Implementation of a fast relative digital temperature sensor to achieve thermal protection in Zynq SoC technology

Charles-Alexis Lefebvre*, Jose Luis Montero, Leire Rubio

IK4-Ikerlan, Paseo J.M. Arizmendiarrieta, Arrasate-Mondragón 20500, Spain

ABSTRACT

More and more industrial embedded systems are developed to undergo hard environmental conditions, especially high temperatures. To prevent this impact, environmental conditions (e.g. the temperature) could be monitored. Plenty of new industrial designs are built around SoCs, and more especially around the Zynq-7000 introduced by Xilinx in 2011. In fact, monitoring the temperature inside the Zynq has become a challenge. While many applications focused on precision, the application proposed here instead is in an industrial context and aims at detecting a temperature excess as fast as possible to achieve the thermal protection of a logic area of the chip. Most of the digital sensors designed require a calibration to be operational. Such a process is not viable for time to market, and a solution must be found to either lighten it (e.g. by doing a simple 3-points calibration) or simply avoiding it. Instead of measuring the temperature in an absolute way, this paper focuses on detecting if the temperature is above or below a threshold. This work exhibits the implementation of three temperature digital sensors with promising results on Zynq technology. Two of the presented sensors are based on a ring-oscillator and another uses a flip-flop as a sensing element. Results show that a temperature increase can be detected in less than 1ms without any calibration protocol and this sensor was found to perfectly fit the targeted application.

© 2017 Published by ELSEVIER COMPANY.

1. Introduction

Since the release of the Zynq-7000 System-on-Chip (SoC) by Xilinx in 2011, powerful adaptable industrial embedded systems have been designed. The Zynq-7000, designed in a 28 nm technology, joins the power of a Processing System (PS) based on a dual-core ARM Cortex-A9 processor and the modularity of a Programmable Logic (PL). Many of those target applications in hostile environments where boards that may endure strong temperature conditions. Even though the Zynq comes with an analog built-in sensor to measure the global temperature of the chip, its fixed position prevents it from protecting a precise region implementing a critical function.

The temperature protection of an area is a serious issue when conforming hardware with industrial norms. In order to comply with them, a new type of sensor is needed. For a temperature sensor which can be implemented anywhere in the PL, a full-digital solution based on hardware logic must be used. Our first researches on this topic were presented in [1] which introduced a ring-oscillator-based

sensor and its working on a Zynq Z-7020 to achieve temperature measurements. However, in order to comply with the industrial norms, measuring the temperature is not necessarily compulsory if hazardous temperature changes can be monitored. In this paper, the application proposed instead consists of monitoring whether the temperature exceeds or not a defined threshold.

In this paper, 3 different digital sensor architectures are presented, implemented on a Zynq Z-7020 and studied. The goal is to obtain a sensor with a fast temperature detection time in case of a strong temperature increase. This way, a specified area of the chip can be protected. The main contribution of this paper consists of implementing a ring-oscillator based sensor on a Xilinx Zynq-7000 chip and using it in an unusual way in order to achieve temperature detection and protection, avoiding calibration processes to enhance both reproducibility and time to market.

The rest of this article is organized as follows. Section 2 presents the state-of-the-art of the actual digital-based temperature monitoring techniques. The methodology used in this paper is presented in Section 3. The different implemented and tested architectures and their abilities are presented in Sections 4, 5 and 6. Section 7 explains how the thermal protection function was obtained, and the conclusions are given in Section 8.

---

* Corresponding author.
  E-mail address: calefebvre@ikerlan.es (C.-A. Lefebvre).

2. Background

The number of available contributions [1-4,6,8-10] makes digital thermal sensors for FPGAs a hot topic. The main advantage over integrated analog sensor has been identified as the possibility of reconfiguration. After introducing the targeted chip, two types of digital sensor principles described in the literature will be presented: one based on a ring-oscillator and the other on the metastable effect in a flip-flop.

2.1. Xilinx-7000 SoC

The applications presented in this article target a Xilinx Zynq 7z020 SoC. The Zynq-7000 series is a 28 nm SoC structured around two blocks as illustrated by Fig. 1. The PS block contains, among others, a dual-core ARM processor. The second block, the PL, provides the Field Programmable Gate Array (FPGA) function. This block is divided into Configurable Logic Blocks (CLB) containing 2 flip-flops and 16 Look-Up Tables (LUT) where logic functions can be implemented. Thanks to these internal primitives, digital thermal sensors were designed and implemented in the PL part. An analog built-in thermal sensor with a $\pm 4°C$ precision called XADC is present in the Zynq-7000. However, its fixed position prevents it from being used to monitor a local area inside the chip.

