

Logic Depth and Power Consumption: A Comparative Study Between Standard Cells and FPGAs

E. Boemo and S. López-Buedo

*E.T.S. Informática, U.A.M.
Ctra. de Colmenar Km.15,
28049 Madrid, España.
<http://www.ii.uam.es>*

C. Santos Pérez

*SIDSA
PTM, c/ Torres Quevedo 1,
28760 Tres Cantos, España.
<http://www.sidsa.es>*

J. Jáuregui and J. Meneses

*E.T.S.I. Telecomunicación, U.P.M.
Ciudad Universitaria,
28040 Madrid, España.
<http://www.etsit.upm.es>*

Abstract.- *In large combinational datapaths, glitches generated in the first stages produce an avalanche effect in the activity of subsequent nodes. This problem can be avoided by registering every few gates. Thus, the propagation of glitches is blocked, and consequently, the datapath power is reduced. However, these extra registers increase the synchronization power. Therefore, for a given technology, the potential power saving will depend on the balance between synchronization and datapath consumption.*

In this paper, the relationship between logic depth and power consumption is analyzed for three semicustom options: Xilinx and Altera FPGAs, and 1m CMOS Standard Cells from ES2. Main results show that, for all technologies, an intermediate value of logic depth that minimizes datapath power consumption exists. In FPGA-based prototypes, such reduction compensates the synchronization power overhead, but the reverse occurs in the cell-based circuits. In addition, the standard cell technology (SC) exhibits a relatively high off-chip consumption, making less significant any action to cut down other causes of power consumption.

Keywords: Low-Power, FPGAs, Standard Cells, Pipelining, Multipliers.

1. Introduction

Main consumption in CMOS technology corresponds to dynamic power: the energy per clock cycle involved in the charge/discharge of all the circuit node capacitances. This power component can be approximated by:

$$P = \hat{\alpha} C_n f_n V^2 \quad [1]$$

where f_n is the effective frequency of each circuit node (usually different from the system clock), C_n is the output capacitance of each node, and V is the power supply voltage. Thus, setting aside the power supply, the consumption can be modified by varying: the topology (which influences all the

variables); the data (which alters f_n); and finally, the interconnection network, which affects both C_n and f_n .

The effective frequency of each node can also be significantly incremented by the appearance of glitches. Although glitches do not produce errors in well-designed synchronous systems, they can be responsible of up to the 70 % of the circuit activity [1].

The useless consumption caused by glitches can be decreased in two ways: equalizing all circuit paths [2]-[4], or inserting intermediate registers or latches to reduce the logic depth [5], [6]. Although the relationship between logic depth and power consumption has been exhaustively analyzed in a wide variety of technologies and topologies [7], [8], the cost of chip fabrication constrained most of these studies to simulation-based results. This paper intends to contribute to this research area by adding some experimental data obtained from a set of circuits constructed and measured using different technologies.

The array multiplier proposed by H. Guild [9], [10] was selected as benchmark circuit. This topology present several benefits considering the objectives of the experiments. First, its high regularity makes straightforward its pipelining; second, a large set of reconvergent paths exists, a feature that contributes to the production of glitches; and finally, some module generators are available for this circuit [11], [12].

2. Experimental Results

In order to quantify the relationship between logic depth and power consumption several multipliers have been constructed. They have been pipelined with five different granularities β , defined as the maximum number of elementary processors (EP) between successive register banks [13]

Chip measurements have been done using random numbers as well as sixteen vectors that toggle, in each clock cycle, at least 93% of the

