Low-Power Low-Voltage Standard Cell Libraries with a Limited Number of Cells

Jean-Marc Masgonty[1], Stefan Cserveny[1], Claude Arm[1], Pierre-David Pfister[1],
Christian Piguet[1,2]

[1]CSEM, Neuchâtel, Switzerland
[2]LAP-EPFL Lausanne, Switzerland

Abstract. Systems on Chip (SoC) in deep submicron technologies contain several millions of transistors and have to work at lower and lower supply voltages to avoid too high power consumption. Consequently, digital libraries have to be designed to work at very low supply voltages and to be very robust while considering wire delays, signal input slopes, noise and crosstalk effects. At the electrical level, digital standard cells have been designed in a robust branch-based logic style, such as hazard-free D-Flip-Flops. At the RTL level, while considering a design flow using a logic synthesizer, a new approach has been proposed that is based on a limited set of standard cells. It results that the logic synthesizer is more efficient as it has a limited set of cells well chosen and adapted to the considered logic synthesizer. With significantly less cells than conventional libraries, the results show speed, area and power consumption improvements for synthesized logic blocks.

1 Introduction

For innovative portable and wireless devices, Systems on Chip (SoCs) containing several processors, memories and specialized modules are obviously required. Performances but also low-power are main issues in the design of such SoCs. In deep submicron technologies, SoCs contain several millions of transistors and have to work at lower and lower supply voltages to avoid too high power consumption. Consequently, digital libraries have to be designed to work at very low supply voltages and to be very robust while considering wire delays, signal input slopes, noise and crosstalk effects.

For SoCs, very important design problems have to be solved. They are mainly the silicon complexity (reliability, power consumption, interconnect delays), the system complexity (logic, MPU, memories, analog, RF, FPGA, MEMS), the design procedures (300 to 800 people for the design of a single chip, I.P re-use, design levels), verification and test. The total number of transistors on a single chip could be over one billion (predicted [2] to be between 4 to 19 billions in 2014 depending in the circuit type). Some partial answers have been given to the complexity problem, such as design re-use. The SIA Roadmap [2] predicts that in 2012, reuse of processors, logic blocks and peripherals, would reach about 90% of the embedded logic on the chip.

Low-power SoCs will be based on low-power components, such as processor cores, memories and libraries that are available with nice performances, i.e. 20’000 to 100’000 MIPS/watt [3] for some cores (using low-power techniques such as gated clocks). However, standard cell libraries are still a key point to the design of such low-power I.P. cores.

2 Conventional Low-Power Standard Cell Libraries

Several years ago, CSEM has designed digital standard cell libraries in a robust branch-based logic style [1, 4]. Such libraries with 60 functions and 220 layouts have been used for some years ago in the design of industrial chips [1, 4]. The low-power techniques used were the branch-based logic style that reduces parasitic capacitances, so-called “low-power” cell schematics with reduced activity and a clever transistor sizing. Instead to enlarge transistors to have more speed, parasitic capacitances were reduced by reducing the sizes of the transistors and therefore the capacitances that slow down the critical paths. Race-free or speed-independent sequential cells [5], such as D-Flip-Flops, were also proposed to be robust against delay variations. It was therefore possible to use such libraries at any supply voltage. If several years ago, power consumption reduction achieved compared to other libraries were about of a factor 3 to 5, it is today only about a factor 2 due to a better understanding of power consumption problems of library designers.

3 Present Situation

Today, logic blocks are automatically synthesized from a VHDL description while considering a design flow using a logic synthesizer such as Synopsys. Furthermore, deep submicron technologies with large wire delays require a better robustness, mainly for sequential cells sensitive to the clock input slope. Fully static branch-based logic style has still been found as the best. Furthermore, other D-Flip-Flop and Toggle structures much less sensitive to the clock input slope have been discovered [6]. However, the main question is the number of cells and which cells are optimal regarding the logic synthesis. It was known that for manually designed logic blocks, a large number of cells was beneficial for logic blocks optimization. A great number of complex gates allows the designer to reduce the number of transistors and the silicon area. Similarly, the activity was reduced while using complex gates. The question is the following: is this statement still true for a logic synthesizer?