2.2. Ring oscillators

Ring-oscillators-based sensors are widely used to measure the temperature inside FPGAs thanks to intrinsic resonance phenomenons. A ring oscillator (RO) is designed by looping an odd number of NOT gates. Its frequency is given by Eq. (1):

$$f = \frac{1}{2N\tau} \tag{1}$$

where N is the number of gates and τ is the propagation delay of a single gate. This type of oscillator is sensible to physicals parameters including temperature, as widely proved in past publications [2,3] and more theoretically deepen in [4]. The variations of the frequency

Fig. 1. Xilinx Zynq-7000 SoC architecture.

caused by the temperature are captured usually with a counter. The value of this counter gives an image of the temperature.

2.3. Flip-Flop Metastability based sensors

Contrary to the RO-based sensors, Flip-Flop-Metastability based sensors (FFM) are not widely studied. Even though the phenomenon and its detection were highlighted by papers from Xilinx and Altera, its use to measure temperature was first exposed in [5].

In a normal way of working, the input of a flip-flop is copied to its output after each rising edge of clock. That is why the data signal must be set up during a t_{setup} time before the rising edge and hold during a t_{hold} time. If these delays are violated, and a transition occurs at the same time, the flip-flop can enter in a transitory state called metastable state. When a flip-flop is in this type of state, different behaviors are observed, as listed in [6] and [7]. One effect cited among others in such a state is the increase of the clock-to-q delay (i.e. the delay between the rising edge of the clock and the edge of the output) which is temperature-dependent as showed in [5]. The idea of FFM-based sensors is to force a flip-flop to be in a metastable state to make it sensible to temperature.

Based on this phenomenon, designs in programmable devices were implemented in order to measure temperature inside a FPGA [5].

3. Methodology

Under this section, the methodology used to construct temperature sensors using specific FPGA primitives is described. All the digital sensors presented were implemented into Xilinx Zynq technology. They follow the same global architecture: a sensing element sensible to the temperature, a counter to count the number of digital events from the sensing part and a control unit to manage the measuring process. In particular, the control unit can activate or deactivate the sensing element thanks to an enable signal to avoid self-heating and over-consumption of the sensor. The sensors were designed as AXI-Lite IPs (Intellectual Property) to be able to communicate with the PS when instantiated in the PL. Registers were used to set the configuration. The number of gates in the ring can be chosen when instantiating the IP.

All the sensors were tested on an evaluation Zedboard equipped with a commercial version of the Zynq Z-7020 chip whose temperature range is [0° C; +85° C] and powered under 1V. They were successively instantiated in the same position of the PL and communicated to the PS. The temperature was monitored with the Zynq's built-in sensor and a K-thermocouple was stuck on its surface. The data from the counter was retrieved to a computer with an Ethernet connection using a software developed for both the computer and the PS part of the Zynq. Similar to [2,8], a thermal chamber was used for the calibration process.

To obtain the calibration curve of each sensor, the temperature inside the thermal chamber was increased from 30° C to 65° C in steps of 5° C every 10 min. 120 samples were recorded and averaged for each step. The temperature was monitorized by the internal analog temperature sensor and by a K-thermocouple stuck on the surface of the Zynq. In order to obtain the same conditions for both sensors, the digital one was instantiated as close as possible to the XADC sensor. Given a $\pm 4°C$ accuracy of the built-in sensor, the calibration of our digital sensors could not be studied more precisely. However, as explained in Section 1, this does not really matter for the targeted application.

4. Flip-Flop Metastability architecture (FFM)

The first designed sensor used the metastable effect. The architecture is given in Fig. 2. To detect a metastable event, the design
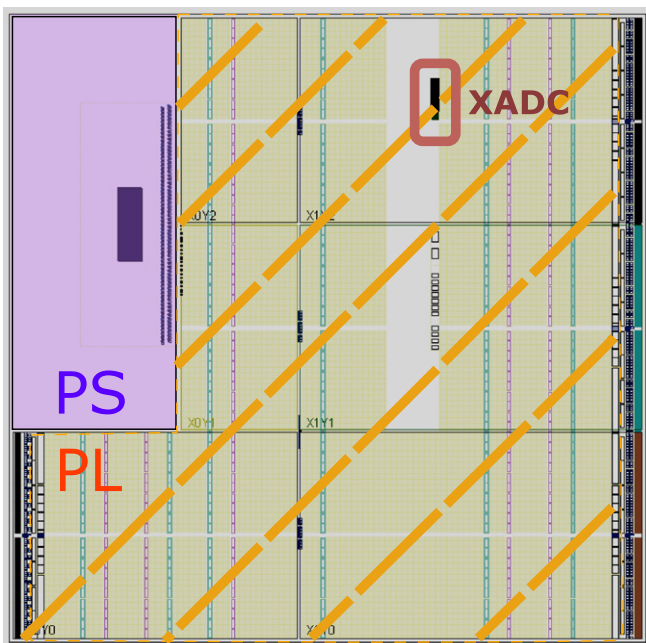
Fig. 2. Architecture of a Flip-Flop metastability based sensor.
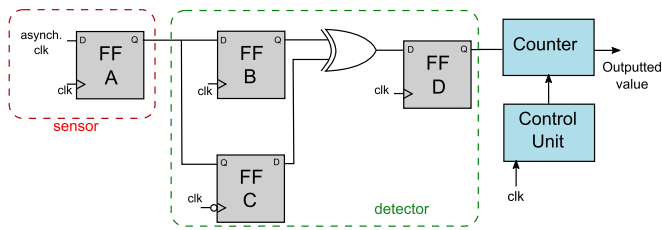

