

2012

Mitigating impacts of workload variation on ring oscillator-based thermometers.

Moinuddin Abrar Sayed
Iowa State University

Follow this and additional works at: <https://lib.dr.iastate.edu/etd>

 Part of the [Computer Engineering Commons](#)

Recommended Citation

Sayed, Moinuddin Abrar, "Mitigating impacts of workload variation on ring oscillator-based thermometers." (2012). *Graduate Theses and Dissertations*. 12860.

<https://lib.dr.iastate.edu/etd/12860>

This Thesis is brought to you for free and open access by the Iowa State University Capstones, Theses and Dissertations at Iowa State University Digital Repository. It has been accepted for inclusion in Graduate Theses and Dissertations by an authorized administrator of Iowa State University Digital Repository. For more information, please contact digirep@iastate.edu.

Mitigating impacts of workload variation on ring oscillator-based thermometers

by

Moinuddin A. Sayed

A thesis submitted to the graduate faculty
in partial fulfillment of the requirements for the degree of
MASTER OF SCIENCE

Major: Computer Engineering

Program of Study Committee:

Phillip H. Jones, Major Professor

Joseph Zambreno

Zhao Zhang

Iowa State University

Ames, Iowa

2012

Copyright © Moinuddin A. Sayed, 2012. All rights reserved.

TABLE OF CONTENTS

LIST OF TABLES	iv
LIST OF FIGURES	v
ACKNOWLEDGEMENTS	viii
ABSTRACT	ix
CHAPTER 1. Introduction	1
CHAPTER 2. Background	4
2.1 Temperature Measurement Techniques	4
2.2 Ring Oscillator-based Thermometers	6
2.2.1 Theory of Operation	7
2.2.2 Benefits of Deployment in FPGAs	8
2.2.3 Summary of Research use in FPGAs	9
CHAPTER 3. Related Work	12
3.1 Ring Oscillator-based Thermometer Concerns	12
3.2 Techniques for Mitigating Ring Oscillator-based Thermometer Concerns	13
CHAPTER 4. Ring Oscillator-based Thermometer Characterization	15
4.1 Architecture	15
4.1.1 Thermal Monitor	15
4.1.2 Thermal Benchmark Circuit	18
4.1.3 External UART/Command Interface	20
4.1.4 Current Measurement	21
4.2 Experimental Setup and Methodology	22

4.3 Results and Analysis	25
CHAPTER 5. Mitigating Impacts of Workload-variation on Ring Oscillator- based Thermometer Behavior	30
5.1 Approach	30
5.2 Results and Analysis	33
CHAPTER 6. Unexpected System Monitor Behavior	36
CHAPTER 7. Conclusions and Future Work	39
BIBLIOGRAPHY	41

LIST OF TABLES

4.1	Size of a single workload unit on the LX110T and LX330	19
4.2	Test Configurations. For each configuration, the following data was collected : 1) current from power supply, 2) FPGA case temperatures from thermal probe, 3) ring oscillator count.	24
5.1	Compensation test table.	34

LIST OF FIGURES

2.1	A ring oscillator with odd number of inverters in a loop. The switch is used to connect/disconnect the last inverter from the first to begin self-sustained oscillations. The output oscillation is observed on the probe point shown after the first inverter.	7
2.2	Ring oscillator period of oscillation. The output of a ring oscillator toggles after delay times number of inverters in the loop, making the period of oscillation twice the sum of delays of all the inverters in the loop.	8
3.1	Ring oscillator frequency dependence on workload. Blocks in red indicate active parts on the FPGA.	13
4.1	Thermal Monitor Architecture with ring oscillator and associated counter circuitry.	16
4.2	Ring oscillator frequency temperature dependence example. The thermally dependent clock's frequency changes with temperature, giving different counts at the output of the thermal monitor during a fixed measurement period.	17
4.3	Ring Oscillator with 23 inverters in a loop. The OR gate is used to initialize oscillations. Output is obtained at the Ring Clock probe point.	17
4.4	Core Block consisting of a chain of D-type Flip Flops connected together. The LUTs between the flip flops in the figure are configured as AND gates.	18

4.5	Thermal Workload Unit consisting of an array of core blocks. The input generator is made up of a NOT gate and a D flip flop and is used to excite the workload unit.	18
4.6	Workload placement on the FPGA. The ring oscillator-based thermometer is placed at the center of the die, surrounded by thermal workload units occupying the entire chip.	19
4.7	Hardware-software setup. The external PC is used to send commands to and receive data from the FPGA over a UART interface.	20
4.8	Difference Amplifier. The voltage across the sense resistor is amplified and provided as input to the System Monitor on the Virtex-5 for conversion.	21
4.9	Complete Instrumentation Sytem.	23
4.10	Actual hardware setup.	23
4.11	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.	27
4.12	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 maximum difference of around 2700 in the count as compared to 300 in (a)	29

5.1	Workload-variation compensation. The plot on the left shows shifts in the counts received from the thermal monitor due to changes in workload. The compensation technique in the middle makes use of the ring oscillator frequency's linear current dependence obtained in Figures 4.12(a) and 4.12(b) to correct the shifts. The plot on the right shows the expected response of the thermal monitor after the compensation technique is applied, with all workload configurations yielding the same response as the baseline configuration.	31
5.2	Compensation applied on Figure 4.11(a). A single response line is obtained after applying the compensation technique, giving a unique value of count at a given temperature in every mode.	32
5.3	Comparison of compensated and uncompensated temperature estimates on LX110T.	33
5.4	Error in temperature measurement between a) uncompensated estimates and the thermal probe, and b) compensated estimates and the thermal probe, for various workload configurations.	35
6.1	Steady-state temperatures for various utilizations. The figures show an increased temperature difference, as the chip temperature rises, between the System Monitor reported values and values obtained using a surface mounted thermal probe.	37

ACKNOWLEDGEMENTS

It gives me immense pleasure to thank all the people directly and indirectly responsible for making this thesis a reality. Without a doubt, the person I feel most grateful towards is my major professor, Dr. Phillip H. Jones. Firstly, I wish to thank him for believing in me, and giving me a chance to work as a research assistant under him. This thesis would not have been what it is today without his constant support and guidance throughout its course. His “students-first” attitude, along with a constant emphasis on generating ideas from them through discussion and fueling them as needed, has had huge impacts on every aspect of this work. I also wish to thank him for finding time from his very busy schedule to perfect the content in this thesis, and transform it into a piece of work I will cherish for the rest of my life. Lastly, I honestly feel I cannot express enough gratitude towards the patience he has shown while working with me, right from the inception to the culmination of this thesis.

I would like to thank Dr. Joseph Zambreno for his excellent (and often witty) teaching style in the Embedded Systems course, and for letting me borrow his FPGA boards for an extended period of time; and Dr. Zhao Zhang for his role as an instructor and later as a guide while I was a teaching assistant for one of his courses. I also wish to thank them for graciously agreeing to serve on my PoS committee. A special note of thanks is in order to a very close friend and colleague at Iowa State University, Pooja Mhapsekar, for working with me on laying the foundations for this work.

Although I’m aware words cannot do justice here, I thank my family (because of whom I am where I am today) and my friends (in Ames, IA and Boulder, CO) for contributing in their own ways towards this thesis, and for making my student life enjoyable and rewarding. Finally, I’d like to thank the Indian Students Association at ISU for providing a platform where I could showcase my extra-curricular skills and give back to the ISU community, and the ISU Cyclone community in general for creating an atmosphere conducive to one’s overall growth.

ABSTRACT

Thermal issues have resulted in growing concerns among industries fabricating semiconductor devices such as Chip Multiprocessors (CMP) and reconfigurable hardware devices. To reduce passive cooling costs and eliminate the need to package for worst-case temperatures, dynamic thermal management (DTM) techniques are being devised to combat thermal effects. Reliable runtime measurement of device temperature is necessary for implementing DTM techniques. Ring oscillators have often been used for on-chip Field Programmable Gate Array (FPGA) temperature measurement 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 that their frequency of oscillation is affected by changes in device core voltage and current distribution, induced by changes in application workload. The need, then, is to have a workload-compensated ring oscillator-based thermometer for reconfigurable devices.

This work performs a characterization of the ideal as well as non-ideal effects of workload variation on ring oscillator frequency response. Where non-ideal refers to impacts on ring oscillator oscillation frequency due to phenomena other than the workload's impact on device temperature. The data obtained from this characterization is used to compensate for these non-ideal effects. A complete hardware-software solution is implemented to collect temperature and power related data along with ring oscillator frequency response to varying workload configurations. The characterization results show an error of approximately 1°C in the estimated temperature for every 8.6mA change in current drawn from the supply on a Xilinx Virtex-5 LX110T FPGA, with respect to the current draw measured while running a baseline workload during thermometer callibration. This lead to a maximum error of $\sim 74^{\circ}\text{C}$ for the workloads evaluated. The compensation technique implemented is shown to reduce this error to $\sim 2^{\circ}\text{C}$.

In addition, a potential issue with using the Xilinx System Monitor to measure die temperature at high temperatures is observed. The System Monitor reported temperatures show a deviation of up to 20°C from temperatures obtained using a case-mounted thermal probe.

CHAPTER 1. Introduction

This chapter provides an overview of issues associated with rising power and temperatures in microprocessors and Field Programmable Gate Arrays (FPGAs). Reasons for the continuous increase in these two inter-related issues are discussed and sources of each outlined. The need for temperature measurement in microprocessors and FPGAs is established, and the use of ring oscillators as thermal sensing elements in such chips is discussed. Then, a concern when using ring oscillators in FPGA-based circuits is described, which forms the motivation for this thesis.

Reasons for rising temperature and power. One of the major reasons for increased power densities in microprocessors is a continuous increase in terms of frequency [ISSCC Trends Report (2011); Rusu (2004)] and die-sizes [Yung et al. (2002); AMD Bulldozer (2011); Intel Processors Online (2012)], while supply voltages have not seen proportional downscaling [Borkar (1999); Huang et al. (2011)]. The number of transistors per chip has been rising according to Moore's law [Moore's Law Inspires Intel Innovation (2011); Bohr (2002); Intel Press Release (2004); Ghani (2009); Intel 22nm Technology (2011)], worsening the power density problem every year.

Sources of power dissipation. The power dissipation of a chip can be categorized as two types, static power and dynamic power. Dynamic power, which is the power consumption due to the switching activity of transistors, has been the dominating component of processor power consumption. However, static power in today's chips cannot be ignored, especially with chips below 100nm, where static power can reach up to 50% of the processor's total power [Kim et al. (2003); Hasan and Bird (2011); Chakraborty and Pradhan (2012)], and continues to be a major design challenge beyond 22nm [ITRS (2009)]. The main source of static power is the leakage current of transistors. Since this leakage current is directly proportional to the temperature of the chip, a positive feedback loop is formed between the two resulting in thermal runaway and

consequent circuit failure [Velusamy et al. (2005); Heo et al. (2003)].

Need for on-chip temperature measurement. Thermal and power issues have become a major bottleneck in the development of next generation semiconductor devices [ITRS (2009)]. This has necessitated devising techniques to keep chip temperature and power dissipation from crossing safe boundaries. One approach is to design a chip's package to tolerate worst case temperatures. However, this method has become increasingly complex and prohibitively expensive over the years [Borkar (1999); Jiang et al. (2008)]. Therefore, instead of designing thermal packages to tolerate maximum power dissipation, they can alternately be designed to handle the *typical* power dissipation, and dynamic thermal management (DTM) techniques can be used to handle scenarios where temperatures exceed predefined thresholds [Brooks and Martonosi (2001)]. To implement DTM, however, it is necessary to accurately measure on-chip temperatures, and several techniques have been proposed to that end.

Ring oscillators as thermal sensors. Ring oscillators are one mechanism for measuring on-chip temperatures of reconfigurable architectures. They are compact design elements that can be implemented using very few reconfigurable logic resources. As a thermal sensor, the linear dependence of their oscillation frequency on temperature is a desirable property [Velusamy et al. (2005); Lopez-Buedo et al. (2000); Zick and Hayes (2010); Jones et al. (2007); Mangalagiri et al. (2008); Franco et al. (2010); Lopez-Buedo et al. (2002); Boemo and López-Buedo (1997)]. However, the oscillation frequency of ring oscillators is sensitive to small changes in a device's core voltage, rendering their application as thermometers cumbersome [Zick and Hayes (2010); Jones et al. (2007)]. One of the reasons for core voltage variations in FPGAs is that supply voltage varies with changes in workload activity [Jones et al. (2007)].