4 A New Approach

A new approach is proposed, which is based on a limited set of standard cells. The number of functions for the new library has been reduced to 22 and the number of layouts to 92. It can be seen that the ratio between the number of layouts and the number of functions is larger (92/22=4.2 instead of 220/60=3.6 for the previous library). It means that the number of cell sizing and buffering is larger. For speed and power optimization achieved by the logic synthesizer, this enlarged ratio is beneficial.

It seems obvious that the logic synthesizer could do a better job if the number of cells in the library is large. With a larger choice, one can think that a better solution could be provided. But this is not the case. Experiments of Tables 1 and 2 show that the delay of some operators is significantly reduced with the new library resulting in a quite small increase in silicon area (Table 1) and that the silicon area is reduced at the same speed with the new library (Table 2).

These results show that the logic synthesizer is more efficient because it has a limited set of well-chosen cells and cell sizing adapted to the considered logic synthesizer. With significantly less cells than conventional libraries, the synthesizer is not lost in some optimizations due to a too large choice of cells.

The applied strategy for the design of the library was the following:

- to have only very fast cells, i.e. to remove all the cells with 3 P-ch transistors in series and to have a very limited number of cells with 2 P-ch transistors in series.
- to increase the number of cell sizing for the same function. However, it is not a simple increase from, for instance, sizing D1 (small transistors), D2 and D3 (medium-sized transistors) to D1, D2, D3, D4 and D5 (very large transistors). The cell sizing performed takes into account how the synthesizer uses the considered cells.
- to consider the combination of a given cell and of a buffer (or sized simple gate) to replace complex gates.

Such a strategy has to be checked through many experiments. The choice of the 22 functions was performed with a large number of experiments with and without a specific cell, and then the decision was made to insert or not this cell in the library. Similar experiments were performed with various sizing and buffering of the same cell. At the end, only 22 functions and 92 layouts were kept in the new library.

	Old Library		New Library	
	Delay [ns]	μm^2	Delay [ns]	μm^2
32-b multiply	16.4	907K	12.1	999K
fp adder	27.7	510K	21.1	548K
CoolRISC ALU	10.8	140K	7.7	170K

Table 1. Delay Comparison (synthesis for maximum speed, 0.5µm process)

	Old Library		New Library	
	Delay [ns]	μm^2	Delay [ns]	μm^2
32-b multiply	17.1	868K	17.0	830K
fp adder	28.1	484K	28.0	472K
CoolRISC ALU	11.0	139K	11.0	118K

Table 2. Silicon Area Comparison (synthesis for a given delay, 0.5µm process)

Furthermore, as the number of layouts is drastically reduced, it takes less time to design a new library for a more advanced process. A lot of time can also be saved for the library characterization, which is known to be often the most time consuming activity in library design. To reduce the number of layouts from 220 to 92 is definitively a great advantage.

It will also be a crucial point in future libraries for which more versions of the same function will be required while considering static power problems. A same function could be realized, for instance, with low or high VT for double VT technologies, or with several cells such as a generic cell with typical VT, a low-power cell with high VT and a fast cell with low VT.

5 Detailed Results

The first library designed with such an approach was realized in a 0.5µm process of a customer. However, the approach is certainly more and more interesting for deep submicron technologies, as explained at the end of Section 4. CSEM is currently working in the design of libraries in deep submicron technologies according to the proposed approach.