Fig. 4. Calibration curve of the FFM sensor.

first proposed by Xilinx [9] was used. The flip-flop A is the sensor and B, C, D constitute the detector. At each rising edge of clock, the asynchronous input is captured by A. Its output Q_A is then captured by the flip-flop C during the falling edge. On the next rising edge, Q_A is copied by the flip-flop B. If A is in a metastable state, its clock-to-q delay will be increased so C will not catch the event, whereas B will. By comparing the result of the two flip-flops with a XOR gate and synchronizing it with D, the metastable event is detected. As the clock-to-q delay depends on the temperature the number of metastable events during a fixed lap of time T_m depends on the temperature. This measurement time T_m allowed to configure the sensibility of the sensor.

4.1. Adjustments

The influence of the measurement time can be noticed in Fig. 3: the faster the sensor, the higher the error. The standard deviation of the ARO (Asynchronous Ring Oscillator) and SRO (Synchronous Ring Oscillator) sensors is discussed in the next Sections. The FFM-based sensor seems to require a higher measurement time (2^{24} clock ticks, i.e. 1 s) to reach the same precision. This measurement time was used for the experiments.

4.2. Calibration

Both RO-based and FFM-based sensors depend on multiple physical factors (voltage, device's silicon, place and route in the PL) and therefore need to be calibrated. For the FFM-based sensor, the results showed some unexpected behavior with previous experiments from [5], as illustrated in Fig. 4. Contrary to [5] where the expression is logarithmic, the curve obtained is lineal and the slope is negative, i.e. the higher the temperature, the smaller the value outputted by the counter. The linear approximation was drawn and the good determination coefficient R^2 proves the linearity of the expression. The experiment was repeated many times and the same tendency was observed. The architecture of the sensor being the
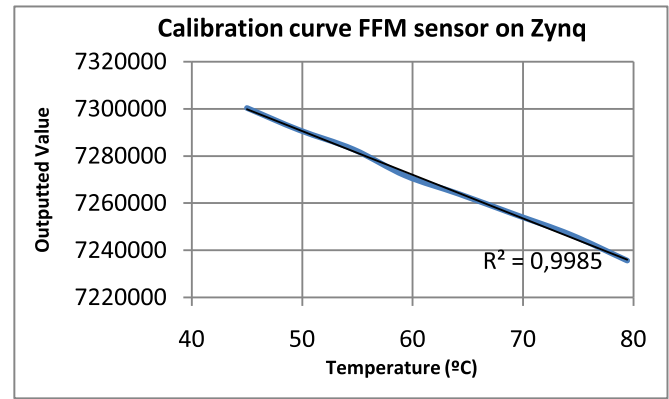
same, these differences may be addressed to the use of a different technology.

4.3. Conclusion

These particularities make the FFM-based sensor a good candidate to measure the temperature inside the Zynq. This sensor has a light hardware impact: the sensor element only requires 4 flip-flops and one logic gate. However, the minimum measurement time is quite high (1 s), and is not suitable for the chip protection application proposed in this paper.

5. Asynchronous Ring Oscillator architecture (ARO)

The second designed sensor was based on the most common architecture encountered in the literature. It consists in a counter sourced on the clock input by a controlled RO. The architecture used in this paper is presented in Fig. 5. The RO was designed in VHDL by looping an odd number of NOT gates and positioning one gate per LUT. As explained in Section 2, the frequency of the RO depends on the temperature. To obtain it, the counter is incremented during a fixed delay adjusted by the control unit. Therefore, the value of the counter is an image of the RO frequency.

As there is no synchronization between the counter and the main clock "clk", this sensor was called Asynchronous Ring Oscillator (ARO). Although this sensor was commonly encountered, this architecture gave strange behaviors with the technology used and the calibration step had to be repeated several time. This could be addressed to the counter which is directly clocked by the RO on its clock input. Feeding a clocked component (i.e. the flip-flops in the counter) on its clock input with an asynchronous element is strongly discouraged by the synthesis tools due to the time violations which may occur. However, to lead a comparative study with past publications, this sensor was tested.
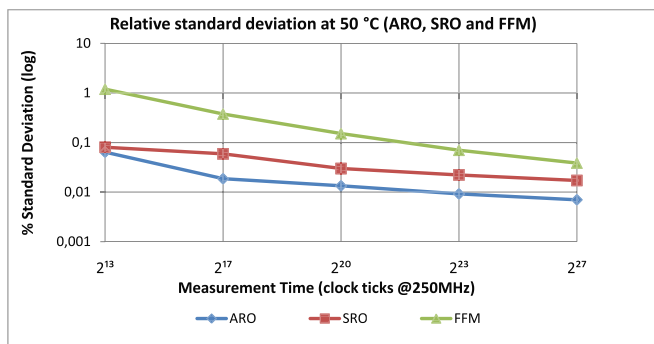