Contributions. The primary contributions of this work are:

- 1) A quantification of ring oscillator based-thermometer dependence on non-ideal effects of workload variation,
- 2) A technique for mitigating ring oscillator-based thermometer workload variation dependence to obtain accurate device temperature estimates,
- 3) A temperature measurement method using ring oscillators that does not require the application to be paused for making stable measurements, unlike previous approaches (described

in Section 3.2), and

- 4) A hardware-software solution to collect temperature- and current-related data from two Xilinx Virtex-5 FPGAs.

In this work, non-ideal refers to impacts on ring oscillator oscillation frequency due to phenomena other than the workload's impact on device temperature, such as variations in the FPGA's core voltage.

Organization. The remainder of this thesis is organized as follows. Chapter 2 gives background on temperature measurement techniques and focuses specifically on using ring oscillators as thermal sensors in FPGAs. Chapter 3 gives an overview of problems associated with using ring oscillators as thermal sensors and discusses work done to date for overcoming these issues. Chapter 4 presents architectural details of the circuits developed for characterizing ring oscillator workload variation dependence. Chapter 4 also describes the evaluation methodology followed by an analysis of the characterization data. The compensation technique used to mitigate workload variation effects on ring oscillator-based thermometer is presented in Chapter 5, along with a discussion of post-compensation results. Chapter 6 describes an issue observed with using the Xilinx System Monitor to obtain temperature measurements at high temperature ranges. Chapter 7 concludes this thesis and suggests areas for future exploration.

CHAPTER 2. Background

This chapter describes temperature measurement techniques available for processors and FPGAs. It also provides an overview of ring oscillators, and how they can be used as thermal sensors in FPGAs.

2.1 Temperature Measurement Techniques

The purpose of this section is to provide a survey of the various techniques available for measuring the temperature of FPGAs and other silicon devices. For the purpose of this thesis, techniques have been categorized into two types: General temperature measurement techniques and ring oscillator-based temperature measurement techniques. This section covers non-ring oscillator-based techniques.

Transducers. Traditionally, thermal transducers such as thermocouples and thermistors have been used to measure temperature of VLSI-based circuits. However, such transducers have a number of disadvantages for measuring temperature. Proper care needs to be exercised when positioning, coupling, and instrumenting such sensors. In addition, measurements using such transducers need to be immune to on-board high frequency signals [Lopez-Buedo et al. (2000)]. Due to these disadvantages, other techniques, including on-chip sensor implementations have been proposed, which are described in the remainder of this section.

CMOS Sensors. Temperature sensing using the voltage and temperature dependence of bipolar and CMOS transistors has been explored since the inception of research in the area of temperature measurement. Although bipolar transistors have greater temperature sensitivity and stability than their CMOS counterparts [Bakker and Huijsing (1996)], they are

not well suited for such solutions in the CMOS-based processes [Szekely et al. (1997)]. Hence, MOS transistors have often been employed as temperature sensors due to their temperature dependence in the weak inversion region. Szekely et al. (1997) describes how this dependence can be exploited to use CMOS technology based transistors as on-chip thermal monitors for VLSI chips. The advantages of these sensors are very low space requirements and low power consumption. The disadvantage, again, is that they exhibit relatively weaker temperature dependence than bipolar transistors.

IR Thermometers. Another method for measuring the temperature of electronic devices is using an infra-red thermometer pointed at a chip's surface. A contactless method to measure temperature with an infra-red thermometer is described in [Becker et al. (2009)], in which the authors characterize the power consumption and temperature of the fine-grain fabric of four different Xilinx FPGAs. No physical contact minimizes the influence of the measurement on the system. Since the effectiveness of such sensors reduces on shiny surfaces, a thin blackened aluminium coating is applied to chips with a metallic surface.

Software Simulators. Thermal simulators are another way of estimating temperature. They are especially useful in measuring temperatures at different granularities of the chip. One of the most widely used simulators, called HotSpot, is described in [Skadron et al. (2003)]. The HotSpot temperature model is based on an equivalent circuit of thermal resistances and capacitances corresponding to microarchitectural blocks present on a chip. Investigations in this work reveal that hotspots typically occur at the granularity of architecture-level blocks present in a chip. Two models are derived for simulating temperatures, a vertical model of heat flow from die to sink and sink to air, and a horizontal model that characterizes heat transfer between microarchitectural blocks. Since validating such a model is difficult due to lack of localized heat measuring techniques, a commercial finite element analysis simulator was used for comparisons, and errors of less than 5.8% between the two models were reported.

Embedded Diodes. An alternative to CMOS sensor-based approaches is to calculate the die temperature by measuring the junction forward voltage of the clamping diode located in

the FPGA pads. Some generations of Xilinx FPGAs provided access to such a diode through external pins, whose output can be connected to an external signal conditioning device to obtain die temperatures. An example of using such a diode for temperature measurements is presented in [Jones et al. (2006)], where the device is reprogrammed when the temperature measured using the output of this diode crosses a threshold. In their work, they create thermal benchmark circuits and monitor the temperature of an FPGA to perform transient and steady-state thermal analysis of the benchmark response. The data is then used to propose a dynamic thermal management technique based on temperature feedback.

System Monitor. More recent versions of Xilinx FPGAs (Virtex-5 onwards) come with the “System Monitor” [System Monitor (2011)], which is a measurement system based on an on-chip analog to digital converter (ADC) and a number of sensors throughout the chip. For temperature measurements, it consists of a built-in temperature sensing diode that produces a voltage output proportional to the die temperature. The ADC is used to convert this to a 10-bit digitized value that can be read by the FPGA logic.

Many works that target measurements or mitigation of on-chip temperatures are based on either directly using the System Monitor, or using it to validate the data obtained through another technique. However, the validity of the measurements obtained by the System Monitor itself remains unquestioned. This could be a problem, since the work in this thesis reports a potential issue with using the System Monitor to obtain die temperatures, especially in the higher temperature ranges. This behavior is discussed with experimental data in Chapter 6.

2.2 Ring Oscillator-based Thermometers

Since the work in this thesis addresses overcoming ring oscillator-based thermometer workload variation dependence, this section is dedicated to providing background on ring oscillator-based thermometers. It first gives a general overview of ring oscillators and fundamentals of oscillation. This is followed by a discussion of why ring oscillators are suitable for use as thermal sensors in reconfigurable architectures. The section concludes with a summary of research performed to date in deploying ring oscillators for temperature measurement in FPGAs.

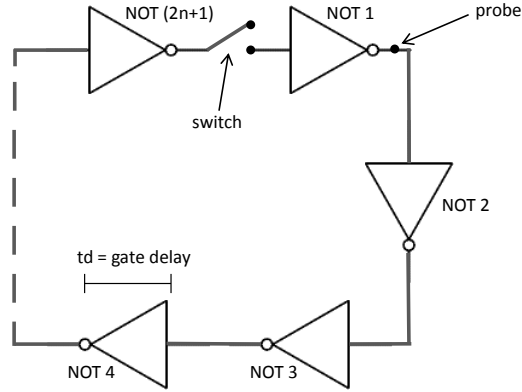


Figure 2.1: A ring oscillator with odd number of inverters in a loop. The switch is used to connect/disconnect the last inverter from the first to begin self-sustained oscillations. The output oscillation is observed on the probe point shown after the first inverter.

2.2.1 Theory of Operation

A ring oscillator is formed by connecting an odd number of inverters to form a loop, as shown in Figure 2.1. As seen in Figure 2.2, a transition in the oscillation occurs after a time period equal to one gate delay times the number of inverters in the loop. The period of oscillation, therefore, is twice the sum of delays of all the inverters in the loop, and is given by Equation 2.1 as,

$$T = 2 * t_d * (2n + 1) \quad (2.1)$$

where T is the period of oscillation, t_d is the delay of one inverter, and $2n+1$ is the number of inverters in the loop.

To begin oscillating, first the last inverter of the loop is disconnected from the first. Next, a constant input is applied to the first inverter. After the input is allowed to propagate through all the inverters in the loop, the output of the last inverter is reconnected to the input of the first. Now, the oscillation is self-sustained.

The frequency of oscillation of a ring oscillator is proportional to changes in temperature, since the switching speed of the transistors that make up each inverter is directly influenced by temperature. Equation 2.2 gives the relation between the carrier mobility in a MOSFET and temperature [Filanovsky and Allam (2001); Ku and Ismail (2007)]. This dependence is close to linear, making the output frequency of ring oscillators easily calibrated to give direct estimates




Figure 2.2: Ring oscillator period of oscillation. The output of a ring oscillator toggles after delay times number of inverters in the loop, making the period of oscillation twice the sum of delays of all the inverters in the loop.

of chip temperature.

$$\mu = \mu_0(T/T_0)^{\alpha_\mu} \quad (2.2)$$

where T is the temperature, T_0 is the nominal temperature, μ_0 is the mobility at T_0 , and α_μ is an empirical parameter referred to as the mobility temperature exponent, with a typical value between -1 and -1.5.

2.2.2 Benefits of Deployment in FPGAs

Some of the advantages of using ring oscillators as temperature sensors are:

- *Junction Temperature Measurement.* Since they are implemented on the FPGA fabric, ring oscillators can be used to measure junction temperatures as opposed to external temperature measurement devices like thermocouples, which can only measure package temperatures.
- *Compact Size.* A ring oscillator can be realized on an FPGA by configuring the available Lookup Tables (LUTs) to implement the structure of Figure 2.1. The size of a ring oscillator is small and takes up only a small fraction of the available chip resources. In terms of number of LUTs, the size of a ring oscillator can be as small as 8 [Zick and Hayes (2010)].
- *Digital Output.* Unlike other temperature transducers (such as a thermal diode), the output from a ring oscillator is fully digital, eliminating the need for an analog to digital

converter after sensing.

- *Finer Granularity.* Localized temperature readings can be obtained with the use of ring oscillators by placing them at various locations of the chip. This can be helpful in implementing strategies that target minimizing hotspots on a chip.
- *Reconfiguration Capability.* Through dynamic reconfiguration, the sensors and associated circuits can be removed from the FPGA when not needed and reinserted as required.

2.2.3 Summary of Research use in FPGAs

Temperature sensing. Perhaps the first notable use of ring oscillators as temperature sensors on programmable logic devices was done by Quenot et al. (1991). They developed a system with a programmable ring oscillator, an 18-bit counter and a control circuit. In addition to temperature, they also conducted experiments to show the dependence of ring oscillator frequencies on supply voltage. The measurements were made in a stable ambient temperature room and the circuit was heated using a built-in resistor. For temperature measurements, they reported an accuracy of around 3°C.

Lopez-Buedo et al. (1998) used Xilinx chips for designing and testing ring oscillator-based thermal sensors. Their basic idea was to count the oscillation of a ring oscillator by making periodic measurements using a control circuit. Along with temperature-dependent frequency data, they also collected data to show the impact of self-heating, the dependence of oscillation frequencies on power supply variations, and the usefulness of other on-chip resources for measuring temperature, such as the IOB clamping diodes. The thermal sensors were calibrated with the help of a temperature-controlled oven and an Iron-Constantan thermocouple. A microcontroller-based mechanism was used to measure the output frequency of the ring oscillator. Their results showed that the sensors exhibited a linear temperature and power supply dependence in the devices' normal temperature and voltage operating range.

To further their work with ring oscillator-based temperature sensors, Lopez-Buedo et al. (2002) proposed a new strategy to build a fully digital temperature transducer based on ring oscillators that could be dynamically inserted, operated and then eliminated from the circuit

using run time reconfiguration techniques. They first inspect the bitstream of the application running on the FPGA to find potential spaces where the ring oscillators and related temperature measurement and control circuits can be placed. The FPGA was then partially reconfigured to insert the thermal sensors. Once measurements were complete, the FPGA was partially reconfigured to remove the thermal sensors.

Another use of ring oscillators as thermal sensors can be seen in [Velusamy et al. (2005)], in which the authors distributed sensors throughout an FPGA and interfaced to a controller for measurements. For cross-validation, they compared their readings with an architecture-level thermal modeling tool named HotSpot [Skadron et al. (2003)]. The temperature sensor design was identical to the one described in [Lopez-Buedo et al. (1998)]. In comparisons with HotSpot, they reported temperature differences of less than 0.2°C .

