, a fixed 33MHz system clock drives a 12-bit incrementer, the most significant bit (MSB) of which is applied to an edge detection circuit. The output from the ring oscillator is applied to a 16-bit incrementer. As soon as an edge is detected on the MSB of the fixed-clock incrementer, a select signal is applied to a multiplexer that places the output of the 16-bit thermally-dependent-clock incrementer on the final output of the circuit, and a ready signal registers this value. The 16-bit incrementer is then reset to 0. As the FPGA die temperature varies, the number of ring oscillator oscillations counted by the 16-bit incrementer in a fixed time period changes due to the effect that temperature has on the period of oscillation.

Figure 7 illustrates the dependence of the ring-oscillator frequency on temperature. As the temperature decreases, the period of oscillation reduces, resulting in an increase in the output frequency, and hence an increase in the incrementer count value. The basic idea of the thermal monitor described above is taken from [7].

Frequency (MHz)	Percentage Of Chip Enabled				
50	0	20	40	60	80
100	0	20	40	60	80
150	0	20	40	60	80
200	0	20	40	60	80

Fig. 8. Test configurations. For each combination, the following data was collected : 1) current from power supply, 2) FPGA case temperatures from thermal probe, 3) ring oscillator count.

IV. EXPERIMENTATION

A. Implementation Details

The design described in Section III was instantiated on two different FPGAs, the XC5V110T residing on a Xilinx XUPV5-LX110T board [16] and the XC5V330 residing on a Hitech Global TB-5V-LX330-DDR2-E board [17]. For the purpose of characterization, tests were run at four different frequencies 50, 100, 150 and 200MHz. For each frequency, the number of workload units was varied from 0% and 80% utilization of the FPGA, in steps of 20%. Figure 8 lists all the frequencies and utilizations that were tested. A thermal probe was used to monitor the temperature of the FPGA case. All data was logged into a file for plotting.

The FPGA case temperatures were collected using an external thermal probe that was kept in contact with the case. This, as opposed to using the Xilinx System Monitor [12], was used for temperature monitoring due to an odd observation with the behavior of the System Monitor at high temperatures. This behavior is described in Section V-A. The difference between the Xilinx Virtex-5 FPGA junction and case temperatures is dictated by the junction-to-case thermal resistance (θ_{jc}) of the chip, which is specified to be 0.10°C/W to 0.15°C/W in [19]. In addition, the maximum power consumed by the LX110T was 5W and by the LX330 was 12W, which translates to a maximum temperature difference between the case and junction of 0.5°C for the LX110T and 1.2°C for the LX330. Thus, relying on a surface mounted thermal probe to measure temperature instead of the System Monitor was deemed reasonable.

B. Test Procedure

Tests were run that collected case temperature, current pulled from the power supply and the ring-oscillator count values. These tests were performed to observe the dependency of ring-oscillator frequency on workload variations for a given temperature. The test procedure for conducting these experiments was as follows:

1) **0% utilization:** The chip was heated to 80°C by enabling the maximum number of workload units and with the assistance of an external heating source. A command was then issued across the UART that disabled all the workload units, and data was collected as the chip cooled to its steady state temperature.

