

POWER CHARACTERISATION FOR THE FABRIC IN FINE-GRAIN RECONFIGURABLE ARCHITECTURES

Tobias Becker, Peter Jamieson, Wayne Luk

Department of Computing
Imperial College London

Peter Y. K. Cheung

Department of EEE
Imperial College London

Tero Rissa

Nokia Devices R&D
Finland

ABSTRACT

This paper proposes a methodology for characterising power consumption of the fine-grain fabric in reconfigurable architectures. It covers active and inactive power as well as advanced low-power modes. A method based on random number generators is adopted for comparing activity modes. We illustrate our approach using four field-programmable gate arrays (FPGAs) that span a range of process technologies: Virtex-II Pro, Spartan-3E, Spartan-3AN, and Virtex-5. We find that, despite improvements through process technology and low-power modes, current devices need further improvements to be sufficiently power efficient for mobile applications.

1. INTRODUCTION

Rapidly evolving standards, convergence of increasingly complex features and growing time to market pressure are pushing manufacturers of mobile consumer devices to consider alternatives to ASICs and microprocessors. There is a clear demand for power efficient circuits that are flexible, while capable of delivering performance through parallelism. Reconfigurable architectures, such as FPGAs, have good potential in meeting the demand for flexibility and performance, but they often miss the power requirements by up to several orders of magnitude. They consume significant active power [1] and have limited low-power capabilities when they are inactive. Moreover, such devices can also cause thermal problems when they heat up during intense processing. A further concern is the expected increase of static power in future process technologies. Because of the strong temperature dependency of static power this can lead to thermal runaway conditions.

One of the problems with optimising for low power in FPGAs is that a designer has to map a particular application onto a range of target devices and evaluate their power consumption by using the power estimation tools provided with some FPGA CAD flows. But these estimation tools have often limited accuracy. Moreover, synthesising, implementing and simulating a design on a range of devices can be very

time consuming. The result of an evaluation is only meaningful for this particular design. Instead, it would be desirable to have a set of test cases that can be used to benchmark the power characteristics of reconfigurable devices. A list of test results would allow users to choose suitable devices for a low-power design and the procedure can also be used to evaluate architectural improvements.

Evaluating the power efficiency can be achieved with application specific or application independent test scenarios. In this paper, we focus on an application independent method that allows a fast and easy evaluation of the power consumption in fine-grain FPGAs. The power is measured in several activity modes which reflect realistic high-power or low-power scenarios a device could be operated in. Another important aspect is thermal characteristics such as the temperature dependency of static power and heating up of the device under different processing scenarios. The proposed methodology is intended to provide a fair comparison of existing devices, and should also be able to capture the improvements in future devices with new low-power techniques. This methodology is part of a larger power benchmarking framework that provides further application specific test cases which are representative for mobile devices.

The remainder of this paper is organised as follows. Section 2 describes previous work. Section 3 proposes a method for power characterisation of fine-grain reconfigurable fabrics. Section 4 describes how this method can be implemented on commercial FPGAs and section 5 shows some results of our power characterisation. Finally, section 6 concludes the paper.

2. BACKGROUND

One of the earliest comparative studies for power consumption on FPGAs is done by George *et. al.* [2]. The authors create a low-power FPGA through architecture and low-level circuit design, and compare their FPGA to Xilinx and Altera devices. The comparison is based on three test circuits that are evaluated with Synopsis Powermill. The three circuits consist of a single flip-flop driving 9 routing segments,

a 1K array of 16 bit counters, and a toggle circuit.

Shang *et al.* [3] measure the dynamic power consumption of a Xilinx Virtex-II FPGA [4] using one Xilinx internal benchmark that represents a large industrial circuit. Using this internal benchmark and input stimuli, they calculate the switching activity of the design. They estimate power based on the calculated switching activity and the effective capacitance of each resource on the FPGA. This is possible since they have access to low-level models of the FPGA, but such models are usually not accessible.

Gayasen *et al.* propose an FPGA architecture with two supply voltages where the lower voltage is used for all non-critical path components [5]. The efficiency of their architecture is evaluated with MCNC benchmarks which provide a range of simple circuits and state machines.