A ring oscillator was used in [Jones et al. (2007)] to design a thermal monitor to make a temperature self-regulated system. Since their work also deals with overcoming an issue with using ring oscillators as thermal sensors, it is discussed in detail in the next chapter, along with other related work.

Thermal validation. Mangalagiri et al. (2008) outlined the effects of temperature variations on long term reliability of FPGAs. They gave an overview of temperature-related failures due to failure mechanisms such as Electromigration and Time Dependent Dielectric Breakdown in 65nm FPGAs. They also studied performance degradation due to Negative Bias Temperature Instability. In their work, they developed a thermal predictive model using HotSpot, and used a ring oscillator-based thermometer to validate the results of their simulation tool.

Run-time temperature prediction. In [Happe et al. (2011)], the authors designed a system with a self-calibrating net of thermal sensors based on ring oscillators distributed over an FPGA. The sensors were calibrated with the help of an in-built thermal diode by activating heat-generating circuits throughout the chip. After calibration, spatial and temporal gradients were generated on the FPGA by activating selected heat-generating circuits. A learning algorithm was then used to determine the parameters of the system's internal thermal model, enabling the system to measure and predict temperature distributions autonomously. The temperature predictions obtained from the model were shown to have an average error of 0.72°C as compared

to run-time temperature measurements.

Hotspot management. Zhang et al. (2010) used ring oscillators and associated counter circuitry to predict heat emissions locally and define a temperature-power relationship in FPGA-based systems. Their aim was to use this relationship to restrict the power consumption in systems such as Cognitive Radio equipment.

In [Franco et al. (2010)], experiments at low core voltages, eg. below 1.0V, were conducted to collect sensor responses to varying voltages and temperatures. Their results showed increased non-linearity of ring oscillator frequencies at low voltages.

The works described in this section presented various uses for ring oscillator-based thermal sensors. However, there are concerns that must be addressed when using ring oscillator-based thermometers. The next chapter discusses these concerns, and presents previously used solutions.

CHAPTER 3. Related Work

This chapter discusses challenges associated with using ring oscillators as thermal sensors in FPGAs, and summarizes the techniques developed to date to overcome these challenges. The work in this thesis is then placed in context with this body of work.

3.1 Ring Oscillator-based Thermometer Concerns

A major issue with using ring oscillators as temperature sensors is that their frequency is not dependent on temperature alone, but also on supply voltage. This section describes this problem and the next section presents work that has been done to date to overcome the voltage-induced effects on ring oscillator frequency.

While running an application on an FPGA that employs ring oscillators as thermometers, changes in workload results in significant abrupt shifts in the frequency of the ring oscillator. This is illustrated in Figure 3.1, which forms the motivation for this work. Two postulated reasons for this frequency shift are:

- 1) an increase in workload causes dips in the core operating voltages, which in turn reduces the ring oscillator frequency, and
- 2) an 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 frequency.

Small variations in voltage (on the order of 1-5mV) can cause significant shifts in the output frequency of a ring oscillator [Zick and Hayes (2010); Jones et al. (2007); Franco et al. (2010); Boemo and López-Buedo (1997)]. Taking a specific example from this work, a change in workload from 0% to 80% utilization of a Virtex-5 LX110T FPGA running at 100MHz, causes the output count value obtained from a ring oscillator-based thermal monitor circuit

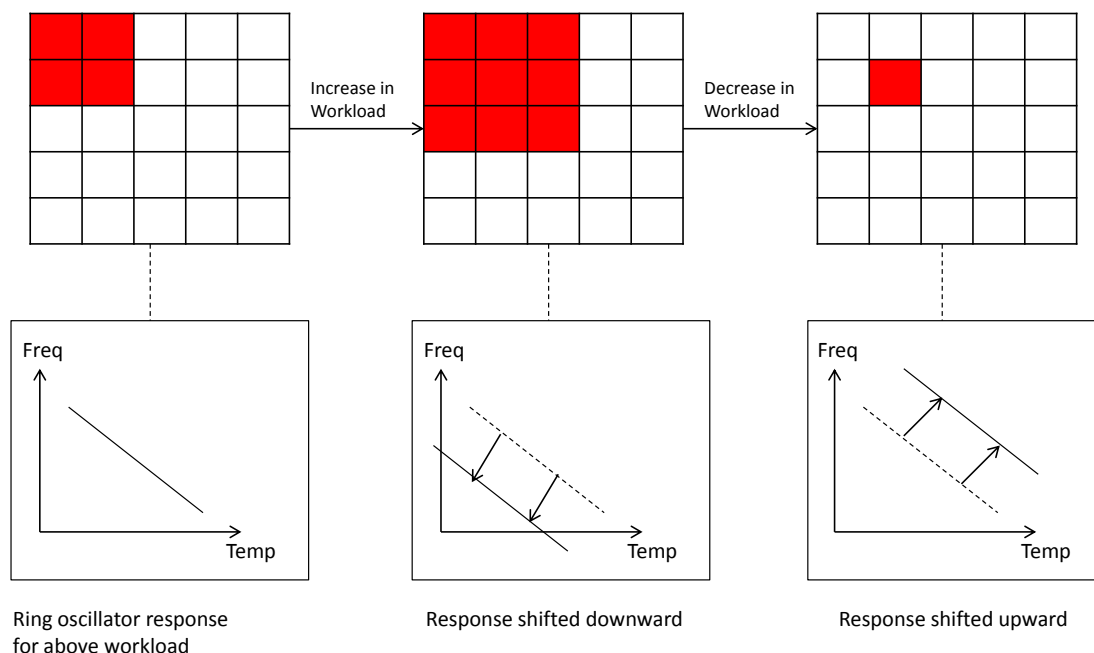


Figure 3.1: Ring oscillator frequency dependence on workload. Blocks in red indicate active parts on the FPGA.

to instantaneously decrease by 295, resulting in a 74°C error in estimated temperature. Thus making it imperative to account for the impact of workload variation.

3.2 Techniques for Mitigating Ring Oscillator-based Thermometer Concerns

Several researchers have made use of ring oscillators as compact thermal sensors on reconfigurable devices [Velusamy et al. (2005); Lopez-Buedo et al. (2000); Zick and Hayes (2010); Jones et al. (2007); Lopez-Buedo et al. (1998)]. The problem of ring oscillator frequency variation due to changes in a device's core voltage has been discussed in [Zick and Hayes (2010); Jones et al. (2007)]. These works in addition develop different approaches to work around this issue.

In [Jones et al. (2007)], the authors switch to a temperature measurement mode to obtain stable temperature measurements. While in the temperature measurement mode, the application is paused, thus reducing the effects of supply voltage variations induced by workload variations. An event counter is used to signal the beginning of the temperature measurement mode, and a sample mode controller is used to pause the application while temperature is being

measured. These temperature measurements are then used to regulate an image recognition application implemented on a Xilinx Virtex-4 FPGA.

Zick and Hayes (2010) use a ring oscillator-based sensor built using only 8 LUTs that is deployed onto a Xilinx Virtex-5 FPGA. To remove the effects of voltage variations, measurements over a range of temperatures and supply voltages are made for the purpose of building an empirical model. Then, the application is paused to make measurements of the frequency and voltage, and finally the frequency and voltage values are plugged into a model that estimates temperature.

A major problem with the workarounds presented in this section is that they require the application be paused to obtain stable temperature measurements. This often leads to negative impacts on application performance. Furthermore, in certain scenarios, it is not feasible to pause the application. The work in this thesis differs from other works to date in that it does not require the application to be paused to obtain accurate temperature estimates. The system can continue normal operation while the mitigation technique proposed in this work is implemented for obtaining reliable temperature measurement using ring oscillator-based thermometers.

CHAPTER 4. Ring Oscillator-based Thermometer Characterization

This chapter first describes the ring oscillator-based thermal monitor used in this work. Next, architectural details of the thermal benchmark circuit used for thermal characterization are presented. This is followed by a description of the characterization experiments that were performed. Lastly, this chapter provides results and analysis of the data obtained during characterization.

4.1 Architecture

4.1.1 Thermal Monitor

The thermal monitor architecture is shown in Figure 4.1. It counts the number of ring oscillator oscillations within a fixed period of time. This is achieved by having the ring oscillator generate the clock, whose frequency is dependent on temperature, of an incrementer circuit. The clock generated by the ring oscillator is called Ring Clock. Another incrementer, which is driven by a fixed system clock, is used to measure how many times the temperature-dependent incrementer counts over a fixed period of time. Since the ring oscillator period is a function of temperature, the count obtained varies with the temperature of the circuit.

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 (Ring Clock) drives 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 incrementer on the final output of the circuit, and a ready signal indicates the output is ready to be read. The 16-bit incrementer is then reset to 0. As the FPGA die temperature varies, the number of ring oscillator oscillations counted by

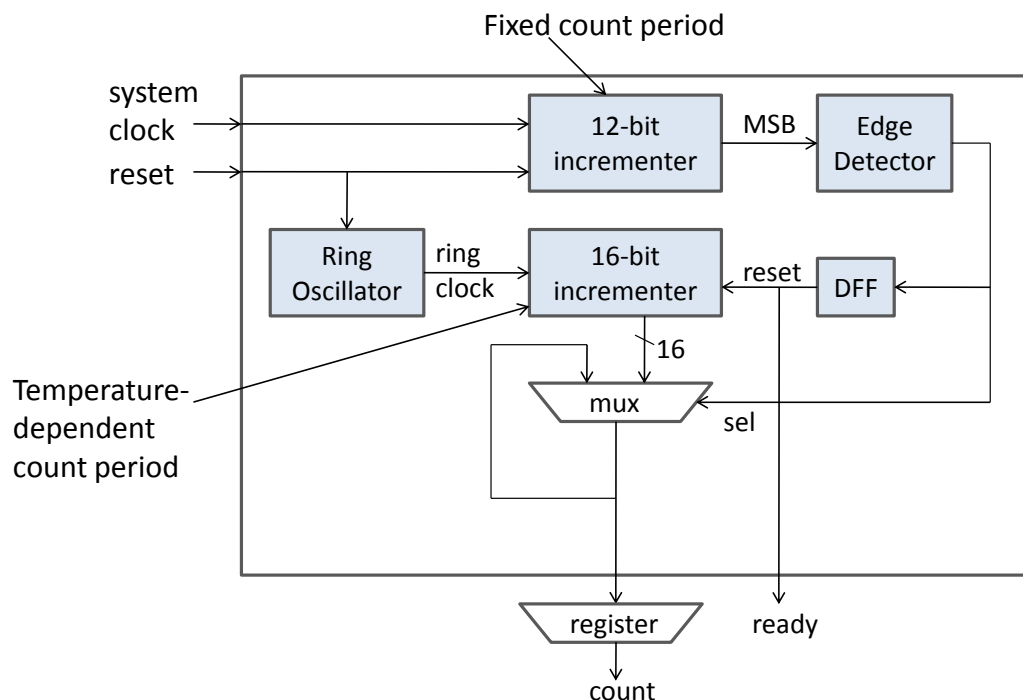


Figure 4.1: Thermal Monitor Architecture with ring oscillator and associated counter circuitry.

the 16-bit incrementer in a fixed time period changes due to the effect that temperature has on the period of oscillation.

Figure 4.2 illustrates the waveform representation of the thermal monitor's operation. It shows the dependence of the ring oscillator frequency on temperature. As temperature decreases, the period of oscillation reduces, resulting in an increase in the output frequency, and hence an increase in the measured incrementer count value. The basic idea of the thermal monitor described above is based on [Jones et al. (2007)].

Figure 4.3 shows the gate-level structure of the ring oscillator used in this work. It consists of 23 inverters connected in a loop. The OR gate is used to initialize the oscillator and the Ring Clock output generates the temperature-dependent clock. The typical frequency of operation at room temperature of this ring oscillator was $\sim 80\text{MHz}$. The number of inverters was chosen such that the count varies by 1 for every 0.25°C change in temperature.




Figure 4.2: Ring oscillator frequency temperature dependence example. The thermally dependent clock's frequency changes with temperature, giving different counts at the output of the thermal monitor during a fixed measurement period.

