

From Design to Design Automation

Jason Cong

Computer Science Department and
Electrical Engineering Department
University of California, Los Angeles
cong@cs.ucla.edu

ABSTRACT

I had the pleasure to work at the Xerox Palo Alto Research Center (PARC) as a summer intern in 1987 with Dr. Bryan Preas. One of the many valuable lessons that I learned through that summer internship is that as a researcher in electronic design automation (EDA), there is a tremendous value to be had from an involvement in VLSI circuit and system designs. This involvement guarantees a first-hand opportunity to discover and formulate new and insightful EDA problems, and to develop practical and impactful solutions. This experience had a great influence on my career as an EDA researcher. In this paper, I would like to share several examples of research projects that I directed at UCLA which followed the principle of starting from design to gain insight and inspiration for design automation.

Categories and Subject Descriptors

B.5.2 [Design Aids]: Automatic synthesis, optimization

General Terms

Algorithms and Design

Keywords

Design, design automation, Boolean matching, buffer planning

1. INTRODUCTION

I began my graduate study at the University of Illinois at Urbana-Champaign in the winter of 1986, and my research area was in VLSI CAD (computer-aided design) or electronic design automation. In my second year as a graduate student, I received an offer to work at the Xerox Palo Alto Research Center (PARC) as a summer intern, and I was fortunate to have Dr. Bryan Preas as my mentor. He was already a highly respected researcher in design

automation at that time, with much pioneering work in IC placement and routing (e.g., [1][2][3][4][5][6]). Prior to my internship at PARC, I had just completed my first research project on three-layer channel routing; this became my MS thesis and was later published in ICCAD'87 [7]. I was in the process of looking for the next research problem, and although I had a paper accepted for publication in a top EDA conference, I had never done any VLSI design. Like many other graduate students in EDA, I was reading the leading journals and conference proceedings in the field and looking for opportunities for improvement of published results on some well-formulated problems. When I began my internship at PARC, the first assignment that Bryan gave me was to create some small circuit designs using the in-house VLSI design automation (DA) system developed at PARC, named DATools. It was a very advanced DA system at that time and allowed the designers to go from schematics to layout in a fairly automated fashion. This was a great learning experience for me, and I was exposed to many important concepts and problems that were totally new, such as the hierarchical design methodology, cell library design and selection, static timing analysis, and layout verification. It also gave me a clear perspective on the important factors that affect the chip area, which was the main optimization objective those days.

After I successfully completed the full design cycle, I was assigned to investigate whether one could improve the standard cell global routing package. At that time, the standard cell design methodology was just becoming popular. The Timberwolf placement package developed at UC Berkeley [9] based on simulated annealing was a big success, and it improved many in-house standard placement tools, including the one used inside of DATools. But there were relatively fewer studies on standard cell global routing. At that time, many standard cell based designs were done with one metal layer, and the two-metal-layer process had just become available. Feedthrough cells were needed to complete routing between cells in different rows. At that time, most of the standard cell global routing tools focused on wirelength minimization. Some considered feedthrough minimization. From my limited design experience gained at the beginning of that summer, I made the observations that (i) the chip area is determined by the maximum cell row length and the total channel density, (ii) the total channel density is a better metric to

optimize than the total wirelength, and (iii) feedthrough assignment affects both the maximum row length and the total channel density. Based on these observations, I developed a new standard cell global routing algorithm that considered the interplay between feedthrough assignment and routing topology optimization, and we also proposed a novel optimization technique called iterative deletion for total channel density optimization¹. This algorithm considerably outperforms both the standard cell global routing tool used in DATools and in the Timberwolf package [9]; it was selected as a highlighted paper for ICCAD'88 [10].

My summer internship with Bryan Preas was a great experience in multiple ways. First, I was able to publish another major research paper, which became a chapter in my PhD dissertation. Second, it was exciting to see my algorithm being used to generate real chip layout (as part of the Xerox PARC DATools System)—all of which was such a gratifying experience to a young PhD student. More importantly, I learned this valuable lesson: As a researcher in EDA, there is a tremendous benefit to be gained from an involvement in VLSI circuit and system designs. That involvement affords the opportunity for a first-hand experience in the discovery and formulation of new and insightful EDA problems, and the development of practical and impactful solutions. This experience has had a great influence on my career as an EDA researcher. In this paper, I will share several examples of the research projects that I directed at UCLA following this principle—from design to design automation.