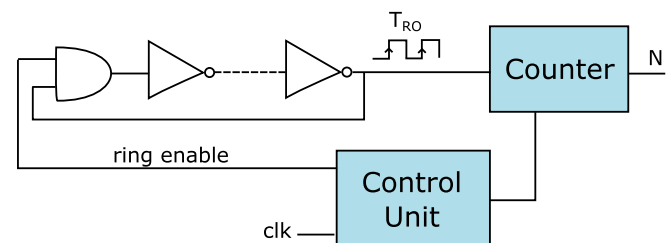**Fig. 3.** Standard deviation against the measurement time for each sensor.

Fig. 5. Architecture of an Asynchronous Ring Oscillator sensor.

5.1. Calibration

The calibration curve of the ARO sensor is given in Fig. 6. Even though the value of the counter decreased with the temperature as expected in Section 2, the curve obtained is not as good as in the past publications and may lead to very inaccurate temperature measurements.

As the Zynq is in a 28 nm technology and past experiments were led in 40 nm(Xilinx Virtex 5/6) [8,10] or in 90 nm(Altera Cyclone II), this difference may be ascribed to the technology. To study this hypothesis and the behavior of the ARO sensor on an old technology, the design was ported on a Xilinx Virtex-5 (chip XC5VFX30T on an Avnet board aes-v5fxt-evl30-g). This chip was powered under the same voltage as the Zynq (1.0 V) but used a 65 nm technology. The design was re-generated by an old tool suite (ISE/XPS 14.7) for both the Zynq-7000 and the Virtex-5 to obtain a viable comparison. The calibration curve obtained on the Virtex-5 is shown in Fig. 7.

5.2. Conclusion

The trend of this curve is clearly different from what was obtained with the Zynq-7000 in Fig. 6. Moreover, no problem were observed during the calibration step with the Virtex-5. The architecture presented here might not be reliable enough when being implemented on the newest die technologies. The ARO sensor was therefore discarded.

6. Synchronous Ring Oscillator architecture (SRO)

The architecture presented in Fig. 8 was introduced in our past paper [1]. In this architecture, the frequency of the RO is not directly measured unlike [2,11]. The signal from the RO is detected and synchronized by a rising-edge detector, rather than directly sent to the counter. This design supposes that "clk" is the sampling clock of the system and thereby respects the Shannon's condition on frequencies (2):

$$f_{clk} \geq 2f_{RO} \tag{2}$$

where f_{clk} is the frequency of the main clock and f_{RO} is the frequency at the output of the RO. The minimum number of NOT gates in the ring is therefore constrained by Eq. (2).

Two design factors influence the performance and the precision of this type of sensors: the number of gates in the ring and the measurement time adjusted by the control unit. These parameters were made adjustable and configurable and different sensor configurations were tested to observe the effects. The number of gates was implemented
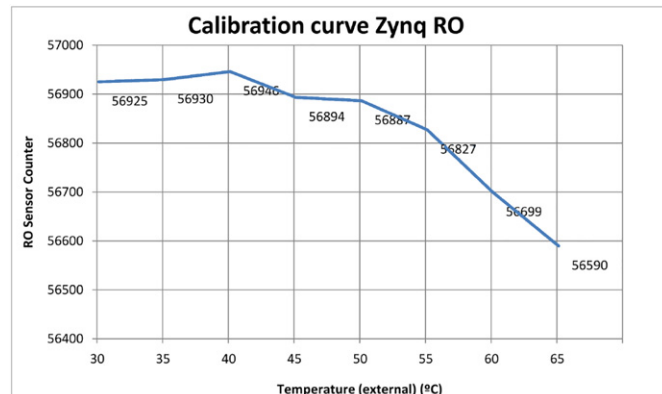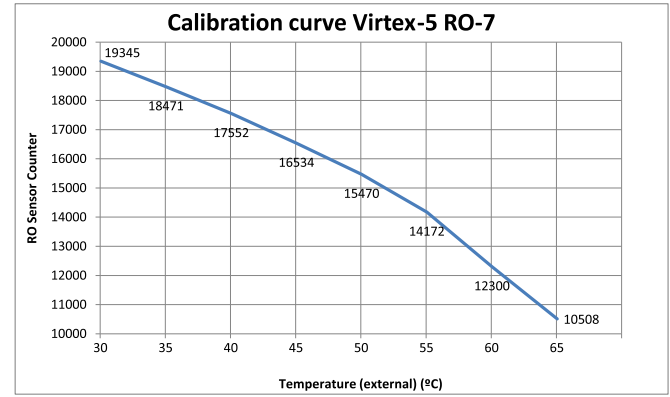
Fig. 7. Calibration curve of the ARO sensor implemented on a Virtex-5.

as a "generic" type and therefore was configured when instantiating the IP. As for the measurement time, it was configured through a register.