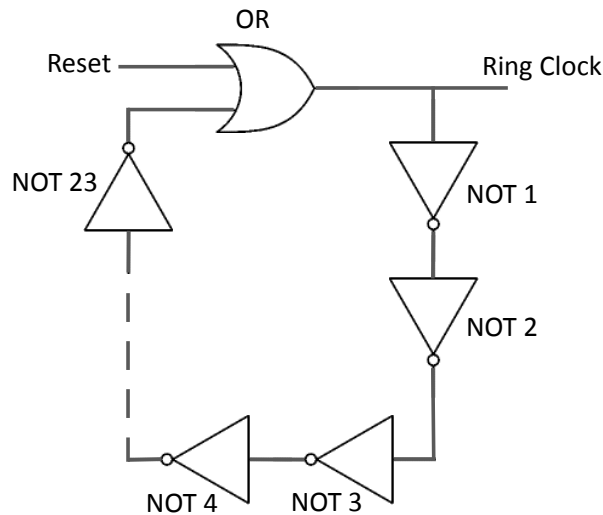


Figure 4.3: Ring Oscillator with 23 inverters in a loop. The OR gate is used to initialize oscillations. Output is obtained at the Ring Clock probe point.




Figure 4.4: Core Block consisting of a chain of D-type Flip Flops connected together. The LUTs between the flip flops in the figure are configured as AND gates.




Figure 4.5: Thermal Workload Unit consisting of an array of core blocks. The input generator is made up of a NOT gate and a D flip flop and is used to excite the workload unit.

4.1.2 Thermal Benchmark Circuit

The work in this thesis was implemented on two different Virtex-5 FPGAs, an LX110T mounted on a Xilinx XUP-V5 board [LX110T User Manual (2011)] and an LX330 mounted on a Hitech Global board [LX330 User Manual (2008)]. For both chips, a ring oscillator-based thermometer was placed in the middle of the chip and a flexible thermal benchmark circuit occupied the chip's resources. This benchmark circuit is based on a thermal benchmark architecture described in [Jones et al. (2006)].

Figure 4.4 shows the implementation of a Core Block using FPGA resources, which is the primary building block of the thermal benchmark architecture. Each core block is a chain of D-type flip-flops connected to each other through logic gates to form an array.

Figure 4.5 illustrates the formation of a Thermal Workload Unit using core blocks. A thermal workload unit consists of a chain of core blocks concatenated to form a Computation

Table 4.1: Size of a single workload unit on the LX110T and LX330

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

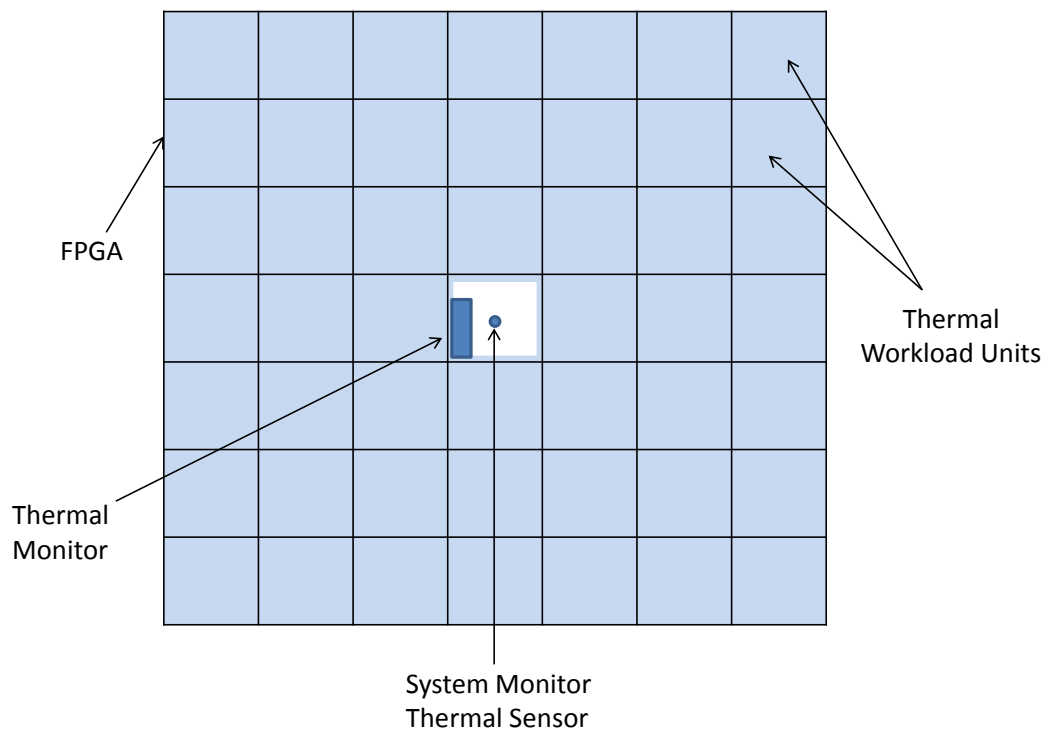


Figure 4.6: Workload placement on the FPGA. The ring oscillator-based thermometer is placed at the center of the die, surrounded by thermal workload units occupying the entire chip.

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. This results in maximum power consumption and heat generation inside the chip.

The benchmark circuit is composed of a number of workload units. They can be selectively enabled to control the amount the FPGA heats. Table 4.1 shows the resource utilization of one workload unit on the LX110T and LX330. The LX110T houses 53 workload units and the LX330 houses 54 workload units to achieve a maximum utilization of 88% and 86% respectively.

The workload units were placed on the FPGA using Xilinx relative location constraints (RLOCs) [Virtex-5 User Guide (2012)]. They were placed in a grid fashion as shown in Figure 4.6. The thermal monitor was placed at the center of the die to minimize the difference between

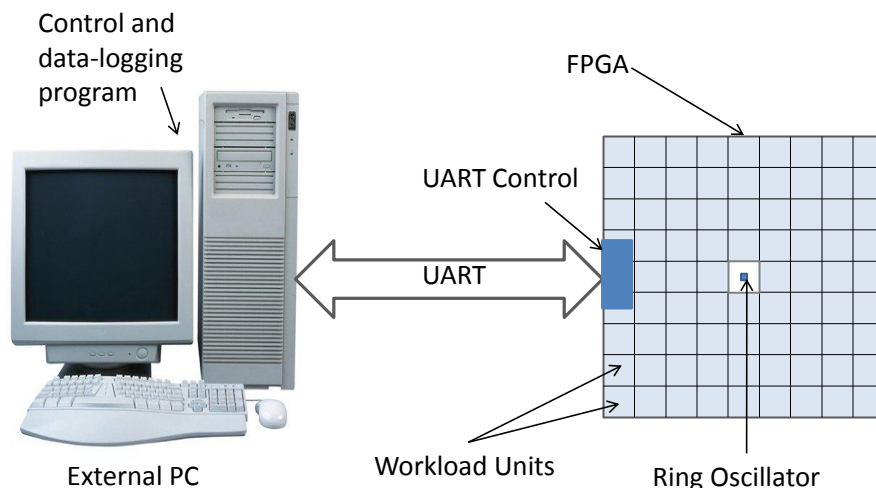


Figure 4.7: Hardware-software setup. The external PC is used to send commands to and receive data from the FPGA over a UART interface.

the temperature given by the ring oscillator-based thermometer and the temperature from the Xilinx System Monitor [System Monitor (2011)], which is located at the center of the die. This was useful when performing comparative studies between the temperature estimates from the ring oscillator calibrated using an external thermal probe and those from the System Monitor.

4.1.3 External UART/Command Interface

A program running on an external PC was used to selectively enable portions of the FPGA 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 commands over a UART interface that were acted upon by a command processing module deployed on the FPGA. Figure 4.7 shows the overall architecture of the measurement system. The pseudo code executed on the external PC is shown in Algorithm 1. Tests that enabled 0% through 80% of the available FPGA resources, in steps of 20%, at various frequencies were performed on both the Virtex-5 LX110T and Virtex-5 LX330.

Algorithm 1 C program to send enable/disable commands to and log measurement data from the FPGA.

```

Start
Initialize read and write UART ports
Initialize commands
while not end of test do
    Write command to UART port
    Read response from FPGA on UART port
    Log data in file
end while
Close port
Exit

```

4.1.4 Current Measurement

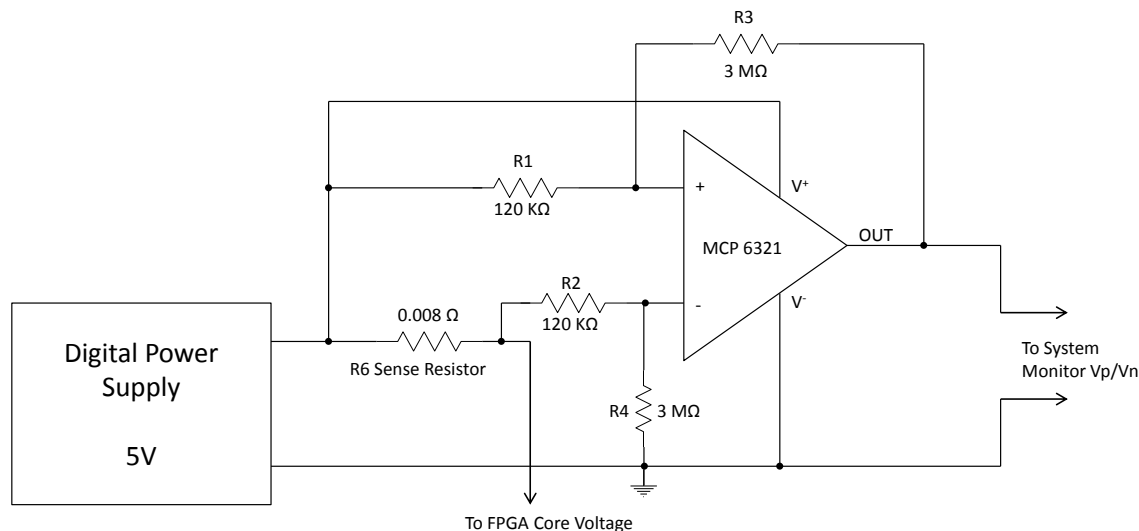


Figure 4.8: Difference Amplifier. The voltage across the sense resistor is amplified and provided as input to the System Monitor on the Virtex-5 for conversion.

An indirect method was used to characterize the effect of workload variation on the measurement obtained from the ring oscillator-based thermal monitor. Changes in power supply current draw were used to track workload changes of the thermal testbench circuit. A sense resistor of 0.008 ohms was connected in series with the power supply to help compute current draw. A Microchip MCP6231 op-amp was used, as shown in Figure 4.8, to amplify the voltage across the sense resistor by a factor of 81. The amplified voltage was then provided to the System Monitor's internal Analog to Digital converter (ADC) via two external Virtex-5

pins (V_p/V_n). The resulting 10-bit digital value from the ADC was then read by the UART controller on the FPGA.

Once received by the ADC, Equation 4.1 was used to convert the analog voltage to a 10-bit digital value [System Monitor (2011)].

$$V_{V_p/V_n} = V_{in}/(977 * 10^{-6}) \quad (4.1)$$

Where V_{V_p/V_n} is the analog to digital converted 10-bit value and V_{in} is the voltage applied at the System Monitor V_p/V_n pins.

The calculation performed to obtain the current draw from the power supply is given by Equation 4.2.

$$I_{measured} = (V_{V_p/V_n}/Gain)/R_{sense} \quad (4.2)$$

Where $I_{measured}$ is the current drawn from the supply in amperes, $Gain$ is the voltage gain of the difference amplifier used (in this case $Gain = 81$), and R_{sense} is the value of the sense resistor (0.008 ohms).

Chapter 5 describes how the measured current ($I_{measured}$) was used to compensate for non-thermal effects of workload variations on the ring oscillator-based thermal monitor measurements.

Figure 4.9 illustrates the architecture of the measurement setup after including supply current measurement instrumentation to the architecture depicted in Figure 4.7. A snapshot of the actual setup used for this work is shown in Figure 4.10.

