Logic Optimization and Synthesis: Trends and Directions in Industry

Luca Amarú*, Patrick Vuillod†, Jiong Luo*, Janet Olson*

*Synopsys Inc., Design Group, Sunnyvale, California, USA
†Synopsys Inc., Design Group, Grenoble, France

***Abstract*—Logic synthesis is a key design step which optimizes abstract circuit representations and links them to technology. With CMOS technology moving into the deep nanometer regime, logic synthesis needs to be aware of physical informations early in the flow. With the rise of enhanced functionality nanodevices, research on technology needs the help of logic synthesis to capture advantageous design opportunities. This paper deals with the synergy between logic synthesis and technology, from an industrial perspective. First, we present new synthesis techniques which embed detailed physical informations at the core optimization engine. Experiments show improved *Quality of Results* (QoR) and better correlation between RTL synthesis and physical implementation. Second, we discuss the application of these new synthesis techniques in the early assessment of emerging nanodevices with enhanced functionality. Finally, we argue that new synthesis methods can push further the progress of electronics, as we have reached a multiforking point of technology where choices are tougher than ever.**

I. Introduction

The field of logic synthesis originally started from the classic work on switching theory [1]. It then took a sharp turn in the eighties with the advent of *Application-Specific Integrated Circuits* (ASICs), enabled by *Very-Large-Scale Integration* (VLSI), demanding for more powerful synthesis techniques and tools. Since then, the continuous progress on logic synthesis has enabled the semiconductor industry to design chips with billions of transistors [2]. Despite being a mature technical field, logic synthesis still offers intellectual challenges of the highest order. In particular, the interaction between synthesis and advanced technologies raises new problems. Innovative solutions for these problems have the potential to prolong the exponential growth of computer power.

Nowadays, the main directions for logic synthesis are scalability and physically-accurate logic optimization. Scalability is essential because state-of-the-art instances of synthesis problems have tens, or even hundreds, of millions of logic nodes. In this scenario, efficient synthesis algorithms must have strictly linear time complexity with the circuit size. The use of physical information during logic optimization, also called logic restructuring, becomes crucial as we move further into deep nanometer technology nodes. Indeed, the impact of placement and routing on the synthesis results grows with the shrinking of digital devices, as the interconnects turn into the dominant delay components. Furthermore, functionally equivalent gates in different technology libraries can have quite diverse area, delay and power characteristics, because of specific logic styles and cell layouts. Embedding as much technology information as possible early in the logic optimization engine is key to make advantageous logic restructuring opportunities carry over at the end of the design flow.

In this paper, we examine the synergy between logic synthesis and technology, from an industrial perspective. We present technology aware synthesis methods incorporating advanced physical information at the core optimization engine. Internal results evidence faster timing closure and better correlation between RTL synthesis and physical implementation. We elaborate on synthesis aware technology development, where logic synthesis enables a fair system-level assessment on emerging nanodevices with enhanced functionality. We conclude that industrial and academic research on logic synthesis is still required to push further the progress of electronics.

The remainder of this paper is organized as follows. Section II presents innovative synthesis methods that are technology aware at different abstraction levels. Section III considers synthesis aware technology development and its growing impact on the nanotechnology research community. Section IV concludes the paper.

II. Technology Aware Synthesis

With the advent of deep nanometer technology nodes and the increasing chip complexity, the separation between logic synthesis and physical design progressively vanishes in favor of a technology aware synthesis framework. While several works, in academia and industry, already considered physically aware logic techniques [3], we aim at pushing the efficacy of this approach to the next level. We do not just consider physical models to compute more accurate timing and area costs during synthesis, but we also create optimization engines and databases capable of capturing the specific logic opportunities distinctive of a technology. We achieve this goal by developing (i) library-dependent exact circuit databases, (ii) specialized logic rewriting techniques and (iii) physically-accurate timing models for optimization. We elaborate more on each topic hereafter.

A. Creation of Library-Dependent Circuit Databases

The relative area, delay and power characteristics of a logic gate can sensibly vary between different technology libraries. For example, MUX gates in some technologies can have comparable delay to NAND gates, while in other technologies