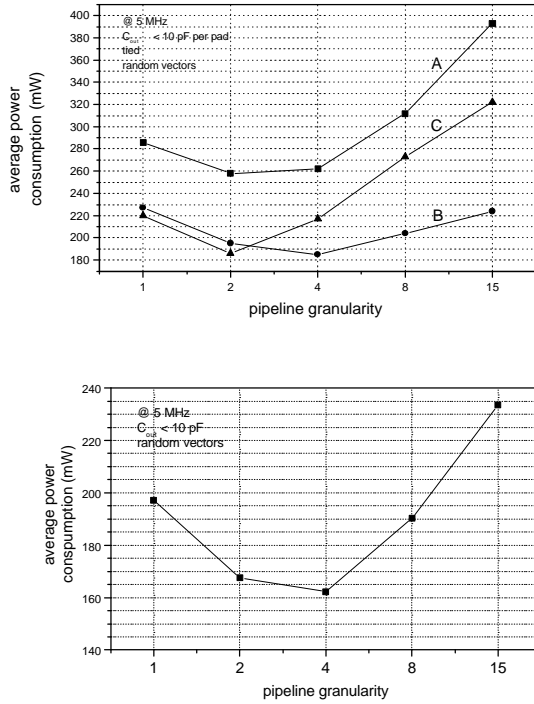


Fig. 1: Experimental results on FPGAs. Top: Xilinx FPGAs (A: XC3090, default compilation; B: XC3090, optimized mapping; C: XC4005, default compilation). Bottom: Altera Flex81188.

multiplier output bits. The last sequence makes easy to detect pipeline problems like double-clocking and zero-clocking [14] but also produces an increment of both datapath and off-chip power consumption. Although each sub-circuit has been tested at a complete set of frequencies, the following analysis have been particularized to 5 MHz for the FPGAs and 50 MHz for the SC circuit.

3.1 FPGA-based circuits

In Xilinx FPGAs, the relationship between pipelining and power consumption has been quantified for the *XC3090PC84-100* and the *XC4005PC84-6* chip models. In order to assess the effect of an efficient LUT utilization, two versions for each circuit have been constructed for the XC3090: a default map, place and route implementation; and another corresponding to a manual mapping optimization. The experiment was also repeated using an Altera *FLEX81188CG232-3*. In this case, the default compilation options were utilized.

Fig.1 shows the average power consumption of the multipliers versus pipeline granularity. The off-chip power fraction was maintained as low as possible in order to avoid masking datapath power effects. In this way, each pad supported just the 10 pF (max.) logic analyzer probe load.

Despite the hardware overhead, fine grain

pipelines not only ran faster than the combinational versions ($\beta=15$), but also exhibited lower consumption if operated at the same frequency. In all cases the minimum power value corresponded to logic depths between two and four LUTs. As a consequence, pipelining allows the FPGA user to trade power consumption for additional logic blocks and latency. For example, for default implementation conditions, the consumption of a $\beta=15$ multiplier can be reduced by 33 % (XC3090) or by 58 % (XC4005) if it is $\beta=4$ pipelined. In both cases the number of registers is increased from 32 to 104, and the latency from one to four clock cycles.

3.2 The Standard Cell Experiment

In SC, a careful planning of the experiments was mandatory to cut down prototyping costs. Thus, just one chip, composed of five independent 8-bit Guild multipliers pipelined with $\beta = 1, 2, 4, 8$ and 15, was constructed. In this case, β is equivalent to the logic depth¹, and indicates the number of processing elements between successive registers.

Each multiplier version was placed and routed independently in different regions of the die, using the *Placement Class* feature of Cadence DFVII tools [15]. All the multipliers have their own clock tree, but they share the I/O, input clock signal and power supply pads. In Fig.2, a block diagram of the chip is depicted.




Fig. 2: Chip block diagram

¹ In RAM-based FPGAs, pipeline granularity (processing elements between consecutive registers) and logic depth (in fact, LUTs between consecutive registers) is usually different, depending on the mapping process and the matching between the processing elements and the LUT characteristics.

	1	2	4	8	15
Pipeline Stages	15	8	4	2	1
NOR2 cells	64	64	64	64	64
FADD2 cells	64	64	64	64	64
REGISTER cells	408	216	120	72	48
Core area (mm ²)	1.9762	1.1942	0.8264	0.6498	0.5706
SC datapath (pF)	427	387	326	308	299
SC clock tree (pF)	183	648	53	39	8

Table 1: 8-bit Guild Multiplier characteristics versus β .