4.2 Experimental Setup and Methodology

The design described in Section 4.1 was instantiated on two different FPGAs, the XC5VLX110T residing on a Xilinx XUPV5-LX110T board [LX110T User Manual (2011)] and the XC5VLX330 residing on a Hitech Global TB-5V-LX330-DDR2-E board [LX330 User Manual (2008)]. Characterization experiments were performed to correlate FPGA temperature, ring oscillator period, and current draw of the FPGA. These three pieces of data were collected over a wide range

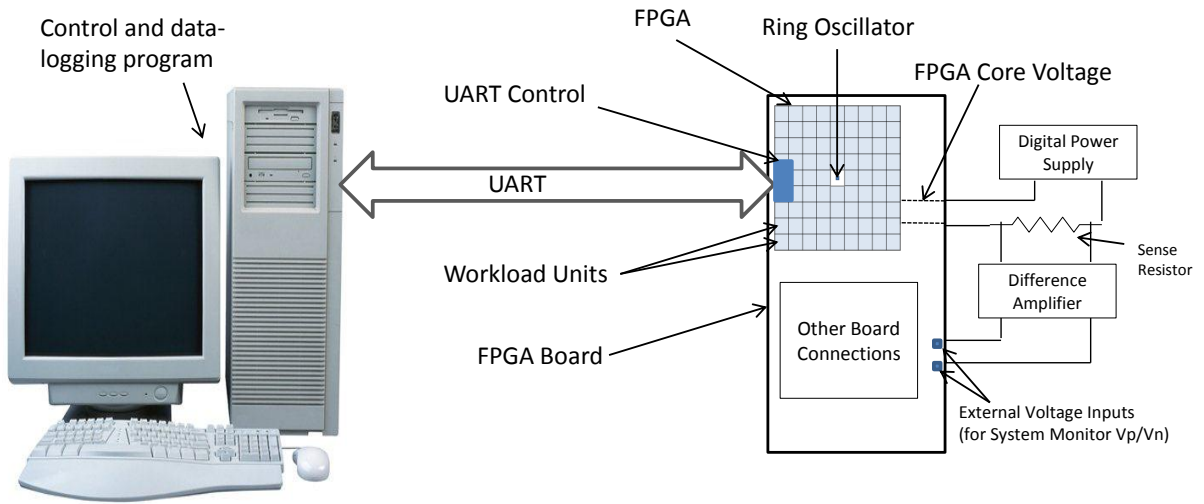


Figure 4.9: Complete Instrumentation System.




Figure 4.10: Actual hardware setup.

Table 4.2: Test Configurations. For each configuration, the following data was collected : 1) current from power supply, 2) FPGA case temperatures from thermal probe, 3) ring oscillator count.

Frequency(MHz)	Percentage of chip enabled				
	0	20	40	60	80
50					
100					
150					
200					

of operating points. Table 4.2 provides these specific operating points, in terms of operating frequency and logic resources.

A thermal probe was used to monitor the temperature of the FPGA case. This, as opposed to using the Xilinx System Monitor [System Monitor (2011)], was used for FPGA temperature monitoring due to an odd behavior of the System Monitor being observed at high temperatures. This behavior is described in Chapter 6.

Relying on a surface mounted thermal probe to measure temperature instead of the System Monitor was deemed reasonable based on the following analysis. 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 $0.10^{\circ}\text{C}/\text{W}$ to $0.15^{\circ}\text{C}/\text{W}$ in [Virtex-5 Packaging/Pinout Spec (2012)]. 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 0.75°C for the LX110T and 1.8°C for the LX330. Given the temperature ranges over which our experiments were performed, this maximum delta is acceptable.

These tests were performed to characterize the dependency of ring oscillator frequency on workload variations for a range of temperatures. The test procedures for conducting these experiments is 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.

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%, 40% & 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 to 0% FPGA utilization after reaching 80°C, the utilization was set to the target utilization 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 target utilization being tested.

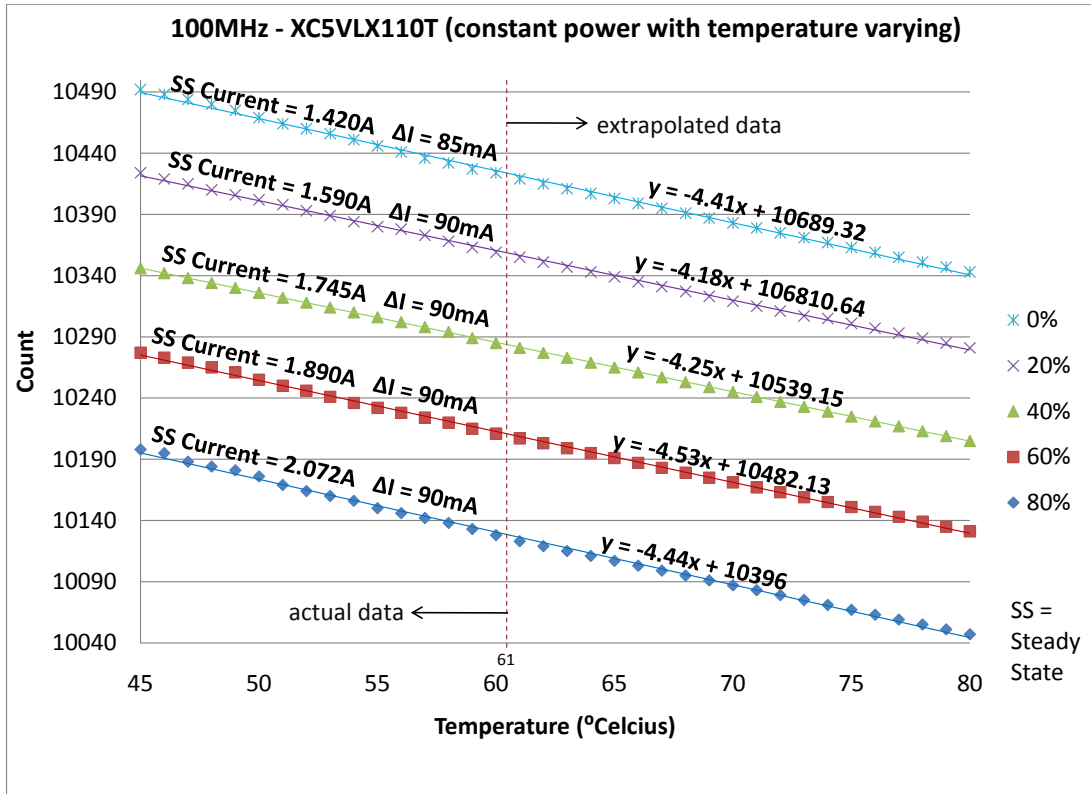
4.3 Results and Analysis

Figures 4.11(a) and 4.11(b) plot temperature versus thermal monitor count value for a subset of the tests conducted (100MHz only). Figure 4.11(a) shows the response for the XC5VLX110T FPGA, and Figure 4.11(b) shows the response for the XC5VLX330.

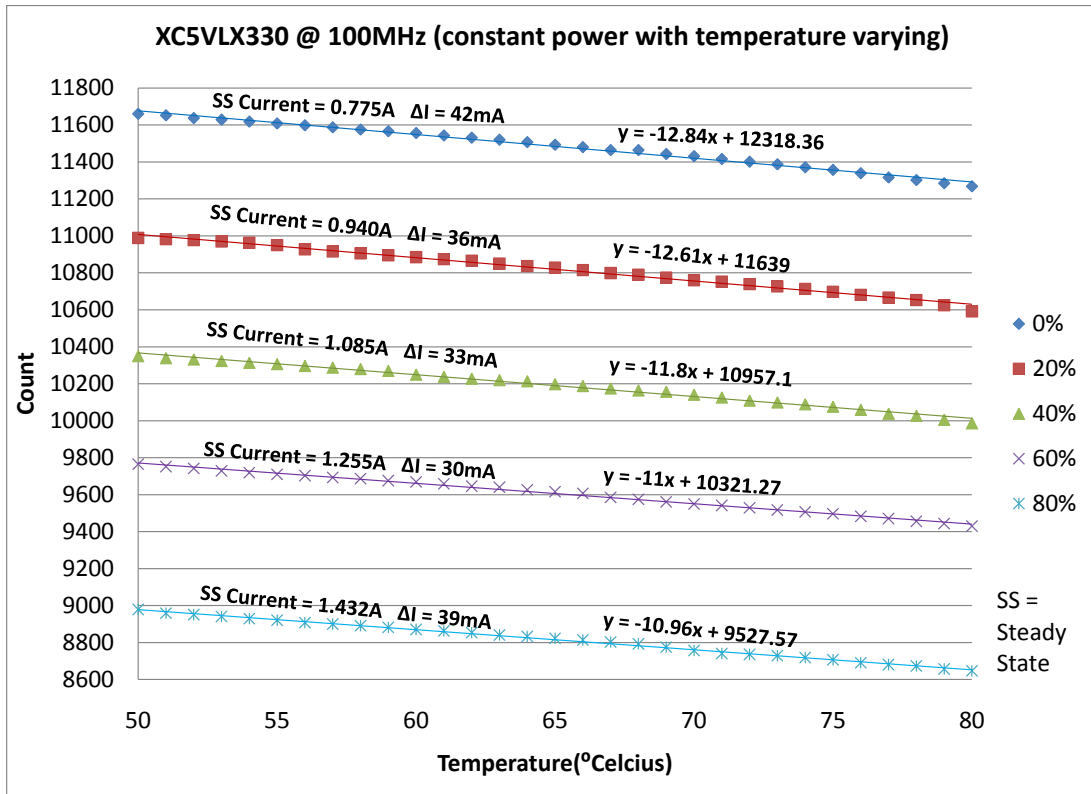
These graphs illustrate how workload variation impacts the relationship between the temperature of the FPGA and the count of the thermal monitor (i.e. ring oscillator period). The SS (Steady-State) Current indicates the current being drawn from the supply used to power the FPGA board after running a particular configuration 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 configurations 0% and 80% utilization causes a change of around 300 in the count obtained from the ring oscillator. The slopes of Figures 4.11(a) give a relation of ~ 4 counts per degree celcius. Assuming the 100MHz/0% utilization configuration was used to calibrate the ring oscillator-based thermometer, a change of 300 counts translates to a discrepancy of $\sim 75^\circ\text{C}$ in the estimated temperature. This emphasizes the necessity to

compensate for the effect of workload variation. It should be noted that although the plot for the XC5VLX330 shows measured values for the entire range, the plot for the XC5VLX110T contains extrapolated values for temperatures greater than 60°C.



(a)



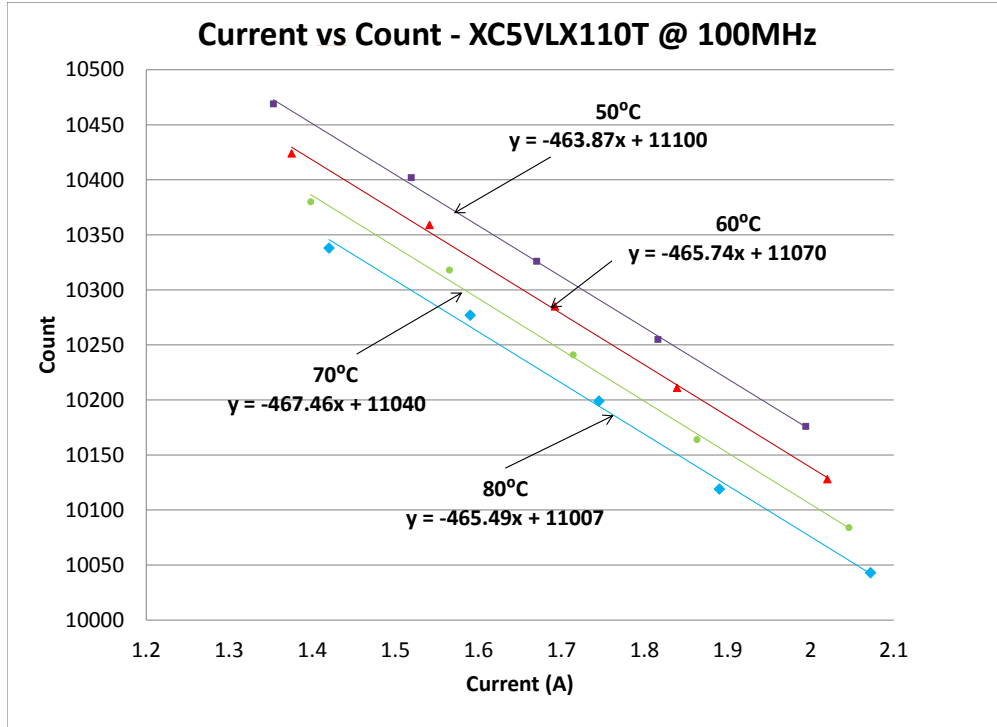
(b)

Figure 4.11: 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.