6.1. Adjustments

The first test consisted of obtaining the ideal number of gates in the SRO ring. The sensor was configured with 5 to 45 NOT gates-based sensors for different measurement times T_m, from 2^{17} (557 us) to 2^{27} (570 ms). The board was placed in the thermal chamber regulated at 50° C and 30 points were recorded, each one separated by 10s. The standard deviation against the measurement time was then computed in Fig. 9. All the measurement times T_m were chosen to be in ratio to the thermal propagation time of a fault across the silicon. For the silicon, the thermal diffusivity, i.e. its ability to transfer the heat from a point to another, is $a = 88mm^2/s = 0.088mm^2/ms$. Therefore, in 10 ms, nearly $1mm^2$ is affected by the temperature. Given the size of the die of the Zynq 7z020, a 10 mm-side square, $1mm^2$ -i.e. 1%- of the chip would be affected in 10 ms. The results are presented in Fig. 9.

As shown on Fig. 9 a high deviation was observed for the SRO sensor between 5 and 17 gates due to the non-respect of the Shannon's condition (2) in this interval, and then tended to stabilize with at least 23 NOT gates. Therefore, a 23-gate-based oscillator was used for the following experiments. Getting a fast sensor is compulsory for a temperature-safe application and that is why the smallest possible measurement time was chosen: 557 μs (2^{17} clock ticks). According to Fig. 3 and Fig. 9, this time does not trade a lot a precision and offers a best reaction time which perfectly fitted the requirements.

6.2. Calibration

The problem observed with the ARO version seemed to be corrected with the SRO version of the sensor, as shown by its calibration curve given in Fig. 10. The calibration process was led in the same conditions.
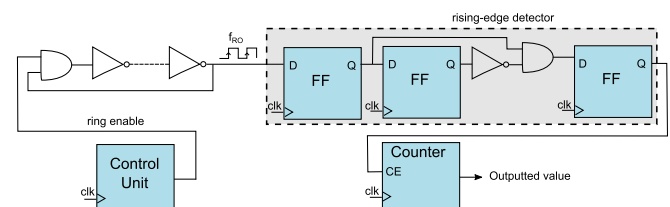
Fig. 6. Calibration curve of the ARO sensor.

Fig. 8. Architecture of a Synchronous Ring Oscillator sensor.

Fig. 9. Relative standard deviation against the measurement time for each sensor.

Fig. 11. 1-stage XOR-based ring oscillator.

6.3. Conclusion

As expected, the output of the counter (representing the frequency) decreased with the temperature. Despite the non-linearity of the curve owing to the low core voltage of the Zynq used (1 V) [8,12], this result is in line with [3]. The SRO calibration curve can be exploited with a polynomial approximation. This approximation was used to obtain the maximum error of this sensor: $\pm 5 °C$. This sensor complied with the requirements and was selected.

7. Application to protection: failure insertion and diagnostic

Until now, the stress was put on measuring the temperature in an absolute way. This section focuses on detecting if the temperature is above or below a threshold, on a relative manner. In Section 6, the dependence in temperature of the output of the SRO sensor was highlighted. In this section, using a SRO digital sensor on a Zynq 7z020, a way to detect a temperature failure in a precise area part of the Zynq PL is studied.

Different causes were identified as responsible for temperature overheating inside a FPGA, such as a short-circuit, a strong computation load, etc. The function affected may be a critical function (e.g. regulator in an aircraft). Such a failure was modeled by a digital heater. Its working is explained and the simulation cases are then presented.

7.1. Failure insertions

To simulate a local heat failure, a digital heater combining 100 switching LUTs and 4000 FFs was designed. Each switching LUT
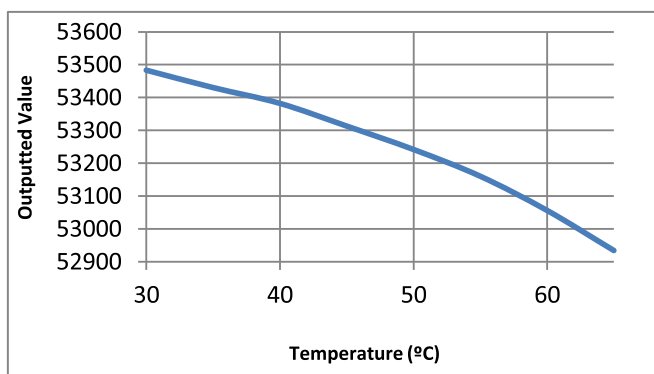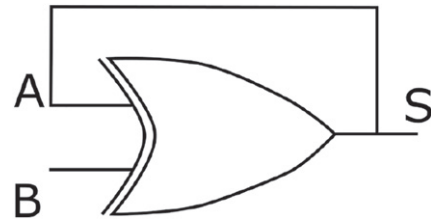
consists in a 1-gate RO based on a single XOR gate, as illustrated in Fig. 11. Its working is based on the self-heating phenomenon of a RO and the heat produced by switching a FF. This version produced a 1° C instant heat when activated and kept the voltage stable. The digital heater was encapsulated in an AXI-Lite IP.