2. BOOLEAN MATCHING FOR LUT-BASED FPGA SYNTHESIS

After I joined the faculty of UCLA in 1990, in addition to my research in VLSI physical designs, I soon developed a parallel interest in FPGA synthesis. Following the circuit model used in most FPGA synthesis papers in the early 1990s, my research team initially assumed that all the programmable logic blocks (PLBs) in SRAM-based FPGAs were uniform-size lookup tables (LUTs). We made significant progress on technology mapping LUT-based FPGAs, including the development of the first depth-optimal polynomial-time mapping algorithm [11]. Then, we began some FPGA design and prototyping efforts on commercial FPGAs. It was interesting to see, however, that some popular FPGAs, in fact, do not have uniform-size FPGAs. For example, the PLB in the widely successful Xilinx XC4K family consists of three LUTs of different sizes, as shown in Figure. 1.




Figure 1. Xilinx XC4000 PLB [11].

It can implement some functions of up to nine inputs. But there is no theory or algorithm to decide what kind of logic can be mapped to such a PLB (not even by Xilinx, even though they designed and manufactured the chip with this PLB). To gain a better understanding of the efficiency of these PLBs, and more importantly to provide a tool to explore new PLBs in the future, we developed a theory and algorithm based on Boolean matching techniques to completely characterize the kind of logic functions that can be mapped to Xilinx XC4K PLBs ([12][13]). The result was exciting news (and surprising) to Xilinx. Overnight, Xilinx gave over 1,868 logic functions for testing. Our algorithm successfully mapped each of them to a Xilinx 4K PLB. It turned out that these test cases were actually customer complaints that had accumulated over the years (e.g., some customers reported that they could map the function manually to a XC4K PLB, but the automated mapping tool failed). Since our theory and algorithm provided the complete characterization of the set of Boolean functions implementable on the PLB², these were straightforward test cases for us [13]. This algorithm, together with other FPGA mapping algorithms that we developed at UCLA, is available for free download as part of the RASP system [16]. These algorithms were used by Aplus Design Technologies, a spin-off from the VLSI CAD Laboratory at UCLA. Aplus was among the first to provide FPGA architecture evaluation tools and physical synthesis tools. Aplus was acquired by Magma Design Automation (now part of Synopsys) in 2003.

3. 3D IC PHYSICAL DESIGN

In the early 2000s, my group was involved in a DARPA program to develop physical design automation tools for three-dimensional integrated circuits (3D ICs). As with most DARPA programs, this was very much a forward-looking project, as there were very few (if any) people doing 3D IC designs at that time. The first step of our research was an in-depth study of various IC design and fabrication technologies, and we worked with our partner CFDRG to come up with detailed technology and circuit

¹ The iterative deletion technique was also used for circuit partitioning [8]

² More general, but less efficient Boolean matching algorithms were later presented in [14] and [15].

models for 3D ICs. These studies gave us a more realistic understanding of the needs and challenges of 3D IC physical designs. We made two key observations. One observation focused on the thermal concerns due to both device stacking and low thermal conductivity of the dielectric layers (proposed or used in some 3D IC processes at that time). To address this problem, we were among the first to develop thermal-driven 3D floorplanning [17], placement [18], and routing [19] tools for 3D IC designs. Our second observation was that through-silicon-vias (TSVs) played a unique role in 3D IC designs. Once the TSV locations are determined, a 3D IC design can be decomposed into a sequence of 2D IC designs using standard 2D IC layout tools. Therefore, our research focused a lot on TSV planning and placement ([20][21]). Based on these research efforts, we jointly developed a 3D IC design environment with IBM and Penn State University. We released a complete 3D IC physical design tool, named 3D-Craft [22] (See Figure 2), in 2009 with interfaces to the OpenAccess design infrastructure [OA]. Our 3D placement tool was used by Prof. Paul Franzon’s team at North Carolina State University for three 3D IC designs, including a 2-point FFT butterfly processing element (PE), an Advanced Encryption Standard (AES) encryption block, and a multiple-input and multiple-output wireless decoder (MIMO) [23]. Their study showed that our 3D placement tool leads to significant wirelength and power reduction compared to manual 3D placement solutions.