Recently, Kuon *et al.* [1] assess the gap between FPGAs and ASICs. This work includes an attempt to measure the dynamic and static power consumption gap between the two technologies. They estimate the static and dynamic power consumption of an FPGA using the power estimation tools provided by the FPGA vendor, and use either included test-benches or estimates of net activity.

FPGA fabrics have also been characterised in terms of their thermal characteristics and die variation. Lopez-Buedo *et al.* [6] use ring oscillators programmed onto an FPGA. These oscillators are placed around an existing mapped design to measure the temperature of the fabric when the design is operating. The local temperature is determined by measuring the frequency of the ring oscillator. In relation to this work, Sedcole *et al.* [7] used the ring oscillator concept to study within-die delay variation.

Most current commercial FPGAs have limited low-power capabilities. Traditionally, the only methods of saving power are to employ clock gating or to turn off the entire device. The first method has limited potential whereas in the latter case, state and configuration are lost. More modern FPGAs have additional low-power capabilities. Xilinx Spartan-3A and Spartan-3AN FPGAs support a suspend mode in which the clock tree can be stopped and powered down [8]. Other examples are Actel IGLOO [9] and Silicon Blue iCE [10] FPGA families, both targeting the portable device market. At the moment however, it is unclear how these advancements can be compared and measured.

3. FABRIC CHARACTERISATION METHOD

When developing a general characterisation method one faces a number of challenges. The method should:

- Be applicable to a wide range of devices.
- Be a fair comparison and results free from implementation tool influences or hand-optimisations.
- Allow to capture different power modes and possible future techniques that are currently not available.

The basic idea of our method is to implement a highly active circuit on the FPGA and measure the power in several activity modes. These activity modes represent active processing, inactivity and dedicated low-power states. This method allows us to characterise power consumption quickly in best-case case and worst-case scenarios and it outlines the suitability of a device for low-power design. Furthermore, we consider the thermal aspects of the device such as heating up under active processing.

The benefit of this characterisation is that it allows us to assess the adequacy of a device for a low-power design. Moreover, it can be used to compare and optimise devices for lower power. The key aspects of our method can be summarised as follows:

- Use random number generators (RNGs) as test circuit with high activity.
- Use 90% of the logic resources in the device.
- Run the test circuit at a fixed clock rate of 100MHz when active.
- Specify the behaviour of activity modes and switch between these with various duty cycles.
- Measure power and temperature in these modes.

To create a worst case active processing scenario, we use pseudo random number generators as test circuits [11]. These random number generators are based on binary linear recurrences where each bit of the next state is generated based on a linear combination of the current state. Compared to linear feedback shift registers (LFSR), the most common type of random number generators, this improves quality of the random numbers. It also results in a circuit where LUTs are heavily interconnected and placement cannot be optimised. This is beneficial for our purposes because the circuit will exercise all different kinds of short and long wires of the routing fabric. It also does not provide any opportunity for logic optimisation or optimised placement and routing. The lack of optimisation potential is an important aspect to reduce the influence of the implementation tools on the result of our characterisation. The random number generator circuit is also characterised by a high and uniformly distributed toggle rate, and therefore suitable to act as worst case scenario of maximum activity in the fine grain fabric. The statistical chance of toggling on the rising clock edge is 50% for all flip-flops in the circuit. Hence, the total toggle rate of the circuit is 50%.

Currently, we use a 512-bit random number generator core that maps to exactly 512 LUTs and 512 flip-flops. Since current FPGAs provide tens to hundreds of thousands of LUTs and flip-flops, we scale the size of the test circuit with the size of the device. To achieve high logic utilisation while still allowing routability, we implement multiple instances of the random number generators so that 90% of

all logic resources are used. The resulting power consumption is normalised to the number of LUTs in order to allow a comparison of differently sized chips.

The cores are driven by a 100 MHz clock when the circuit is active. This frequency simply acts as a reference point for a typical FPGA clock frequency. The power characteristics for different clock frequencies can be estimated by scaling the power consumption linearly to the clock rate.