A previous version was designed by instantiating 2000 switching LUTs. In less than 10min, the maximum temperature for the commercial version of the chip was reached (85° C). The heating curve is shown in Fig. 12.

However, this type of design led to significant voltage drops ($-20mV$) which affect the frequency of the RO, and the temperature effect could not been separated from the voltage effect.

7.2. Failure detection

Two test cases (scenarios 1 and 2) aimed at studying the detection of an internal failure, in the die of the chip.

7.2.1. Scenario 1

Both the heater and the SRO sensor were placed near the XADC. This way the detection speed of the sensor and its ability to protect an area of the chip were studied. The SRO and the built-in sensors were read every 10 ms through an UART connection. Some samples were recorded in a normal state and then the digital heater was activated to simulate a local failure. During all the experiment, the internal voltage of the PL part was monitored to ensure the absence of any voltage drop when the heater was activated. As its value remained stable (993 mV), the influence of the voltage was not significant. The results are exposed in Fig. 13.

The increase of the temperature was detected by the SRO-23 sensor when the heater was activated at t = 200 (i.e. 2 s). The raw output from the sensor was then analyzed to obtain a clearer view. The moving average on 15 samples was computed and plotted on the same figure. This type of transformation is a convolution of the signal with a N-width gate (with N = 15) and give a smoother curve. By
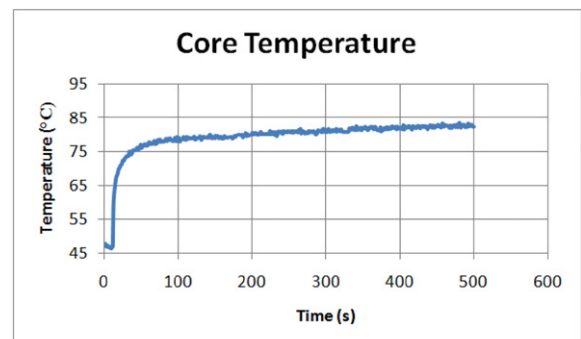

Fig. 10. Calibration curve of the SRO sensor.

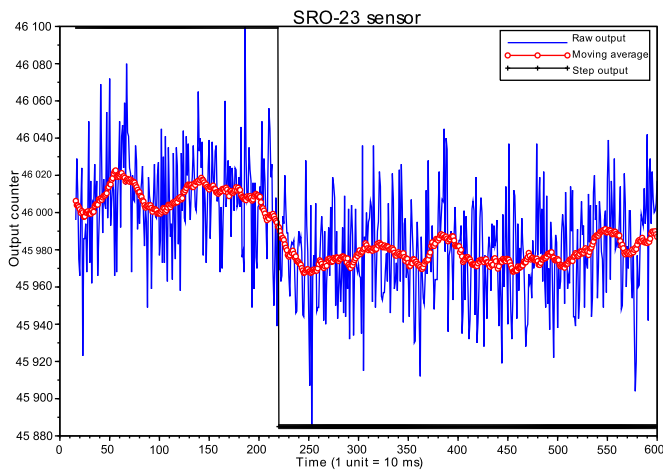Fig. 12. Heating curve with the designed digital heater.

Fig. 13. SRO-23 output in temperature detection and post-treatment.

defining a threshold value, a step is produced when the temperature changes. Even though this operation was computed off-line (after the experiment), this could be as well implemented in hardware: the time needed to compute the moving average on 15 samples with the defined SRO sensor in Section 6 which can measure a point each 600 μs would be $0.6 * 15 = 9ms$ which is completely in the range of 10 ms due to thermal propagation, exposed in the same section. Therefore, this transformation does not prevent it from protecting an area.

7.2.2. Scenario 2

The second scenario aimed at studying the ability of the SRO-23 sensor to differentiate two spaced thermal failures across the die. Two sensors were instantiated: one near the heater and the other further, as shown in Fig. 14. The heater was switched on at t = 250 (250 ms) The digital sensors were read every ms and the same mathematical processing was done with the same parameters for both sensors. The curves are given in Figs. 15 and 16.
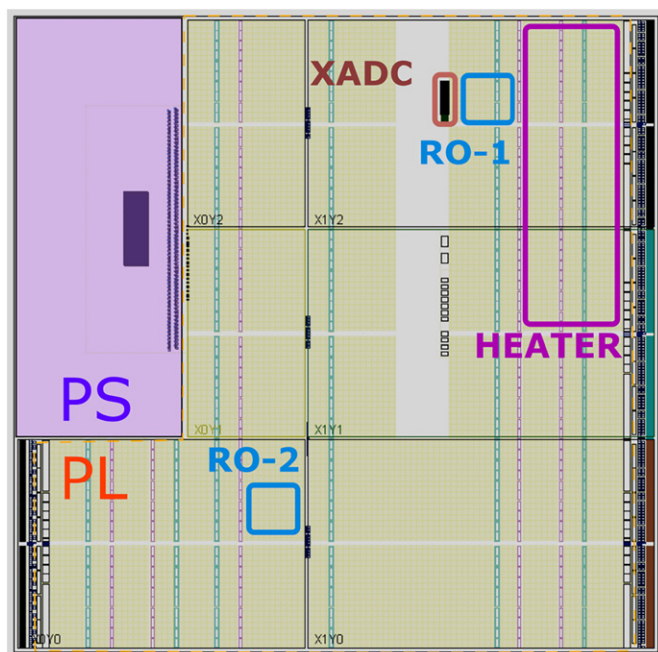