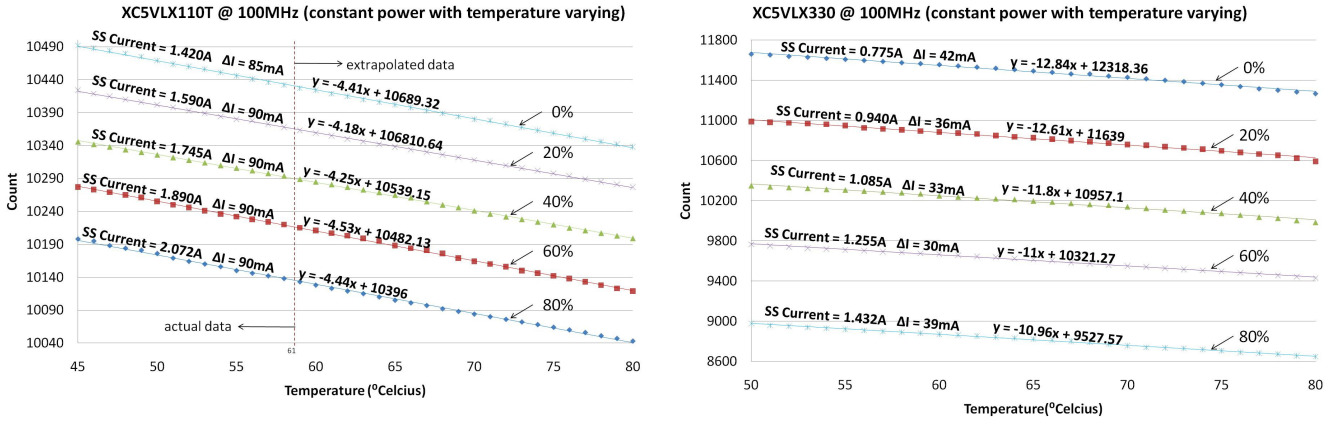

(a) (b)

Fig. 9. Temperature versus ring oscillator count data for utilizations from 0% through 80%. (a) shows temperature vs count values for constant lines of power for LX110T. (b) shows temperature vs count values for constant lines of power for LX330.

2) **80% utilization:** The chip was cooled to a minimum steady state temperature by disabling all workload units. Then a command was issued that enabled 80% of the FPGA resources, and data was collected as the case temperature moved towards 80°C.
3) **20% to 60% utilization:** These tests were run in two phases to collect data over a wide temperature range. Phase 1 is identical to test procedure 1, except instead of going from 80% to 0% FPGA utilization, the utilization was set from 80% utilization to the percentage of chip being tested. Phase 2 is identical to test procedure 2, except instead of going from 0% to 80% utilization, the FPGA utilization was set from 0% utilization to the percentage of chip being tested.

V. RESULTS AND ANALYSIS

Figures 9(a) and 9(b) plot temperature versus ring oscillator count value for a subset of the tests conducted (100MHz only). Figure 9(a) shows the response of the XC5VLX110T FPGA, and Figure 9(b) shows the response for the XC5VLX330.

The graphs show how different workloads affect the relationship between the temperature of the FPGA and the count of the ring oscillator. The SS (Steady-State) Current indicates the current being drawn from the supply used to power the FPGA board after running a particular configuration on the chip for about 15 minutes, and the ΔI value indicates the total change in current from the supply while the workloads are in a particular configuration, for the duration of the test.

The real significance of this data is that between different configurations, different ring oscillator counts are obtained for the same temperature. For the LX110T, a change in current pull of approximately 650mA between the configurations of 0% and 80% utilization causes a change of around 300 in the count obtained from the ring-oscillator. Observing from Figure 9(a) that the precision of count values is 4 counts per degree Celsius, this translates to a discrepancy of about 75°C in the estimated temperature. This emphasizes the fact that the response of the ring-oscillator is shifted by a large amount for a particular change in workload, which necessitates compensating for this dependency. Within a configuration, since the change in current due to temperature is negligible, as indicated by the ΔI value, the dependence of the ring oscillator frequency on temperature can be easily observed to be linear. It should be noted that although the graph for the XC5VLX330 shows measured values for the entire range, the plot for the XC5VLX110T contains extrapolated values for temperatures greater than 60°C.

Alternately, the data collected above can be represented to show how the ring-oscillator count varies with current draw for constant values of temperature. Figures 10(a) and 10(b) show this representation for the two chips. The points on each of the lines correspond to different configurations of the chip, from 0% through 80% in steps of 20%. The count values for both FPGAs show an almost linear variation with changes in current draw. Since each line corresponds to a constant temperature, the dependence of count values on current can be quantified.