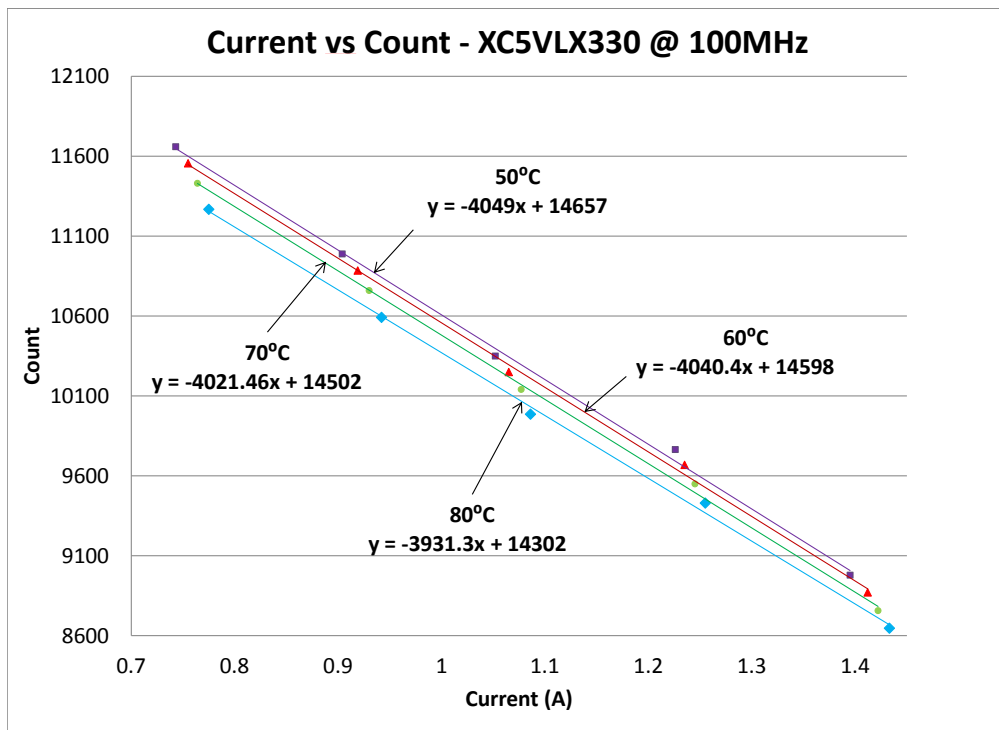
Alternately, the collected data can be plotted to illustrate how the ring oscillator count varies with current draw for constant values of temperature, shown in Figures 4.12(a) and 4.12(b). The points on each of the lines correspond to different configurations of the chip, from 0% utilization through 80% utilization in steps of 20%. The count values for both FPGAs show a near 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 4.12(a) and 4.12(b) are 465.64 (Counts/A) for the LX110T and 4010.49 (Counts/A) for the LX330. Given ~ 4 counts/ $^{\circ}\text{C}$, this translates to a temperature error of $\sim 1^{\circ}\text{C}$ per 8.6mA of current change (relative to a baseline workload configuration used to calibrate the ring oscillator-based thermometer) for the LX110T, and given ~ 12 counts/ $^{\circ}\text{C}$ this translates to a temperature error of $\sim 1^{\circ}\text{C}$ 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 thermometer 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)

Figure 4.12: 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 maximum difference of around 2700 in the count as compared to 300 in (a)

CHAPTER 5. Mitigating Impacts of Workload-variation on Ring Oscillator-based Thermometer Behavior

This chapter presents a technique to compensate for the impact of workload variation on ring oscillator-based thermometers. An approach that exploits the observed linearity of the ring oscillator-based thermometer response to FPGA current draw is described. The experiments performed to validate the effectiveness of the approach are then described. A discussion of the experimentation results obtained after applying this compensation technique concludes the chapter.

5.1 Approach

The technique proposed in this section aims to compensate for the shifts occurring in the thermal monitor response caused by workload variations. These shifts were characterized in Figure 4.11. Figure 5.1 shows the expected result from the proposed compensation technique. The goal is to obtain, for a given temperature, the same count value irrespective of the FPGA's application mode. In other words, for a given temperature in Figures 4.11(a) and 4.11(b), the count value measured should be independent of the FPGA workload. We assume the 0% utilization configuration was used to calibrate the ring oscillator-based thermometer. The remainder of this section derives compensation equations from the data obtained during the characterization phase (Chapter 4).

The plots in Figures 4.12(a) and 4.12(b) show the variation of the ring oscillator counts with respect to current draw of the chip for constant temperatures, thus the slopes of these lines correlate changes in current draw to changes in count values ($Slope_{count_per_amp}$). This correlation was used to compensate for workload variation as follows. First the difference

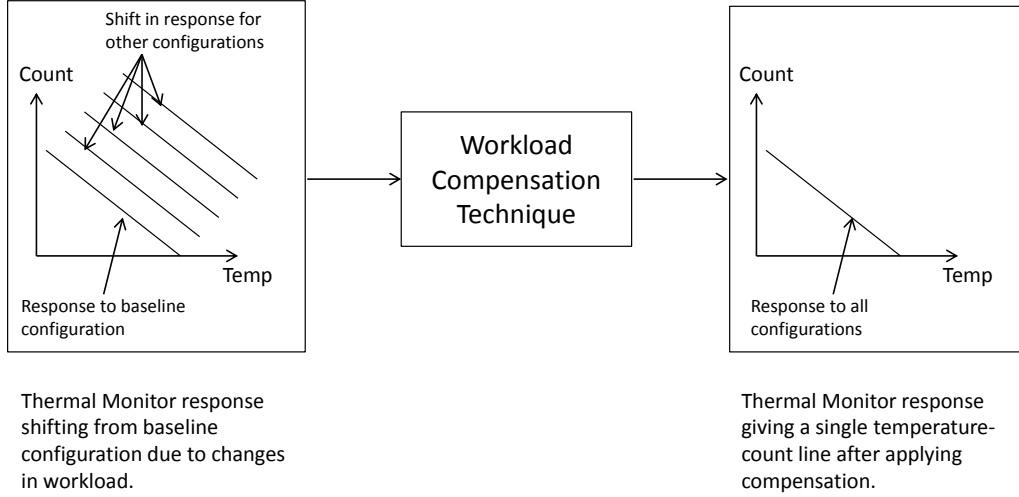


Figure 5.1: Workload-variation compensation. The plot on the left shows shifts in the counts received from the thermal monitor due to changes in workload. The compensation technique in the middle makes use of the ring oscillator frequency's linear current dependence obtained in Figures 4.12(a) and 4.12(b) to correct the shifts. The plot on the right shows the expected response of the thermal monitor after the compensation technique is applied, with all workload configurations yielding the same response as the baseline configuration.

between the baselines configuration's current draw ($I_{base.config}$) and the presently executing configuration's current draw ($I_{measured}$) was computed.

$$\Delta I = I_{measured} - I_{base.config} \quad (5.1)$$

For the LX110T chip, $I_{base.config} = 1.42A$ for the 100MHz/0% utilization configuration.

Then, using the data from Figure 4.12(a), the correlation between count values and current draw ($Slope_{count.per.amp}$) allows us to compute the compensated count value (C_{comp}) as:

$$C_{comp} = C_{measured} + (\Delta I * Slope_{count.per.amp}) \quad (5.2)$$

Where $C_{measured}$ is the thermal monitor count for the configuration executing. For the LX110T chip, the $Slope_{count.per.amp}$ was found to be 465 counts per ampere.

The following example shows how the set of equations above was used to obtain the C_{comp} values for a 100MHz/80% utilization: For the 100MHz/80% utilization, the measured current value, $I_{measured} = 2.044A$. Using the $I_{base.config}$ value of 1.42A, we obtain the difference in

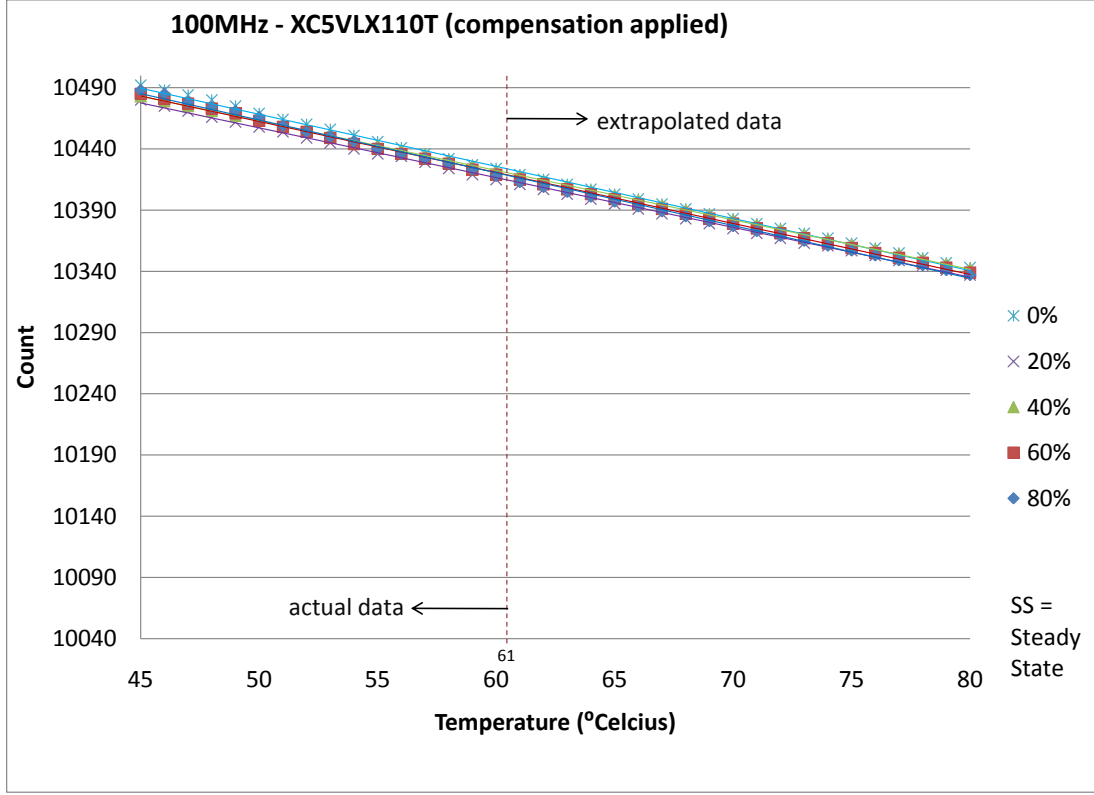


Figure 5.2: Compensation applied on Figure 4.11(a). A single response line is obtained after applying the compensation technique, giving a unique value of count at a given temperature in every mode.

current draw using Equation 5.1 as:

$$\Delta I = 2.044 - 1.42 = 0.624 \quad (5.3)$$

Plugging this value in Equation 5.2 and using the thermal monitor count values $C_{measured}$ and the $Slope_{count_per_amp}$ value of 465 for the LX110T, the compensated count values can be obtained as follows.

$$C_{comp} = C_{measured} + (0.624 * 465) \quad (5.4)$$

Therefore, for the 100MHz/80% configuration on the Virtex-5 LX110T, the compensated count values are given by,

$$C_{comp} = C_{measured} + 290 \quad (5.5)$$

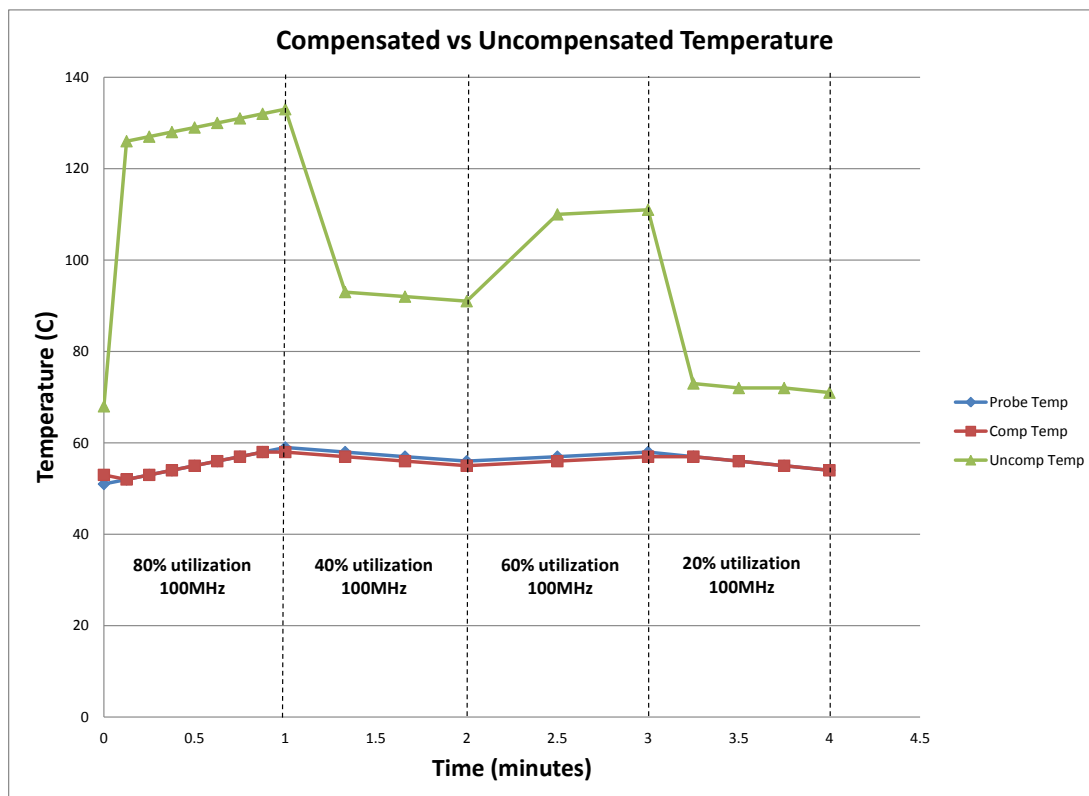


Figure 5.3: Comparison of compensated and uncompensated temperature estimates on LX110T.