To enable a comparison of devices with different low-power capabilities, we define the behaviour of activity modes. These activity modes specify how the device behaves in a certain mode rather than by which means this mode is implemented. The two basic modes that are applicable to all devices are *active* and *inactive*. In *active* mode, the test circuit continuously generates random numbers at 100 MHz clock frequency. The power consumption in this mode is a combination of static and high dynamic power. In *inactive* mode, the circuit does not generate random numbers. However, its state is preserved and it can be instantly brought back into active mode. These are the most basic activity modes and a transition between these two modes can usually be implemented with a simple clock gating approach as illustrated in figure 1. However, we are only concerned about the power profile of the inactive device with preservation of state and instant wake-up capability, and not the details of its technical implementation. This mode corresponds to static power only, if clock gating is chosen as a method to implement this mode. However, depending on the implementation details, there might be supporting circuitry such as clock managers that are still operating and drawing dynamic power. It is important to point out that in this context, we are not necessarily interested in measuring pure static power but rather the minimal power to implement the *inactive* mode.

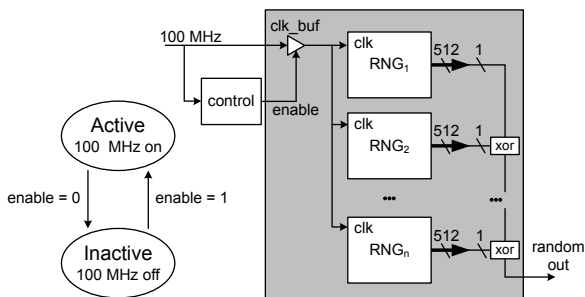


Fig. 1. Fabric characterisation circuit implemented with random number generators.

In order to compare devices with advanced low-power modes, we characterise their behaviour and compare the power consumption based on the two basic modes. Table 1 illustrates an example with our two basic activity modes *active* and *inactive*, and two hypothetical advanced modes *sleep* and *hibernate*. The behaviour of the basic modes is

	activity mode			
	standard		device-specific	
	active	inactive	sleep	hibernate
generate output	yes	no	no	no
retain state	-	yes	yes	no
wakeup time	-	instant	500 μ s	50ms

Table 1. Examples of activity modes. The first two modes are fixed, further modes can be defined based on the device capabilities

fixed, while the behaviour of the advanced modes depends on the device capabilities.

To enable a simple comparison of devices, we measure the power consumption in each activity mode. Since the static power depends on temperature, all measurements of low-power modes are taken at 25°C. In a more detailed analysis, we characterise the heating up of the device based on its activity.

Even though mobile applications are characterised by strict low-power constraints, thermal considerations are also important. In addition to power constraints, devices also have a thermal budget. Hence, we want to analyse how a device heats up for a given amount of activity.

Thermal characterisation of a device usually involves determining the thermal resistance of the device between die (also called junction) and case or junction and ambient air. For this purpose, a special test die is usually mounted in a case and heated up to a defined junction temperature T_j by applying power P . The junction-to-case thermal resistance θ_{jc} can then be calculated based on the following equation:

$$\theta_{jc} = \frac{T_j - T_c}{P} \quad (1)$$

Equation 1 can also be used to calculate the case temperature T_c if θ_{jc} is specified by the device manufacturer. This however requires the knowledge of the junction temperature T_j . A further disadvantage is that equation 1 is not very accurate since it does not consider heatflow into the PC board or the ambient air. As an alternative, we propose to measure the case temperature as well as power consumption in a well-defined environment over the device activity. The activity in the device is adjusted by periodically switching the device between active and inactive state with various duty cycles. For each duty cycle we measure active power, inactive power and temperature. To take these measurements, we first wait until the temperature has converged to its final value as illustrated in figure 2. We record this temperature and then take a reading of active and inactive power. As mentioned earlier, inactive power does not necessarily have

to be equivalent to static power and can have dynamic components as well. Nonetheless, we expect a strong temperature dependency on inactive power because of its close relation to static power. Likewise, we record the active power that we expect to be less temperature dependent, since it is largely based on dynamic power. These measurements are taken for duty cycles from 0% to 100% with 5% increments.




Fig. 2. Depiction of power measurement sample points relative to temperature

In order to reduce other influences on the measurement, a well-defined environment is used. We specify the test environment as a chip mounted on a PC board surrounded by ambient air with a temperature of 25°C. The board is placed in a large cardboard enclosure which is supposed to reduce airflow to natural convection and reflect infrared radiation.