The input data multiplexing is performed by enabling just one (among the five) multiplier clock buffers; consequently, the other circuits remain inactive (all multipliers have registered I/O). At the output, a registered multiplexer delivers the results. Considering that multipliers are pad-limited circuits (see Fig.5), the strategy adopted in this experiment, an optimization of Calaway and Swartzlander ideas [16], leads to a significant reduction in the fabrication costs.

In Table 1, the main characteristics of each multiplier versus β are summarized. All versions have the same number of gates but different register count. Both datapath and clock tree node capacitances have been extracted from the post-layout reports.

In Fig.3, datapath, synchronization and off-chip power consumption versus pipeline logic depth is depicted. Off-chip power was determined by measuring the average input current corresponding to the pad ring (core and I/O have separated power supply inputs). This component was almost constant for all multipliers as a consequence of the common I/O pad utilization. However, in all circuits the off-chip power resulted dominant, even

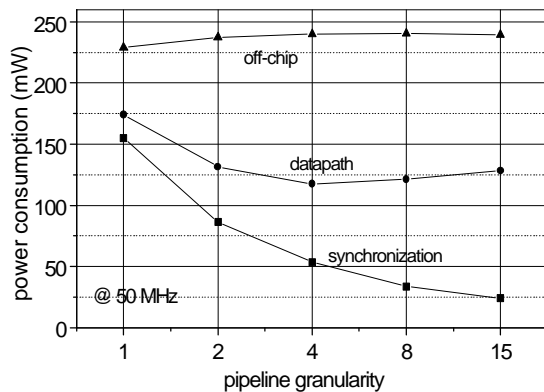


Fig. 3: Datapath, synchronization and off-chip power consumption at 50 MHz versus pipeline

considering that the load of each pad is not high. For example, it constitutes as much as the 60% of the overall consumption for the worst case (the combinational array), a percentage that is close to the pointed in [17] for other applications. Moreover, considering that the outputs have been previously registered, making insignificant the spurious activity at pads, the off-chip power would be still greater in these cases in which output registers were eliminated.

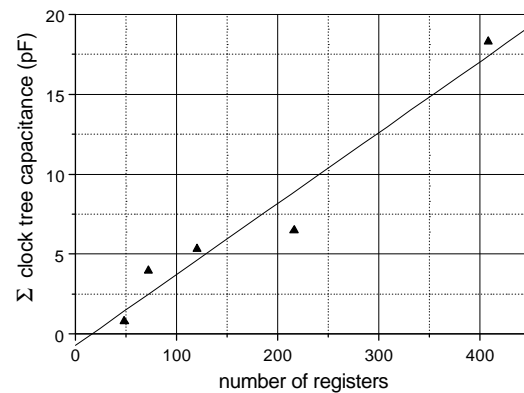



Fig. 4: Synchronization power (top) and overall clock tree capacitance (bottom) of each multiplier vs. number of registers.




Fig. 5: Chip photomicrograph

In Fig.4 the average synchronization power and clock tree capacitance versus the number of registers (NR) are shown for each pipeline multiplier. The graph evidences that registers rather than clock tree capacitance is the primary parameter to predict the synchronization power. In effect, a less linear function is obtained if NR is substituted by the sum of the overall load of each clock tree branch.

In Fig.3 can be observed that, as occurs on FPGAs, the datapath power consumption is minimum for the $\beta=4$ version (four processing elements between registers). However, the power saving is not as high as occurs on FPGAs. For example, $\beta=4$ pipelining allows datapath power consumption to be reduced by 12% (respect to the combinational array), but this gain can not compensate the synchronization power overhead corresponding to the 72 additional registers inserted in the datapath. As a consequence, the

combinational version ($\beta=15$) exhibited the lower overall consumption of all the set.

3. Conclusions

This paper shows that the efficacy of logic depth manipulation as low-power design strategy presents a strong dependence with the technology utilized; it results useful on FPGAs, but produce an opposite effect on Standard Cells. Both FPGAs and Standard Cells present almost equal off-chip power consumption, but this fraction is dominant on Standard Cells. For all circuits and technologies analyzed, the synchronization power exhibited a linear increment with the product *number of registers x clock frequency*, in coincidence with [8].