The average slopes obtained from Figures 10(a) and 10(b) are 465.64 (Counts/A) for the LX110T and 4010.49 (Counts/A) for the LX330. Given 4 counts/°C, this translates to a temperature error of 1°C per 8.6mA of current change (relative to a baseline workload configuration) for the LX110T, and given 12 counts/°C this translates to a temperature error of 1°C per 3mA of current change for the LX330. Thus changes in workloads executing on an FPGA can have a large impact on a ring oscillator's measurement of temperature.

The sensitivity of ring oscillator-based temperature error due to current change appears to be about 2.7 times greater for the LX330 because the Hitech Global board is powered off of a 12V power supply, while the XUP-V5 board runs on a 5V power supply (a 2.4 factor difference). Ideally current values should be measured directly from the 1V voltage regulator supplying the FPGA's core voltage.

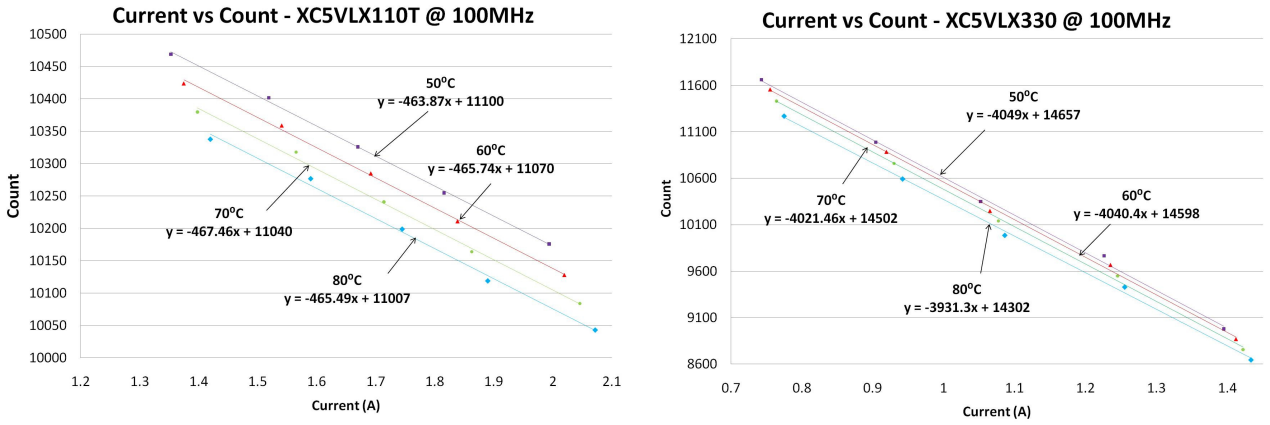

(a) (b)

Fig. 10. Current versus ring oscillator count data for utilizations from 0% through 80%. (a) shows current vs count values for constant lines of temperature for LX110T. (b) shows current vs count values for constant lines of temperature for LX330. Although the lines in (b) appear close together, they are actually much farther apart than those in (a), spanning a difference of around 2700 in the count as opposed to 300 in (a)