5.2 Results and Analysis

Applying the compensation technique described in the previous section gives a temperature versus thermal monitor count plot as shown in Figure 5.2. This plot shows as expected, that the proposed technique compensates for the effect of FPGA workload variation on thermal monitor response.

An experiment was also performed to evaluate the effectiveness of this approach in an environment where over time the FPGA workload varies. During this experiment, the percentage utilization of the LX110T FPGA was changed every minute while the system ran at a constant frequency, and response count values were collected. The application was run using 4 different utilizations as shown in Table 5.1.

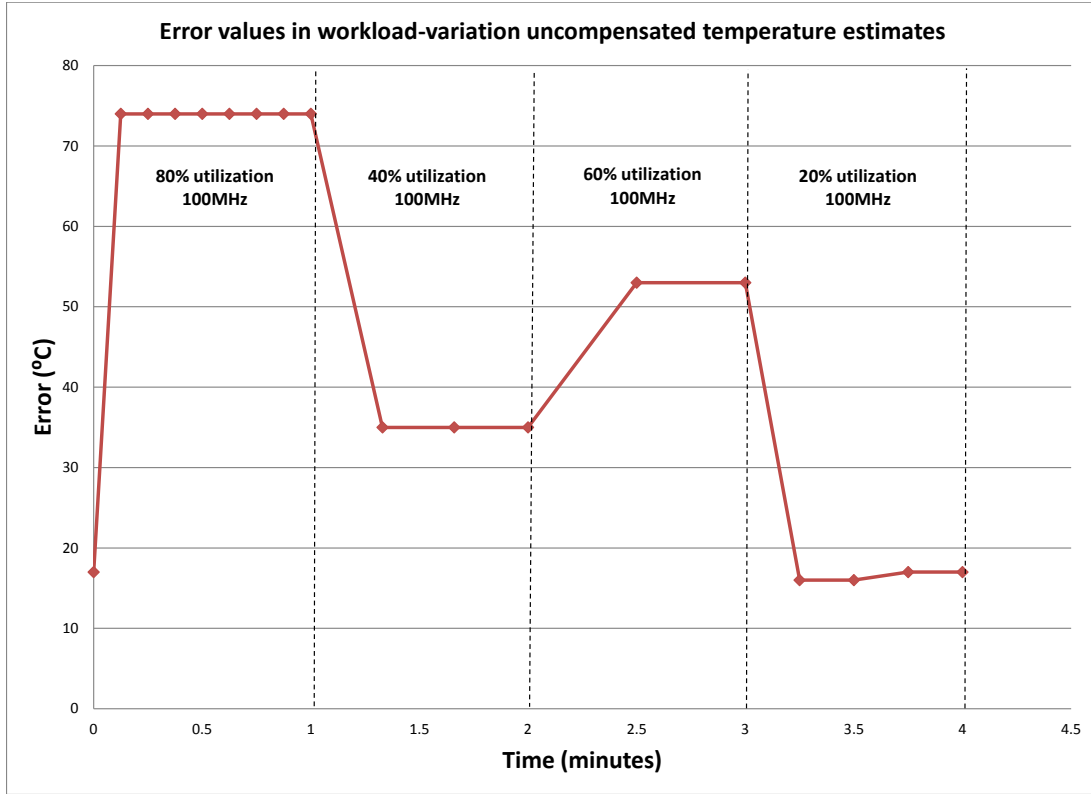
The data obtained after running this experiment is shown in Figure 5.3, which compares the following over time: 1) the temperature from the thermal probe mounted on the FPGA (Probe Temp), 2) the temperature estimate from the uncompensated counts obtained during

Table 5.1: Compensation test table.

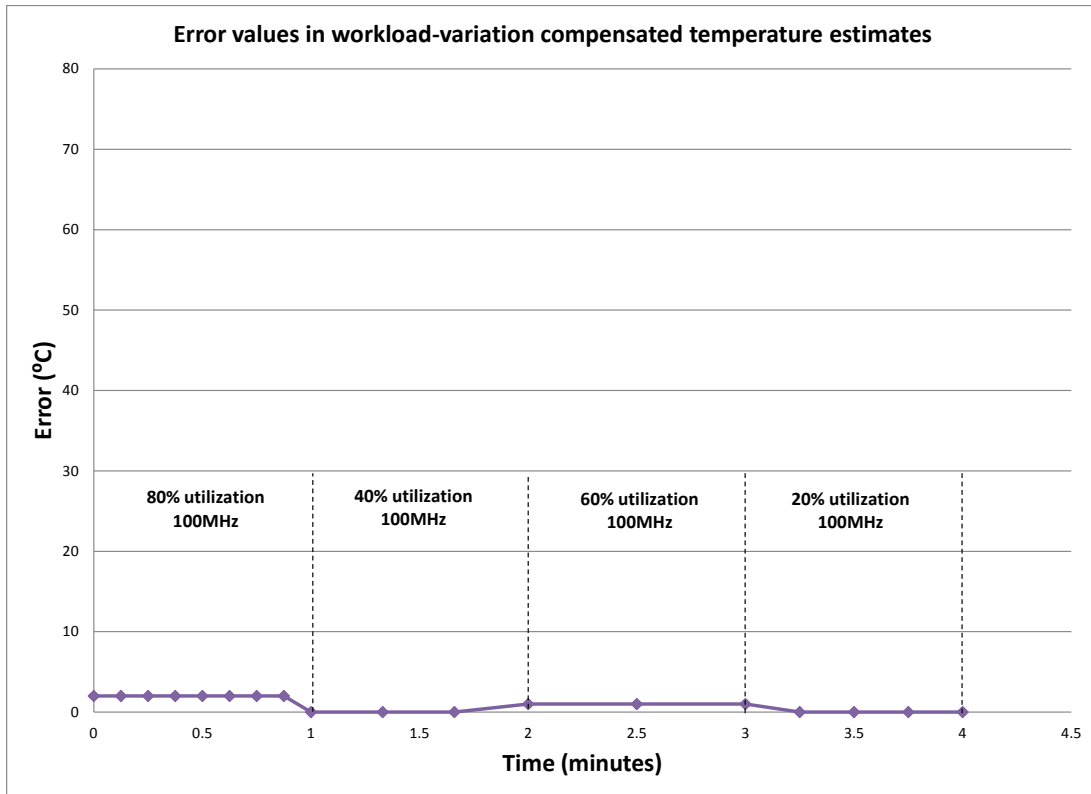
Time(minutes)	% Utilization
0-1	80
1-2	40
2-3	60
3-4	20

the characterization phase (Uncomp Temp), and 3) the temperature estimates obtained after applying the compensation technique (Comp Temp). To obtain the temperature values for the Uncomp Temp plot, the count values obtained in Figure 4.11(a) were plugged into the straight line equation of the baseline 0% utilization configuration, since that is used to calibrate the thermometer. To obtain the workload-compensated temperature estimates, the compensated count values were plugged into the straight line equation for the baseline configuration.

Figure 5.4(a) shows the error (in °C) in temperature values estimated using the uncompensated counts from the thermal probe-measured temperatures, and Figure 5.4(b) shows the error (in °C) in temperature values estimated using the compensated counts from the thermal probe-measured temperatures. It is evident from the plots that greater utilizations of the chip result in larger errors in estimated temperatures. Before compensation, the 100MHz/80% utilization configuration on the Virtex-5 LX110T gives a maximum error of 74°C in the estimated temperature. This is understandable since the current draw is the maximum for this configuration. After applying the compensation technique, this error is reduced to 2°C.



(a)



(b)


Figure 5.4: Error in temperature measurement between a) uncompensated estimates and the thermal probe, and b) compensated estimates and the thermal probe, for various workload configurations.

CHAPTER 6. Unexpected System Monitor Behavior

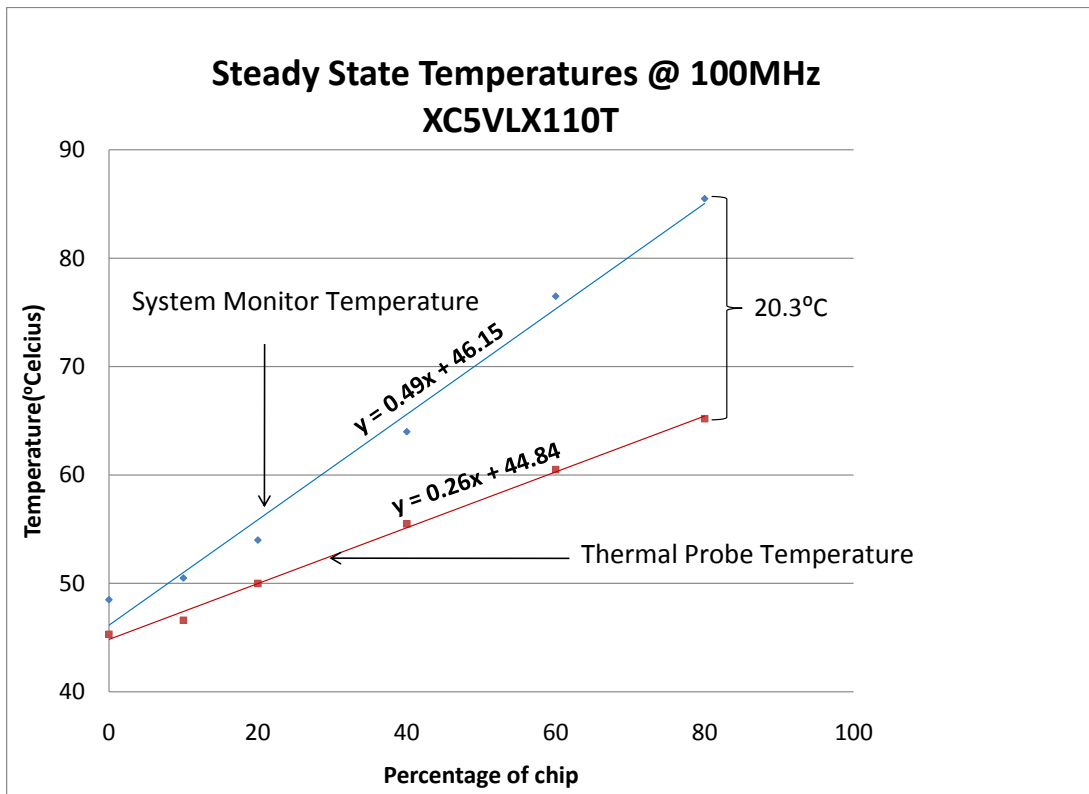
This chapter describes an odd behavior observed when using the Xilinx System Monitor for measuring high temperatures on two Xilinx Virtex-5 FPGAs. It compares the temperatures obtained from the System Monitor to those obtained using a thermal probe kept in contact with the FPGA case. Rationale on why the case temperature can be relied to question the validity of the System Monitor reported temperature is also provided.

Figures 6.1(a) and 6.1(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 (junction temperature) and by 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.