978-3-9815370-8-6/17/$31.00

can have much larger delay, close to XNOR gates. This can be due to different gate design styles, e.g., complementaty static *vs.* pass transistor style, to different technology nodes, e.g., 45 nm *vs.* 22 nm, to different layouts, e.g., relative widths of standard cells, or any combination of those and other effects. As a consequence, conventional technology assumptions at the basis of logic minimization algorithms may not be accurate at all times. For this reason, we develop circuit databases which contain optimum circuit realizations specific to a given technology. Such databases can be complete, i.e., containing circuits for all the 2^{2^n} functions of n variables, for some limited n, or may be partial, i.e., containing only the most frequently occuring functions of n variables. Thanks to smart encoding and data mining techniques, we have nowadays access to complete databases for $n = 4$ and 5, and we are able to extend partial databases to $n = 7$ and 8. The application of these databases in large-scale synthesis is discussed in the next subsection. The circuits stored in the databases are highly optimized and, in most cases, exact, i.e., achieving the global minimum for area, delay and power metrics. While exact solutions for area can be computed and stored with existing methods, exact solutions for delay and power are inherently more difficult to find and to enumerate. This is because of the additional boundary conditions typical of delay and power goals, such as arrival times at the inputs and switching activities at the inputs. Thanks to our recent breakthroughs on this topic, we are able to overcome these issues and create the most complete exact circuit databases for all key design goals. It is important to mention that technology libraries are not known a priori so these databases must be populated online during design time. Complete databases for 4 variables, fed with modern technology libraries, can be computed in less than a minute. Partial databases can be progressively computed and stored during design time. Fig. 1 shows a sample entry for an exact-delay database, corresponding to (i) a specific 4-input Boolean function, (ii) figurative library delay characteristics and (iii) input arrival times. Please note that in this illustrative example the input arrival times are fixed, but in the actual database we have access to exact results for any other arrival time pattern.

B. Logic Rewriting with Exact Databases

Logic rewriting is a fast greedy algorithm for optimizing a logic circuit [4]. It iteratively selects sub-circuits rooted at a node and replaces them with pre-computed sub-circuits with minimum local area, delay, power, or any combination of these metrics. The pre-computed sub-circuits are stored in a database. It is possible to build circuit databases for any design goal and with arbitrary optimization quality. Of course, the more optimized a database is, the better the logic rewriting results will be. Through the use of the library-dependent exact databases discussed in the previous subsecton, it is possible to improve significantly the efficacy of logic rewriting. Our internal results show improved QoR and better correlation between RTL synthesis and physical implementation. The tight link to technology and guaranteed local optimality are the key

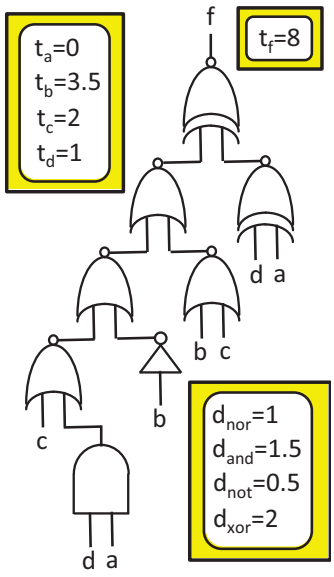

Fig. 1. Exact-delay circuit for given arrival times and gate library values.

factors enabling the success of this approach.

The scalability of logic rewriting makes it applicable to very large designs, with many millions of nodes. However, the impact of each replacement must be evaluated at the global level. While, locally, an exact database replacement is advantageous by construction, its interaction with the rest of the circuit may globally degrade some design metric. This is mostly due to global placement and routing adjustement, in particular to modified wire lengths and wire capacitances. For this reason, accurate physical models are still needed to evaluate the impact of each local optimization move. Delay, commonly referred to as *timing*, is very sensible to circuit alterations. We elaborate more on physically-accurate timing models for optimization in the following subsection.

C. Physically-Accurate Timing Models for Optimization

Physical effects, such as assignment of placement areas, wire length, wire capacitance and routing in general, can have significant impact on a design delay and need to be accounted early in the design flow, already during logic synthesis. Modern synthesis and optimization tools take these and other effects into account via physical models to achieve superior quality of results. It is important to note that the accuracy of a physical model may have a noticeable impact on the runtime complexity of a synthesis tool. While it is always desirable to use the most accurate physical model available, from a QoR perspective, more relaxed tradeoffs are pursued in practice, to alleviate the runtime costs. For example, the wire load model is a popular instrument to estimate the effect of wire length and fanout on the resistance, capacitance, and area of nets [5]. Semiconductor vendors develop wire load models, based on statistical information specific to the vendors process. Standard wire load models include coefficients for area, capacitance, and resistance per unit length, and a fanout-

to-length table for estimating net lengths. However, such models are usually generated on a design basis and do not account for detours and nets in congested regions. Nowadays, we have access to enhanced wire estimation models which overcome the aforementioned limitations. In our advanced design frameworks, wire capacitance, resistance and delay are estimated within synthesis on a net per net basis. This enables a more accurate assessment of the new interconnects introduced by local optimizaton moves. Globally, it translates in better correlation between RTL and physical design.

III. Synthesis Aware Technology

In the quest to build increasingly powerful computers, several emerging nanotechnologies are under consideration. Some of them mainly offer devices with enhanced performance over CMOS while others also offer enhanced functionality over standard switches. Enhanced functionality devices have the potential to build complex digital systems with fewer physical resources than in standard technologies. To fully exploit this potential, logic synthesis must adapt to harness a technology unique expressive power [6].