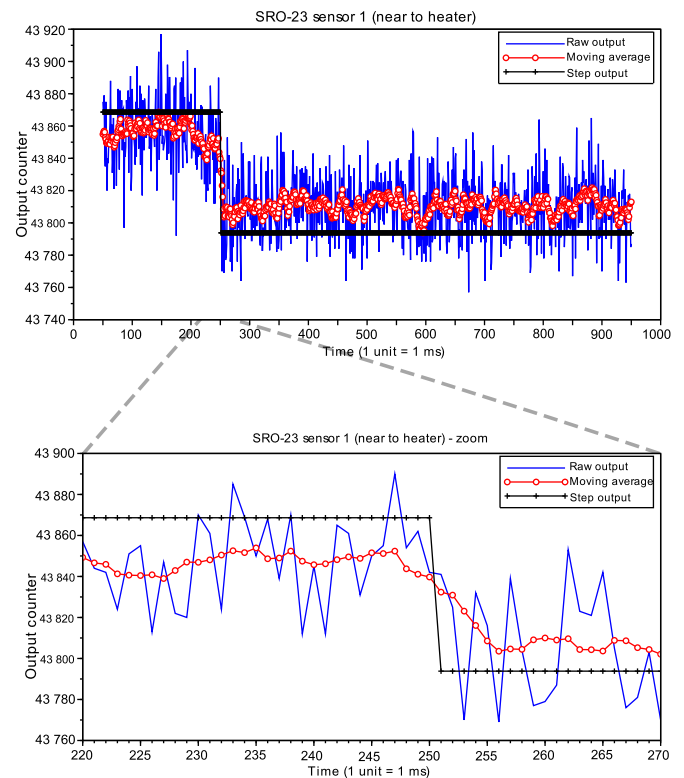**Fig. 14.** Placement of the two SRO-23 sensors and the heater.

Fig. 15. SRO-23 sensor near to the digital heater.

In the case "close to the heater", Fig. 15, the drop occurred at t = 251 (251 ms), whereas in the case "far from the heater" in Fig. 16 it did not occur because the raw signal was not clear enough to be
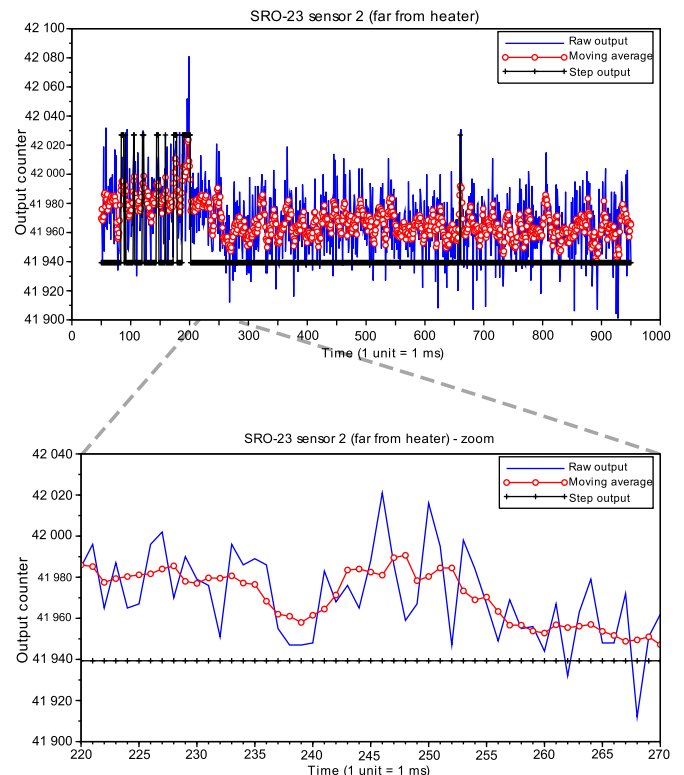

Fig. 16. SRO-23 sensor far from the digital heater.

detected as a thermal failure. Only artifacts were observed, before the heater was activated at t = 250. This means the sensibility of the sensor is reduced enough and it is worth considering a system where a sensor is instantiated for each area to protect.