4. IMPLEMENTATION OF THE FABRIC CHARACTERISATION METHOD

We implement our fabric characterisation in Xilinx Virtex and Spartan series FPGAs. The synthesis and mapping of the random number generator circuit for these devices is straightforward. Xilinx Virtex series FPGAs do not support any low-power modes other than simply stopping the clock. The most efficient way of setting the device into inactive mode as specified in section 3 is to disable the clock tree using an internal clock buffer. This buffer is located at the root of the clock tree and therefore reduces the dynamic component of inactive power to a minimum. The enable signal of the clock buffer is connected to an external pulse generator with variable duty cycle. Spartan-3A and Spartan-3AN FPGAs also feature an advanced low-power mode called *suspend* [12]. This mode is controlled by an external signal pin and does not require modifications to the logic design itself.

Figure 1 illustrates the implementation of our test circuit. The RNG cores only have a clock input and a 512-bit wide output. Each core is initialised with a different seed and one output pin of each core is connected via an XOR chain to avoid logic optimisation of the circuit. The reported LUT usage of the implementation tools should match the expected value of $512 \cdot n$, where n is the number of instantiated RNGs.

The power consumption of an FPGA can be obtained by measuring the current on the supply rails. However, it

should be considered that FPGAs are sensitive to core voltage variations and usually have to stay within 50 mV of the nominal value. The voltage drop caused by the shunt resistor of a current meter can exceed this value if high currents are measured. As an alternative, we insert precision current sense resistors directly into the rail and measure the voltage drop over the resistor. The resistor value is selected so that the voltage drop at maximum current is less than 50 mV.

We use an Optris MiniSightPlus infrared thermometer to determine the surface temperature of the device. This contactless method minimises the influence of the measurement on the system. Infrared thermometers work very well on matt, black plastic cases. However, they are inaccurate on shiny metal surfaces since these emit less infrared radiation. We therefore apply a thin blackened aluminium plate to devices with a metallic surface.

5. RESULTS AND MEASUREMENTS

We perform a fabric characterisation as described in section 3 on four commercially available FPGAs. Our measurements cover for the following Xilinx devices and boards listed in table 2. The tested devices vary notably in their process technology and core voltage and thus have significantly different power profiles. Table 2 also shows the logic capacity of each device, the number of RNG cores used and the resulting device usage which should be as close as possible to 90%. Both Spartan devices have less than half of the logic capacity than the Virtex devices. The Spartan devices also have plastic packages which have higher thermal resistances. This leads to increased junction temperatures for a given surface temperature. All boards have current measurement facilities except the ML505 board which was specifically modified to allow equally precise measurements. The designs are implemented with Xilinx ISE 9.2.

Figure 3 shows the surface temperature of the devices over the duty cycle. The temperature is measured with an infrared thermometer as explained in section 4. For all devices the temperature increases almost linearly. Virtex-II Pro has the steepest slope and cannot be run beyond 60% duty cycle where the device reaches a surface temperature of 74°C. Based on the thermal resistance and the power consumption, we estimate that this surface temperature corresponds to the maximum allowed junction temperature of 85°C. Virtex-5 and Spartan-3 can be run at full duty cycles without overheating. These devices reach a final temperature of 54°C and 47°C respectively.

Figure 4 illustrates the inactive power consumption normalised per LUT over duty cycle. Each point on the line is measured when the temperature of the FPGA has stabilized for a given duty cycle as illustrated in Figure 2. We observe that inactive power consumption, which in our implementation is almost equivalent to static power, increases

FPGA Family	Device	Board	Process Technology	Core Voltage	Thermal Resistance Θ_{JC}	Number of LUTs/FFs	Number of RNGs	Percent Logic Utilisation
Virtex-II Pro [4]	V2P30	XUP	130nm	1.5 V	0.6 °C/W	27,392	48	89.7%
Spartan-3E [8]	S3E500	Spartan-3E Starter Kit	90nm	1.2 V	9.8 °C/W	9,312	16	88%
Spartan-3AN [8]	S3AN700	Spartan-3AN Starter Kit	90nm	1.2 V	5.3 °C/W	11,776	21	91.3%
Virtex-5 [13]	V5LX50T	ML505	65nm	1 V	0.2 °C/W	28,880	50	88.8%

Table 2. Details of FPGAs used for our fabric characterisation experiment.