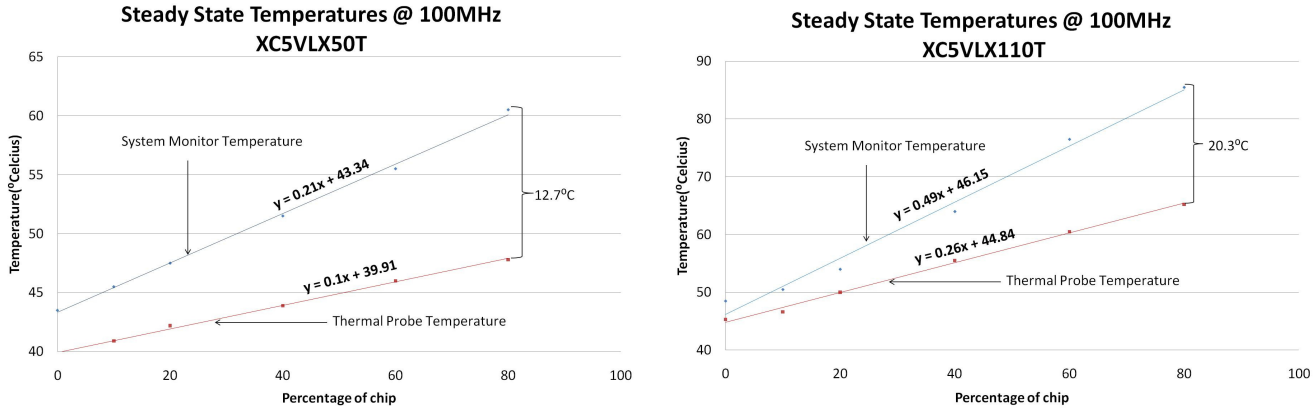

(a) (b)

Fig. 11. Steady-state temperatures for various utilizations.

A. *Unexpected System Monitor Behavior*

Figures 11(a) and 11(b) show the steady state temperatures for different configurations of workloads running at a frequency of 100MHz on two chips, the Xilinx XC5VLX50T and the XC5VLX110T respectively. The two plots show the temperatures as reported by the System Monitor (die temperature) and a temperature probe contacting the top center of the device (case temperature). It was observed that the steady state temperatures as reported by the System Monitor were higher than that obtained from the thermal probe, and this difference increased with the FPGA temperature (by much more than could be accounted for by θ_{jc}).

Examining the plot for the XC5VLX50T (Figure 11(a)) shows when 0% of the chip was enabled that the System Monitor reports the temperature as 43.5°C, while the thermal probe reports 39.7°C (a difference of 3.8°C). This is within the range of error for System Monitor measurements (4°C), as specified in [12]. However, as the temperature of the FPGA increases, the temperature reported by the System Monitor rises much faster than that shown by the thermal probe. At 80% utilization, a steady state temperature of 60.5°C is shown by the System Monitor, while the probe shows the case temperature as 47.8°C, thus increasing the initial difference of 3.8°C to 12.7°C. For the XC5VLX110T FPGA (Figure 11(b)), this difference in even larger at higher temperatures. For this chip, at 0% utilization, the difference between System Monitor and probe temperatures is 3.2°C (System Monitor showing 48.5°C and probe showing 45.3°C) and at 80% utilization is 20.3°C (System Monitor showing 85.5°C and probe showing 65.2°C). The potential reasons for this unexpected behavior are still under investigation.

VI. CONCLUSIONS AND FUTURE WORK

This paper has described a method to collect temperature and power related data for three different Xilinx Virtex-5 FPGAs, for the primary purpose of characterizing the effects of workload-variations on ring oscillator response in FPGAs. A complete hardware-software setup has been developed to log measurement data from an FPGA to an off-chip computer

in real time. Also, unexpected behaviors when using the Xilinx System Monitor at high on-chip temperatures has been discussed.

It is shown that the sensitivity of errors in temperature measurements using a ring oscillator-base thermometer can be as great as 1°C per 3mA change in current drawn induced by changes in the FPGA workload. This strong dependence of the ring oscillator response to workload variation makes apparent the need for compensating for this impact to increase the robustness of ring oscillator-based thermometers. The data obtained from this characterization is a starting point for compensating for such effects. Ideally an auto-calibrated workload variation-compensated ring oscillator-based thermometer for FPGAs is desired.

REFERENCES

[1] S. Borkar, "Design challenges of technology scaling," *Micro, IEEE*, vol. 19, no. 4, pp. 23 –29, jul-aug 1999.

[2] M. Flynn and P. Hung, "Microprocessor design issues: thoughts on the road ahead," *Micro, IEEE*, vol. 25, no. 3, pp. 16 – 31, may-june 2005.