Obviously, the new library with 22 functions provides simple gates (NOR, NAND) and less complex gates. It is known that replacing complex gates with simple gates or decomposing complex gates in simple gates is beneficial for speed [7]. There are less transistors in series, less parasitic capacitances attached to the gate outputs, even if the critical paths contain more logic gates (but faster gates) in series. NAND gates are interesting as they provide only one P-ch transistor in their P-ch networks. The new proposed D-Dlip-Flop [6] is based only on NAND gates and inverters. However, generally, more transistors are required with simple gates. It is therefore less intuitive to see that logic synthesizers, with the same speed performances, are capable sometimes of reducing silicon area with only simple gates and very few complex gates as shown in Table 2.

Table 3 shows other experiments for which an easy to reach delay has been specified for the synthesis. With Tables 2 and 3, one can see that the results in silicon area are quite similar for both libraries.

	Old Library		New Library	
	Delay [ns]	µm^2	Delay [ns]	µm^2
16-b multiply	12.0	251K	12.0	234K
32-bit multiply	20.0	918K	19.8	943K
Decoder/comp	2.5	24K	2.5	21K
CoolRISC dec	1.4	9K	1.2	8K

Table 3. Silicon Area Comparison (synthesis for a easy to reach delay, 0.5µm process)

However, as shown in Table 1, if one asks for the fastest logic block, one can see that the synthesis with the new library is roughly 30 to 40% faster with a very moderate increase in silicon area.

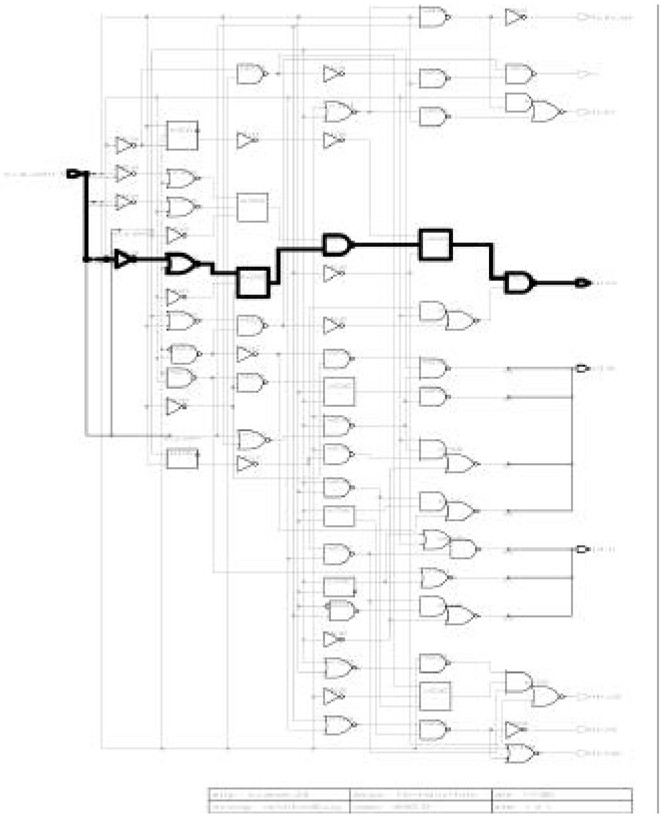

Figure 1. Small decoder with its critical path implemented with the old library

Speed, area, and power consumption comparison were performed on several VHDL models available at CSEM. The models were mainly arithmetic components.

The signed multiplier model is a Radix-4 Booth Multiplier which description is found in [8]. The floating-point adder is a model partially based on [9], with a few architectural modifications. Other components were extracted from the CoolRISC processor [3, 12]. These are the CoolRISC ALU, a decoder and a comparator.

Figure 1 shows a small decoder with its critical path implemented with the old CSEL-LIB 4 library. It contains 65 gates or 73 if the hidden inverters included in complex gates are taken into account. The total number of inverters is 26. The critical path contains 7 gates and, according Table 3, last row, achieves 1.4 ns.

Synthesized with exactly the same constraints, the same decoder implemented with the new library is shown in Figure 2. It contains less gates (53) and less inverters (13) than the previous implementation. The critical path, which is now situated between another input-output path, contains less gates (5) and is faster (1.2 ns).