In the following, we discuss how the technology-aware synthesis techniques introduced in Section II can find application in the efficient optimization of emerging nanotechologies with enhanced functionality. Moreover, we evaluate the enabling role of logic synthesis in the early phases of technology exploration, in order to direct research efforts on the most promising nanodevices.

A. Synthesis for Enhanced Functionality Devices

From a logic standpoint, emerging nanodevices with enhanced functionality lead to technology libraries with complex type of gates and unconventional characteristics. For example, controllable polarity transistors, such as double gate silicon nanowire FETs, enable XNOR gates with half of the devices required by traditional CMOS technology [7]. Spin wave devices realize MAJ gates at the same cost as AND/OR gates [8]. Disregarding interconnections as a first approximation, it is possible to generate equivalent gate libraries for these nanotechnologies. In this scenario, the gate area, delay and power characteristics reflect the logic expressive power of the elementary devices employed. From there, we can build exact databases for an emerging nanotechnology using the same rationale presented in the previous section. Logic rewriting techniques, set up accordingly, are suitable to exploit the logic expressive power of a given nanotechnology.

Although the aforementioned approach paves the way for a generally applicable *technology-dependent technology-independent* optimization methodology, there are still many challenges to be addressed. First, logic rewriting is a powerful technique but its logic restructuring ability remains local. Global logic restructuring techniques are still needed, especially for random logic portions of a design. When targeting enhanced functionality devices, the global restructuring techniques must be devised to extract the native logic primitives of a technology. Second, interconnection schemes for logic gates in emerging technologies can be very different than in CMOS. For instance, information transmission in many spin-based technologies is achieved by propagating spin waves in a clocked fashion [8]. Without considering this information, logic optimization for spin wave circuits may be ineffective. Interconnect models and physical constraints need to be included in the optimization engines for emerging nanotechnologies, as is currently done for CMOS.

B. Synthesis-guided Technology Exploration

The exploration of new technologies traditionally aims at improving performance, reducing the area footprint, decreasing leakage and switching energy of primitive devices. Intuitively, primitive devices with superior physical properties lead to more efficient digital systems. However, we have seen that this is not the only way to increase the efficiency of digital systems. Devices with enhanced functionality can implement complex systems with few physical resources [6]. Therefore, it is important to probe early in the exploration of new devices both the physical and logic properties. Highly-flexible synthesis approaches, such as the one discussed in this paper, can be used to guide the research efforts of technologists. Indeed, without taking into account their logic expressive power, many promising emerging devices would be prematurely discarded from consideration.

IV. Conclusions

In this paper, we discussed the interaction between logic synthesis and advanced technologies, from an industrial perspective. We introduced technology aware synthesis methods incorporating advanced physical information at the core optimization engine. Our internal results show faster timing closure and better correlation between RTL synthesis and physical implementation. We considered synthesis aware technology development, where logic synthesis enables a fair system-level assessment on emerging nanodevices with enhanced functionality. We concluded that industrial and academic research on logic synthesis is still required to push further the progress of electronics.

References

[1] E. J. McCluskey, *"Minimization of boolean functions"*, Bell Syst. Techn. J., vol. 35, no. 6, pp. 14171444, 1956.

[2] G. De Micheli, *"Chip Challenge"*, in IEEE Solid-State Circuits Magazine, vol. 2, num. 4, p. 22-26, 2010.

[3] W. Gosti, S. P. Khatri, A. L. Sangiovanni-Vincentelli, *"Addressing the timing closure problem by integrating logic optimization and placement"*, IEEE/ACM international conference on Computer-aided design (ICCAD), 2001.

[4] A. Mishchenko, S. Chatterjee, R. Brayton, *"DAG-aware AIG rewriting a fresh look at combinational logic synthesis"*, Design Automation Conference (DAC), 2006.

[5] K. D. Boese, A. B. Kahng, S. Mantik, *"On the relevance of wire load models"* International workshop on System-level interconnect prediction, 2001.

[6] L. Amaru, P.-E. Gaillardon, S. Mitra, G. De Micheli, *New Logic Synthesis as Nanotechnology Enabler*, Proceedings of the IEEE, 103.11, (2015): 2168-2195.

[7] M. De Marchi, et al. *"Polarity control in double-gate, gate-all-around vertically stacked silicon nanowire FETs"*, IEEE International Electron Devices Meeting (IEDM), 2012.

[8] K. Bernstein, R. K. Cavin, W. Porod, A. Seabaugh, J. Welser, *"Device and architecture outlook for beyond CMOS switches"* Proceedings of the IEEE, 98.12 (2010): 2169-2184.