Fig. 3. Variation of device surface temperature with duty cycle.

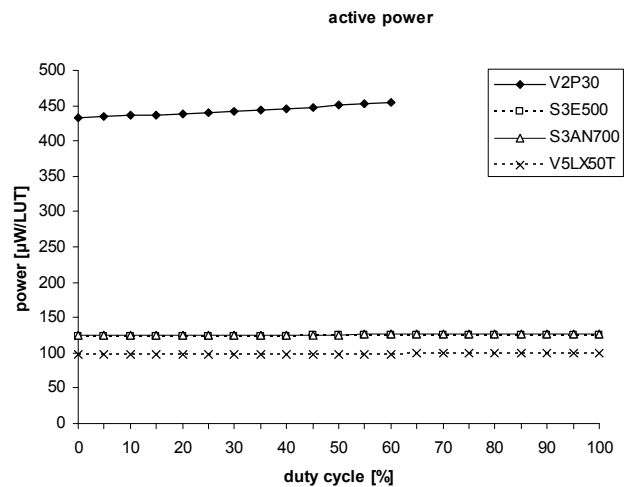


Fig. 5. Variation of instantaneous active power consumption with duty cycle.

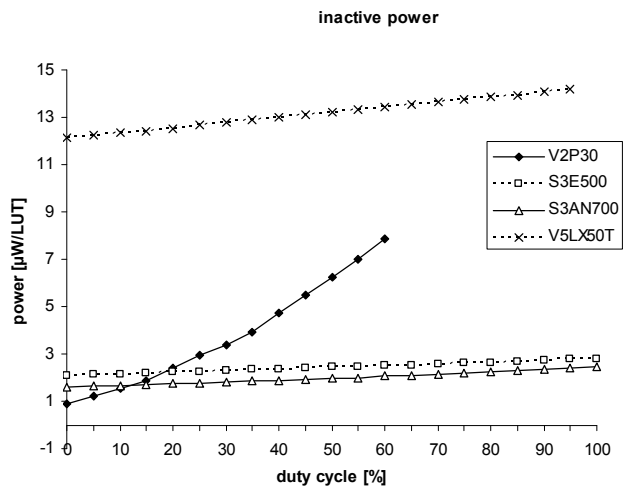


Fig. 4. Variation of inactive power consumption with duty cycle.

with the duty cycle. This is due to the rising temperature of the device caused by heat dissipated during the active part of the duty cycle. The Virtex-5 device, which is manufactured in 65 nm process, has a 12 times higher inactive

power consumption under cold conditions (0% duty cycle) than the 130 nm Virtex-II Pro. However, the inactive power in Virtex-II Pro deteriorates quickly with higher duty cycles because of the high device temperature as illustrated in figure 3. Virtex-5 and the Spartan-3 devices develop less heat which leads to a less progressive increase in inactive power. The Spartan-3 devices show the overall best efficiency during inactive phases.

Figure 5 shows the normalised active power consumption for all devices. This is the instantaneous active power during on-phase of the duty cycle as illustrated in figure 2. We can observe a notable improvement in active power for newer devices which is due to feature and voltage scaling in the process technology. The improvement from Virtex-II Pro (130 nm) to Spartan-3 (90 nm) is especially noteworthy. Compared to Virtex-5, the active power is reduced by more than a factor of 4. The active power consumption is relatively independent of duty cycle and temperature, although, this could change if static power becomes a more dominant component in active power. In current devices, we find that inactive power is considerably less than active power although the ratio increases from 0.2% to 14% between Virtex-II Pro and Virtex-5 in cold conditions, or from 1.5% to 16% in hot conditions.

The only device in our test that features an advanced low-power mode is Spartan-3AN that provides a suspend mode. This mode reduces the power consumption of all auxiliary circuits powered on the VCC_{aux} rail [12]. The logic state is preserved during suspend mode and the wake-up time ranges between 100 μs and 500 μs . Table 3 illustrates the core power consumption in all modes. Compared to the inactive mode, the suspend mode reduces the power consumption by a factor of 3. IO power is not listed because it is highly dependent on pin loads of a given design. IOs can usually be powered down if the device is not used and are not important for low-power modes.