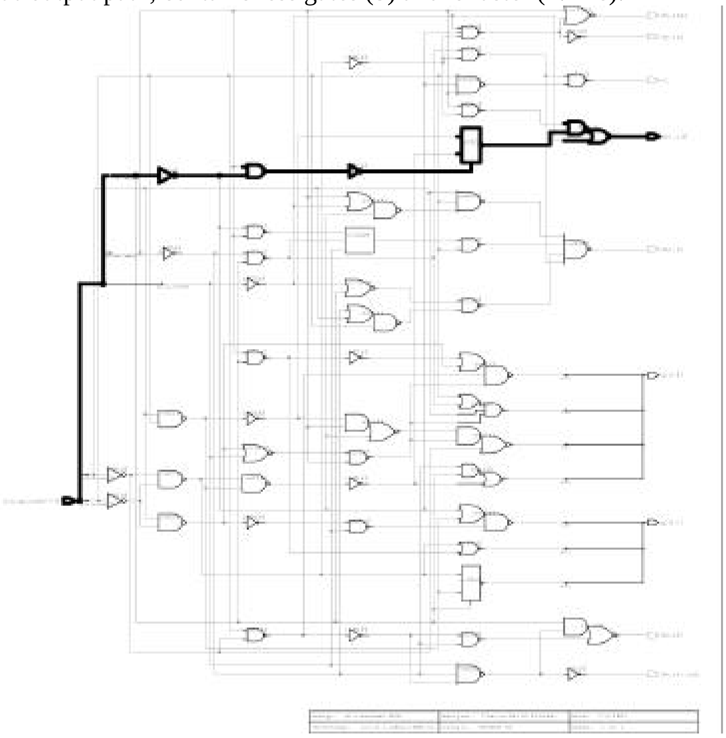

Figure 2. Same small decoder with its critical path implemented with the new library

6 Power Consumption

Power consumption was evaluated using the following tools:

- Synopsys Design Compiler for synthesis
- ModelSim to generate a set of random test vectors
- Calibre to convert the Verilog netlist into a Spice netlist
- Mach TA to analyse the power consumption.

The first step in the design flow was to convert the Verilog netlists synthesized for both the old and new libraries into Spice netlists, which can be used with Mach TA.

This was performed using Calibre's v2lvs tool. Furthermore, a Modelsim simulation of the source VHDL models was performed in order to obtain a set of random test vectors. In our case, we simulated 100 random vectors on each component. These vectors were converted into Mach TA's format using Mach TA's tssi2vect tool. Some minor editing of the generated files was required in order to suppress errors and warnings at compilation time: this was done using UNIX scripts, and tools such as sed. Once both Spice netlist and Mach TA test vectors were ready, the Mach TA simulation was performed using Mach TA's mpa tool. The same transistor models from the c05a (0.5µ) technology were used for both old and new library. The simulation was performed at the same voltage (2.5V) and temperature (25°C). The average current during simulation was extracted. This is shown in Table 4.

	Old Library		New Library	
	Delay [ns]	Iavg [mA]	Delay [ns]	Iavg [mA]
16-b multiply	12.0	8.61	12.0	7.17
32-bit multiply	20.0	19.21	20.0	15.95
Fpadd 32-bits	30.0	6.17	30.0	4.94

Table 4. Average Current Comparison (designed for the same speed, 0.5µm)

The results for the new library show an improvement of about 15% over the old library for components synthesized at the same speed. The improvement even reaches 20% in the case of the floating-point adder. The difference can be explained by the fact that the multipliers are mainly built up with one type of cell (full-adder) whereas the floating-point adder instantiates a greater diversity of cell types.