8. Conclusions

This work presented the implementation of 3 different temperature digital sensors on a Xilinx Zynq 7z020. Differences were observed between the ARO behavior on old technologies and on the Zynq. These disparities were attributed to the 28 nm-technology used in the Zynq. Therefore the ARO sensor was dismissed. The FFM sensor was identified as a good candidate to measure the temperature thanks to its linear calibration curve. However the required measurement time of 1 s is too high to achieve a fast temperature sensing. For a thermal protection application, the SRO offers more abilities. If it is calibrated, a $\pm 5°C$ precision was measured using a 23-gate RO and a $557\mu s$ measurement time.

Besides, this sensor can be used for temperature detection with no calibration at all. This ability was proved by simulating a temperature failure inside the Zynq and detecting it. A signal processing routine is required to clean the raw signal coming from the sensor. This lasts less than 10 ms which is, according to the thermal diffusivity inside the silicon, the time taken by a thermal failure to affect 1% of the chip. The detection is also discriminative: due to the physical impact of temperature on ring oscillators, the signal observed depends on the strength of the failure. Finally, the impact on the hardware (56 LUTs and 59 flip-flop, i.e. 0.1% of the available resources) is low enough to consider implementing such a sensor in every critical area to protect.

Acknowledgments

References

[1] C.-A. Lefebvre, L. Rubio, J.L. Montero, Digital thermal sensor based on ring-oscillators in Zynq SoC technology, 2016 22nd International Workshop on Thermal Investigations of ICs and Systems (THERMINIC), IEEE. 2016, pp. 276–278. http://dx.doi.org/10.1109/THERMINIC.2016.7749065.

[2] S. Lopez-Buedo, J. Garrido, E.I. Boemo, Dynamically inserting, operating, and eliminating thermal sensors of FPGA-based systems, IEEE Trans. Compon. Packag. Technol. 25 (4) (2002) 561–566. http://dx.doi.org/10.1109/TCAPT.2002.808011.

[3] C. Ruething, A. Agne, M. Happe, C. Plessl, Exploration of ring oscillator design space for temperature measurements on FPGAs, Proceedings - 22nd International Conference on Field Programmable Logic and Applications, FPL 2012, IEEE. 2012, pp. 559–562. arXiv:1011.1669v3. http://dx.doi.org/10.1109/FPL.2012.6339370.

[4] P. Chen, M.C. Shie, Z.Y. Zheng, Z.F. Zheng, C.Y. Chu, A fully digital time-domain smart temperature sensor realized with 140 FPGA logic elements, IEEE Trans. Circuits Syst. I, Reg. Papers 54 (12) (2007) 2661–2668. http://dx.doi.org/10.1109/TCSI.2007.906073.

[5] G. Tarawneh, M. Terrence, A. Yakovlev, Intra-chip physical parameter sensor for FPGAS using Flip-Flop metastability School of Electrical and Electronic Engineering Newcastle University, Newcastle upon Tyne, NE1 7RU, United Kingdom Department of Computer Science and Engineering The Chinese U, 22nd International Conference on Field Programmable Logic and Applications (FPL), 2012, pp. 373–379. http://dx.doi.org/10.1109/FPL.2012.6339207.

[6] B. Plds, C. Plds, D.U. Test, ispLSI / GAL Metastability Report, October 2001, 1–14.

[7] I.W. Jones, S. Yang, M. Greenstreet, Synchronizer behavior and analysis, Proc. IEEE Int. Symp. Circuits Syst. (2009) 117–126. http://dx.doi.org/10.1109/ASYNC.2009.8.

[8] J.J.L. Franco, E. Boemo, E. Castillo, L. Parrilla, Ring oscillators as thermal sensors in FPGAS: experiments in low voltage, 6th Southern Programmable Logic Conference, SPL 2010 - Proceedings, IEEE. 2010, pp. 133–137. http://dx.doi.org/10.1109/SPL.2010.5483027.

[9] P. Alfke, B. Philofsky, Xilinx Application Note XAPP 094: Metastable Recovery, 1997.

[10] P. Weber, M. Zagrabski, B. Wojciechowski, K.S. Berezowski, M. Nikodem, K. Kepa, Toolset for measuring thermal behavior of FPGA devices, 19th International Workshop on Thermal Investigations of ICs and Systems (THERMINIC) 2013, 2013, pp. 48–53. http://dx.doi.org/10.1109/THERMINIC.2013.6675251.

[11] S. Xie, W.T. Ng, Delay-line temperature sensors and VLSI thermal management demonstrated on a 60nm FPGA, Proc. IEEE Int. Symp. Circuits Syst. (2014) 2571–2574. http://dx.doi.org/10.1109/ISCAS.2014.6865698.

[12] M. Happe, A. Agne, C. Plessl, Measuring and predicting temperature distributions on FPGAs at run-time, Proceedings - 2011 International Conference on Reconfigurable Computing and FPGAs, ReConFig 2011 (2011) 55–60. http://dx.doi.org/10.1109/ReConFig.2011.59.