[3] K. Krewell. (2004) Intel cancels 4ghz p4. [Online]. Available: http://www.cs.virginia.edu/ mc2zk/cs451/184401.pdf

[4] S. Velusamy, W. Huang, J. Lach, M. Stan, and K. Skadron, "Monitoring temperature in fpga based socs," in *Computer Design: VLSI in Computers and Processors, 2005. ICCD 2005. Proceedings. 2005 IEEE International Conference on*, oct. 2005, pp. 634 – 637.

[5] S. Lopez-Buedo, J. Garrido, and E. Boemo, "Thermal testing on reconfigurable computers," *Design Test of Computers, IEEE*, vol. 17, no. 1, pp. 84 –91, jan-mar 2000.

[6] K. M. Zick and J. P. Hayes, "On-line sensing for healthier fpga systems," in *Proceedings of the 18th annual ACM/SIGDA international symposium on Field programmable gate arrays*, ser. FPGA '10. New York, NY, USA: ACM, 2010, pp. 239–248. [Online]. Available: http://doi.acm.org/10.1145/1723112.1723153

[7] P. Jones, J. Moscola, Y. Cho, and J. Lockwood, "Adaptive thermoregulation for applications on reconfigurable devices," in *Field Programmable Logic and Applications, 2007. FPL 2007. International Conference on*, aug. 2007, pp. 246 –253.

[8] P. Mangalagiri, S. Bae, R. Krishnan, Y. Xie, and V. Narayanan, "Thermal-aware reliability analysis for platform fpgas," in *Computer-Aided Design, 2008. ICCAD 2008. IEEE/ACM International Conference on*, nov. 2008, pp. 722 –727.

[9] J. Franco, E. Boemo, E. Castillo, and L. Parrilla, "Ring oscillators as thermal sensors in fpgas: Experiments in low voltage," in *Programmable Logic Conference (SPL), 2010 VI Southern*, march 2010, pp. 133 –137.

[10] S. Lopez-Buedo, J. Garrido, and E. Boemo, "Dynamically inserting, operating, and eliminating thermal sensors of fpga-based systems," *Components and Packaging Technologies, IEEE Transactions on*, vol. 25, no. 4, pp. 561 – 566, dec 2002.

[11] E. I. Boemo and S. López-Buedo, "Thermal monitoring on fpgas using ring-oscillators," in *Proceedings of the 7th International Workshop on Field-Programmable Logic and Applications*. London, UK: Springer-Verlag, 1997, pp. 69–78. [Online]. Available: http://portal.acm.org/citation.cfm?id=647924.738726

[12] *ug192 Virtex-5 FPGA System Monitor User Guide*, Xilinx Inc.

[13] A. Coskun, T. Rosing, K. Whisnant, and K. Gross, "Static and dynamic temperature-aware scheduling for multiprocessor socs," *Very Large Scale Integration (VLSI) Systems, IEEE Transactions on*, vol. 16, no. 9, pp. 1127 –1140, sept. 2008.

[14] D. Brooks and M. Martonosi, "Dynamic thermal management for high-performance microprocessors," in *High-Performance Computer Architecture, 2001. HPCA. The Seventh International Symposium on*, 2001, pp. 171 –182.

[15] S. Lopez-Buedo, J. Garrido, and E. Boemo, "Thermal testing on programmable logic devices," in *Circuits and Systems, 1998. ISCAS '98. Proceedings of the 1998 IEEE International Symposium on*, vol. 2, may-3 jun 1998, pp. 240 –243 vol.2.

[16] *XUPV5-LX110T User Manual*, Xilinx Inc.

[17] *Virtex 5 LX330 DDR2 II Memory image processing ASIC prototyping Board User Manual*, Hitech Global.

[18] P. H. Jones, J. W. Lockwood, and Y. H. Cho, "A thermal management and profiling method for reconfigurable hardware applications," in *Field Programmable Logic and Applications, 2006. FPL '06. International Conference on*, aug. 2006, pp. 1 –7.

[19] *ug195 Virtex-5 FPGA Packaging and Pinout Specification*, Xilinx Inc.