7 Static Power

The main issue is the design of future libraries will be the static power. For Vdd as low as 0.6 to 0.3 Volt in 2014, as predicted by the Roadmap [2], VT will be reduced accordingly in very deep submicron technologies. Consequently, the static power will increase significantly due to these low VT [10]. Several techniques with double VT, source impedance, well polarization, dynamic regulation of VT [11] are today under investigation and will be necessarily used in the future. This problem is crucial for portable devices that are often in stand-by mode in which the dynamic power is reduced to zero. It results that the static power becomes the main part of the total power.

As mentioned earlier, these static power problems will necessarily imply multiple versions of the same cell or function. In some advanced technologies, there are three different cells with high, typical and low VT for low-power, generic and fast cells. It means that three different layouts have to be designed. For techniques aiming at an electrical VT regulation, one has to polarize substrate or wells. These cells are larger than conventional cells, so probably standard cell libraries will provide both versions. It is why standard cell libraries with a limited number of functions, i.e. 22 as proposed, is a very important point for future libraries as well as for library characterization. This last point has to be stressed for libraries working at any voltage, for instance for microprocessors or logic blocks with a dynamic regulation of Vdd and VT.

8 Conclusion

A new approach has been proposed for standard cell libraries with a limited set of functions reduced of a factor 3 compared to conventional libraries and adapted to logic synthesizers. The results show that speed is improved with a similar silicon area and that silicon area is reduced at same speed performances. Furthermore, the work to design and characterize a new library is significantly reduced, which is a very important point for future libraries that have to take into account static power problems by providing many versions of the same cell with, for instance, high, typical and low VT.

References

[1] C. Piguet, J.-M. Masgonty, S. Cserveny, E. Dijkstra, "Low-Power Low-Voltage Digital CMOS Cell Design", PATMOS'94, Fourth International Workshop Power and Timing Modelling, Optimization and Simulation, October 17-19, 1994, Barcelona, Spain, pp. 132-139.

[2] http://notes.sematech.org/ntrs/rdmpmem.nsf

[3] C. Arm, J-M. Masgonty, C. Piguet, “Double-Latch Clocking Scheme for Low-Power I.P. Cores”, PATMOS 2000, Goettingen, Germany, September 13-15, 2000.

[4] "Low Power Design in Deep Submicron Electronics", NATO ASI Series, Series E: Applied Sciences - Vol. 337, Editors W. Nebel and J. Mermet, Kluwer Academic Publishers, 1997.

[5] C. Piguet, J. Zahnd, "Signal-Transition Graphs-based Design of Speed-Independent CMOS Circuits", ESSCIRC'98, September 21-24, 1998, Den Haag, The Netherlands, pp. 432-435.

[6] C. Piguet, “Robustness of Asynchronous Sequential Standard Cells in a Synchronous Environment”, AINT’2000, “Asynchronous Interfaces”, Delft, The Netherlands, July 19-20, 2000.

[7] C. Piguet, "Design of Low-Power Libraries", Invited Paper at ICECS'98, September 7-9, 1998, Lisbon, Portugal, Vol. 2 of 3, pp. 2.175-2.180.

[8] Gary W. Bewick, "Fast Multiplication: Algorithms and Implementation", Stanford University Dissertation, February 1994.

[9] Nhon Quach, "Design and Implementation of the SNAP Floating-Point Adder", Stanford University Technical Report, CSL-TR-91-501, December 1991.

[10] D. Liu, C. Svensson, "Trading Speed for Low Power by Choice of Supply and Threshold Voltages", JSSC-28, No 1, Jan. 1993, pp. 10-17.

[11] V. von Kaenel et al. "Automatic Adjustment of Threshold & Supply Voltage Minimum Power Consumption in CMOS Digital Circuits", 1994 IEEE Symposium on Low Power Electronics, San Diego, October 10-12, 1994, pp. 78-79.

[12] C. Piguet et al. "Low-Power Design of 8-bit Embedded CoolRISC Microcontroller Cores", IEEE JSSC, Vol. 32, No 7, July 1997, pp. 1067-1078.