The profits of logic depth reduction from a power perspective is reinforced on FPGAs by some peculiarities of these devices: a) neither

synchronization nor off-chip power fraction are high in comparison with the datapath component; b) the die size and clock trees are fixed for a given chip model; and c) the interconnection delay is dominant.

Finally, the parallel experimentation on both technologies shows that the power consumption is significantly better in Standard Cells than in FPGAs. For example, the XC3090-based combinational versions exhibited a consumption (measured at 5 MHz) of 560 mW, versus the 8,07 mW corresponding to the equivalent SC circuit.

Acknowledgments

The authors wish to thank Fernando González Sanz for his assistance in the realization of the chip photomicrograph.

References

- [1] A. Shen, A. Gosh, S. Devadas y K. Keutzer, "On average Power Dissipation and Random Pattern Testability of CMOS Combinational Logic Networks", *Proc. ICCAD-92 Conf*, pp.402-407. IEEE Press, 1992.
- [2] T. Sakuta, W. Lee and P. Balsara, "Delay Balanced Multipliers for Low-Power/Low-Voltage DSP Cores", in *"Low-Power CMOS Design"*, A. Chandrakasan and R. Brodersen (Eds.), IEEE Press, 1998.
- [3] M. Pedram, "Power Minimization in IC Design: Principles and Applications", *ACM Trans. On Design Automation of Electronic Systems*, vol.1, nº1, pp.3-56, January 1996.
- [4] E. Boemo, S. Lopez-Buedo, and J. Meneses, "Some Experiments about Wave Pipelining on FPGAs", *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, Vol.6, N°2, June 1998.
- [5] Lemnios and K. Gabriel, "Low-Power Electronic", *IEEE Design & Test of Computers*, pp. 8-13, Winter 1994.
- [6] A. Chandrakasan, S. Sheng and R. Brodersen, "Low-Power CMOS Digital Design", *IEEE Journal of Solid-State Circuits*, Vol. 27, No. 4, pp. 473-484. April 1992
- [7] E. Mussol and J. Cortadella, "Low-Power Array Multiplier with Transition-Retaining Barriers", *Proc. PATMOS'95, Fifth Int. Workshop*, pp. 227-235, Oldenburg, October 1995.
- [8] J. Leiten, J. van Meerbegen and J. Jess, "Analysis and Reduction of Glitches in Synchronous Networks", *Proc. 1995 ED&TC*

Conference, pp.398-403. Los Alamitos: IEEE Press 1995.

- [9] H. Guild, "Fully Iterative Fast Array for Binary Multiplication and Addition", *Electronic Letters*, pp. 263, Vol. 5, No. 12, June 1969.
- [10] T. Hallin and M. Flynn. "Pipeline of Arithmetic Functions". *IEEE Trans. on Computer*, pp. 880-886. August 1972.
- [11] P. Ruíz, T. Riesgo, Y. Torroja, E. de la Torre and J. Uceda, "A Library of Reusable Arithmetic Components", *Proc XII DCIS*, pp.559-564, Universidad de Sevilla: 1997.
- [12] S. López-Buedo and E. Boemo, "Web-based Parameterized Module Generator", available at <http://www.ii.uam.es/~eda/>
- [13] C. Hauck, C. Bamji and J. Allen, "The Systematic Exploration of Pipelined Array Multiplier Performance", *Proc. ICASSP 85*, pp. 1461-1464. New York: IEEE Press, 1985.
- [14] J. Fishburn, "Clock Skew Optimization", *IEEE Trans. on Computers*, pp.945-951, July 1990.
- [15] Cadence, "Preview Cell Ensemble. Reference Manual". Cadence 1992.
- [16] T. Callaway and E. Swartzlander, "Estimating the Power Consumption of CMOS Adders". *Proc. IEEE 11th Symposium on Computer Arithmetic*, pp. 2210-216, Windsor, Ontario. July 1993.
- [17] M. Afghagi y C. Svensson, "Performance of Synchronous and Asynchronous Schemes for VLSI Systems", *IEEE Trans. on Computers*, vol.41, N°7, pp.858-872, July 1992.