Examining the plot for the XC5VLX50T (Figure 6.1(a)) shows when 0% of the chip was enabled that the System Monitor reported the temperature as 43.5°C, while the thermal probe reported 39.7°C (a difference of 3.8°C). This was within the range of error for System Monitor measurements (4°C), as specified in [System Monitor (2011)]. However, as the temperature of the FPGA increased, the temperature reported by the System Monitor rose much faster than that reported by the thermal probe. At 80% utilization, a steady state temperature of 60.5°C was given by the System Monitor, while the probe measured the case temperature as 47.8°C, a difference of 12.7°C. For the XC5VLX110T FPGA (Figure 6.1(b)), this difference was even larger at higher temperatures. For this chip, at 0% utilization, the difference between System Monitor and probe temperatures was 3.2°C (System Monitor showing 48.5°C and probe showing 45.3°C) and at 80% utilization was 20.3°C (System Monitor showing 85.5°C and probe



(a)



(b)

Figure 6.1: Steady-state temperatures for various utilizations. The figures show an increased temperature difference, as the chip temperature rises, between the System Monitor reported values and values obtained using a surface mounted thermal probe.

showing 65.2°C).

A temperature difference between the System Monitor and probe is to be expected based on the thermal resistance between the junction and case, denoted by θ_{jc} . θ_{jc} is specified to be 0.10°C/W to 0.15°C/W for the Virtex-5 FPGA in [Virtex-5 User Guide (2012)]. Given the maximum power utilized by either of these chips was 5W for these experiments, the thermal resistance only accounts for a maximum expected error of 0.75°C. Thus, the temperature deviations of 12.7°C and 20.3°C are significantly larger than expected. Potential reasons for this unexpected behavior are still under investigation.

CHAPTER 7. Conclusions and Future Work

This work has described a technique to mitigate the impact of workload variation on ring oscillator-based thermometers. First, a method to characterize the effects of workload variation on ring oscillator frequency response has been proposed and implemented. Temperature and current related data for two Xilinx Virtex-5 FPGAs was obtained at frequencies of 50MHz, 100MHz, 150MHz and 200MHz, and at resource utilizations of 0%, 20%, 40%, 60% and 80%. A ring oscillator-based thermal monitor and a flexible thermal benchmark circuit have been designed and deployed on the two Xilinx FPGAs to obtain characterization data. A mechanism to measure the current draw from the power supply has also been described. A complete hardware-software solution was developed to log measurement data from the FPGA to an off-chip computer in real time. On the Virtex-5 LX110T FPGA, workload variation dependence of the ring oscillator-based thermometer resulted in $\sim 1^\circ\text{C}$ error in estimated temperature for every 8.6mA change in current draw of the FPGA. The maximum error in estimated temperature on this device was observed to be 74°C .

The characterization results were used to derive a set of equations that compensated for the non-ideal effects of workload variation on ring oscillator response. The compensation technique reduced the maximum error in temperature estimates from 74°C to 2°C .

Lastly, an unexpected behavior when using the Xilinx System Monitor at high on-chip temperatures has been discussed. When measuring high temperatures, a maximum temperature difference of 20.3°C is observed between the System Monitor reported temperature and the temperature obtained from a case-mounted thermal probe.

Future work in this area includes: 1) placing the ring oscillator in different locations of the FPGA to characterize the impact of process variations within the FPGA on a ring oscillator-based thermometer, 2) measuring voltage variations directly across the voltage regulator used

to power the FPGA to obtain direct characterization of voltage variation-induced effects on ring oscillator-based thermometers, and 3) characterizing the impact of hard blocks (such as Block RAMs, DSPs, Embedded Processors, etc.) present on the FPGA and external I/O peripherals present on the board, on ring oscillator frequencies.

BIBLIOGRAPHY

- AMD Bulldozer (2011). The bulldozer review: Amd fx-8150 tested. <http://www.anandtech.com/show/4955/the-bulldozer-review-amd-fx8150-tested>.
- Bakker, A. and Huijsing, J. (1996). Micropower cmos temperature sensor with digital output. *Solid-State Circuits, IEEE Journal of*, 31(7):933 –937.
- Becker, T., Jamieson, P., Luk, W., Cheung, P., and Rissa, T. (2009). Power characterisation for the fabric in fine-grain reconfigurable architectures. In *Programmable Logic, 2009. SPL. 5th Southern Conference on*, pages 77 –82.
- Boemo, E. I. and López-Buedo, S. (1997). Thermal monitoring on fpgas using ring-oscillators. In *Proceedings of the 7th International Workshop on Field-Programmable Logic and Applications*, pages 69–78, London, UK. Springer-Verlag.
- Bohr, M. (2002). Intel’s 90nm technology: Moore’s law and more. <http://www.cs.virginia.edu/~mc2zk/cs451/184401.pdf>.
- Borkar, S. (1999). Design challenges of technology scaling. *Micro, IEEE*, 19(4):23 –29.
- Brooks, D. and Martonosi, M. (2001). Dynamic thermal management for high-performance microprocessors. In *High-Performance Computer Architecture, 2001. HPCA. The Seventh International Symposium on*, pages 171 –182.
- Chakraborty, A. and Pradhan, S. (2012). A technique for power reduction of cmos circuit at 65nm technology. In *Recent Advances in Information Technology (RAIT), 2012 1st International Conference on*, pages 576 –580.

- Filanovsky, I. and Allam, A. (2001). Mutual compensation of mobility and threshold voltage temperature effects with applications in cmos circuits. *Circuits and Systems I: Fundamental Theory and Applications, IEEE Transactions on*, 48(7):876 –884.
- Franco, J., Boemo, E., Castillo, E., and Parrilla, L. (2010). Ring oscillators as thermal sensors in fpgas: Experiments in low voltage. In *Programmable Logic Conference (SPL), 2010 VI Southern*, pages 133 –137.
- Ghani, T. (2009). Challenges and innovations in nano-cmos transistor scaling. http://www.cse.nd.edu/courses/cse60547/www/lectures/03_Neikei_Presentation_2009_Tahir_Ghani.pdf.
- Happe, M., Agne, A., and Plessl, C. (2011). Measuring and predicting temperature distributions on fpgas at run-time. In *Reconfigurable Computing and FPGAs (ReConFig), 2011 International Conference on*, pages 55 –60.
- Hasan, M. and Bird, M. (2011). Energy reductions for embedded processors in reconfigurable hardware. In *Electro/Information Technology (EIT), 2011 IEEE International Conference on*, pages 1 –8.
- Heo, S., Barr, K., and Asanovic, K. (2003). Reducing power density through activity migration. In *Low Power Electronics and Design, 2003. ISLPED '03. Proceedings of the 2003 International Symposium on*, pages 217 – 222.
- Huang, W., Rajamani, K., Stan, M., and Skadron, K. (2011). Scaling with design constraints: Predicting the future of big chips. *Micro, IEEE*, 31(4):16 –29.
- Intel 22nm Technology (2011). 3-d, 22nm: New technology delivers an unprecedented combination of performance and power efficiency. <http://www.intel.com/content/www/us/en/silicon-innovations/intel-22nm-technology.html>.
- Intel Press Release (2004). Intel drives moores law forward with 65nm process technology. http://download.intel.com/museum/Moores_Law/Articles-Press_Releases/Press_Release_Aug2004.pdf.

- Intel Processors Online (2012). Compare intel products online. <http://ark.intel.com/compare/27492,58664>.
- ISSCC Trends Report (2011). International solid-state circuits conference. <http://isscc.org>.
- ITRS (2009). International technology roadmap for semiconductors. <http://public.itrs.net>.
- Jiang, C., Xu, X., Wan, J., and You, X. (2008). Energy management for microprocessor systems: Challenges and existing solutions. In *Intelligent Information Technology Application Workshops, 2008. IITAW '08. International Symposium on*, pages 1071 –1076.
- Jones, P., Moscola, J., Cho, Y., and Lockwood, J. (2007). Adaptive thermoregulation for applications on reconfigurable devices. In *Field Programmable Logic and Applications, 2007. FPL 2007. International Conference on*, pages 246 –253.
- Jones, P. H., Lockwood, J. W., and Cho, Y. H. (2006). A thermal management and profiling method for reconfigurable hardware applications. In *Field Programmable Logic and Applications, 2006. FPL '06. International Conference on*, pages 1 –7.
- Kim, N., Austin, T., Baauw, D., Mudge, T., Flautner, K., Hu, J., Irwin, M., Kandemir, M., and Narayanan, V. (2003). Leakage current: Moore's law meets static power. *Computer*, 36(12):68 – 75.
- Ku, J. C. and Ismail, Y. (2007). On the scaling of temperature-dependent effects. *Computer-Aided Design of Integrated Circuits and Systems, IEEE Transactions on*, 26(10):1882 –1888.
- Lopez-Buedo, S., Garrido, J., and Boemo, E. (1998). Thermal testing on programmable logic devices. In *Circuits and Systems, 1998. ISCAS '98. Proceedings of the 1998 IEEE International Symposium on*, volume 2, pages 240 –243 vol.2.
- Lopez-Buedo, S., Garrido, J., and Boemo, E. (2000). Thermal testing on reconfigurable computers. *Design Test of Computers, IEEE*, 17(1):84 –91.

- Lopez-Buedo, S., Garrido, J., and Boemo, E. (2002). Dynamically inserting, operating, and eliminating thermal sensors of fpga-based systems. *Components and Packaging Technologies, IEEE Transactions on*, 25(4):561 – 566.
- LX110T User Manual (2011). *ug347 ML505/ML506/ML507 Evaluation Platform User Guide*. Xilinx Inc.
- LX330 User Manual (2008). *Virtex 5 LX330 DDR2 II Memory image processing ASIC prototyping Board User Manual*. Hitech Global.
- Mangalagiri, P., Bae, S., Krishnan, R., Xie, Y., and Narayanan, V. (2008). Thermal-aware reliability analysis for platform fpgas. In *Computer-Aided Design, 2008. ICCAD 2008. IEEE/ACM International Conference on*, pages 722 –727.
- Moore’s Law Inspires Intel Innovation (2011). Moore’s law inspires intel innovation. <http://www.intel.com/content/www/us/en/silicon-innovations/moores-law-technology.html>.
- Quenot, G., Paris, N., and Zavidovique, B. (1991). A temperature and voltage measurement cell for vlsi circuits. In *Euro ASIC '91*, pages 334 –338.
- Rusu, S. (2004). Trends and challenges in high-performance microprocessor design. <http://www.eda.org/edps/edp04/submissions/presentationRusu.pdf>.
- Skadron, K., Stan, M., Huang, W., Velusamy, S., Sankaranarayanan, K., and Tarjan, D. (2003). Temperature-aware microarchitecture. In *Computer Architecture, 2003. Proceedings. 30th Annual International Symposium on*, pages 2 – 13.
- System Monitor (2011). *ug192 Virtex-5 FPGA System Monitor User Guide*. Xilinx Inc.
- Szekely, V., Marta, C., Kohari, Z., and Rencz, M. (1997). Cmos sensors for on-line thermal monitoring of vlsi circuits. *Very Large Scale Integration (VLSI) Systems, IEEE Transactions on*, 5(3):270 –276.

- Velusamy, S., Huang, W., Lach, J., Stan, M., and Skadron, K. (2005). Monitoring temperature in fpga based socs. In *Computer Design: VLSI in Computers and Processors, 2005. ICCD 2005. Proceedings. 2005 IEEE International Conference on*, pages 634 – 637.
- Virtex-5 Packaging/Pinout Spec (2012). *ug195 Virtex-5 FPGA Packaging and Pinout Specification*. Xilinx Inc.
- Virtex-5 User Guide (2012). *ug190 Virtex-5 FPGA User Guide*. Xilinx Inc.
- Yung, R., Rusu, S., and Shoemaker, K. (2002). Future trend of microprocessor design. In *Solid-State Circuits Conference, 2002. ESSCIRC 2002. Proceedings of the 28th European*, pages 43 –46.
- Zhang, X., Jouini, W., Leray, P., and Palicot, J. (2010). Temperature-power consumption relationship and hot-spot migration for fpga-based system. In *Green Computing and Communications (GreenCom), 2010 IEEE/ACM Int'l Conference on Int'l Conference on Cyber, Physical and Social Computing (CPSCoM)*, pages 392 –397.
- Zick, K. M. and Hayes, J. P. (2010). On-line sensing for healthier fpga systems. In *Proceedings of the 18th annual ACM/SIGDA international symposium on Field programmable gate arrays*, FPGA '10, pages 239–248, New York, NY, USA. ACM.