In an overall comparison, Spartan-3AN device is the most power efficient device from our range of tested devices. On the other hand, a power consumption of 23.9mW in suspend mode is still unacceptably high for most power budgets in mobile applications. To bring reconfigurable technology into these devices, the effectiveness of low-power modes needs further improvement.

	active mode	inactive mode	suspend mode
P_{int} [mW]	1349	18.7	18.1
P_{aux} [mW]	44	43.6	5.8
P_{total} [mW]	1393	62.3	23.9

Table 3. Core power in a Xilinx Spartan-3AN 700 FPGA for active, inactive and suspend modes. All values are measured at 25°C.

6. CONCLUSIONS AND FUTURE WORK

In this paper, we provide a new, application independent methodology for the fabric characterisation of fine-grain FPGAs. This methodology is useful in evaluating the active and inactive power consumption as well as advanced low-power modes. We describe procedures for measuring active and inactive power and temperature on FPGAs using a simple experimental setup. The key to this setup is the use of random number generators as a highly active circuit.

To illustrate our methodology, we perform the fabric characterisation for four Xilinx FPGAs. Our measurements show how advances in process technology reduce the active power by more than a factor of 4. We also observe an increase of inactive power by up to one order of magnitude. However, modern devices generate less heat per activity and suffer less from temperature-based deterioration of inactive power. Additionally, we measure one specific low-power mode that reduces the inactive power by a factor of 3. These improvements are noteworthy, but advances through process technology alone are not enough to meet the strict power constraints in mobile devices. In particular, the power consumption during inactive periods needs to be addressed. To meet

mobile power requirements, future devices need flexible and more effective low-power modes.

Our proposed methodology is part of a power benchmarking framework and further application specific test cases that are representative of computations in mobile devices are currently being developed. Current and future work also includes studying a wider range of devices and characterising hardened IP blocks.

7. REFERENCES

- [1] I. Kuon and J. Rose, "Measuring the gap between FPGAs and ASICs," in *ACM/SIGDA 14th international symposium on Field programmable gate arrays*. ACM Press, 2006, pp. 21–30.
- [2] V. George, H. Zhang, and J. Rabaey, "The design of a low energy FPGA," in *Intl. symposium on Low power electronics and design*, 1999, pp. 188–193.
- [3] L. Shang, A. S. Kaviani, and K. Bathala, "Dynamic power consumption in Virtex-II FPGA family," in *ACM/SIGDA tenth international symposium on Field-programmable gate arrays*. ACM, 2002, pp. 157–164.
- [4] *Virtex-II Pro and Virtex-II Pro X Platform FPGAs: Complete Data Sheet*, Xilinx Inc., May 2007.
- [5] A. Gayasen, K. Lee, V. Narayanan, M. Kandemir, M. J. Irwin, and T. Tuan, "A dual-Vdd low power FPGA architecture," in *Field-Programmable Logic and its applications*. Springer, 2004, pp. 145–157.
- [6] S. Lopez-Buedo and E. Boemo, "A method for temperature measurement on reconfigurable systems," in *Proceedings of XII Design of Circuit and Integrated Systems Conference*, 1997, pp. 727–730.
- [7] N. P. Sedcole and P. Y. K. Cheung, "Within-die delay variability in 90nm fpgas and beyond," in *International Conference on Field Programmable Technology*, 2006, pp. 97–104.
- [8] *Spartan-3 Generation FPGA User Guide v.1.2*, Xilinx Inc., April 2007.
- [9] *Igloo Handbook*, Actel, January 2008.
- [10] *iCE65 Ultra Low-Power Programmable Logic Family Data Sheet*, SiliconBlue, May 2008.
- [11] D. B. Thomas and W. Luk, "High quality uniform random number generation using LUT optimised state-transition matrices," *VLSI Signal Processing*, vol. 47, no. 1, pp. 77–92, 2007.
- [12] *Using Suspend Mode in Spartan-3 Generation FPGA*, Xilinx Inc., May 2007.
- [13] *Virtex-5 Family Platform Overview LX and LXT Platforms v2.2*, Xilinx Inc., January 2007.