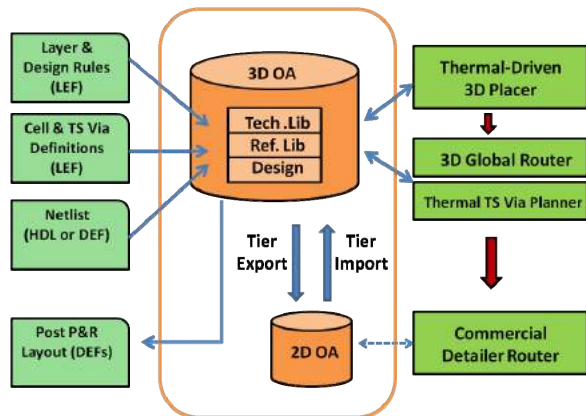


Figure 2. Overview of 3D-Craft [22].

4. HIGH-LEVEL SYNTHESIS

In the early 2000s, Xilinx introduced its Virtex-II Pro FPGAs with embedded IBM PowerPC 405 processor cores. Like many other researchers, we were very interested in using such a highly integrated field-programmable system-on-a-chip (FP-SoC) for prototyping and accelerated computing. After some initial exploration, it was evident to me that to design such an FP-SoC, one should raise the level of design abstraction so that the input language can specify both the function to be implemented on the FPGA fabric as well as the computation to be

executed on the embedded processors. For this reason, we selected C/C++ as the input languages (even though they were not ideal candidates for hardware synthesis) and started a research project on high-level synthesis (HLS) from C/C++ to cycle-accurate hardware descriptions in VHDL or Verilog. The concept of HLS was not new when we started, but prior HLS systems often failed to match the quality of human RTL designs. Therefore, we focused on optimization for a better quality of results and developed a number of novel algorithms to improve the quality of the HLS results, such as scheduling using systems of difference constraints [24], efficient pattern-mining [25], scheduling with soft constraints [26], optimization using behavior-level don’t-cares ([27][28]), and automatic memory partitioning [29]. The project led to the development of xPilot, a novel platform-based HLS system [30] (See Figure 3), which was licensed to AutoESL Design Technologies, another spin-off from our lab at UCLA, for commercialization in 2006.

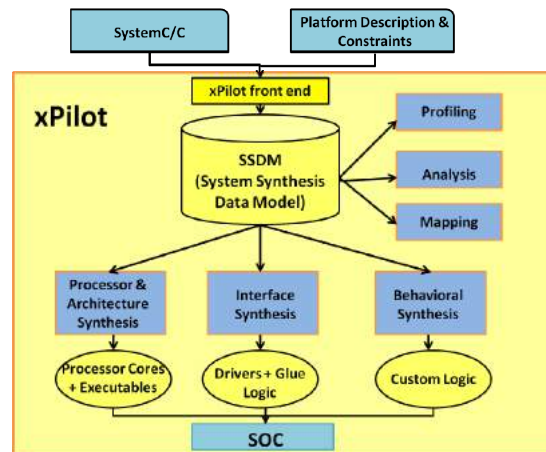


Figure 3. xPilot system-level synthesis framework [30].

Throughout the development of xPilot and AutoPilot (the HLS tool from AutoESL), my research group has been an active user of the HLS tools mainly for accelerated computing. In turn, these design activities helped to introduce interesting new research problems for HLS. One good example is memory partitioning. It is quite common for a HLS tool to map an array in the input C/C++ language to a memory block in the hardware implementation. However, most memory blocks have a very small number of ports (in fact, only two ports for most embedded memory blocks on FPGAs), which greatly limits the degree of parallelism in the final implementation. To overcome this limitation, the designer has to go through a tedious process to re-write the input C/C++ code to partition an array into a set of smaller arrays to enable multiple concurrent memory accesses [31]. To spare designers from such manual code rewriting, we quickly developed the theory and algorithms for automatic memory partitioning ([29][32]), and our work in [28] received a Best Paper Award from the ACM Transactions on Design Automation of Electronic Systems in 2013.

AutoESL's HLS tool was adopted by some of the world's largest software and semiconductor companies. AutoESL was acquired by Xilinx in 2011. The AutoESL's HLS tool was renamed as Vivado-HLS and is now available to tens of thousands of Xilinx FPGA designers worldwide, becoming the most widely deployed and used HLS tool in the history of EDA.

5. CONCLUDING REMARKS

I was very fortunate to have that summer research opportunity at Xerox PARC in 1987 and to have Bryan Preas as my mentor. I built up a long-term friendship with Bryan, and I benefited immensely from Bryan's vast experience, deep insight, and remarkable wisdom in the subsequent years. I considered Bryan as my second PhD advisor. The truly important lesson that I learned from Bryan is to go from design to design automation—always trying to gain first-hand knowledge and experience about needs from the designer's perspective, so that one can discover and formulate new and insightful EDA problems and develop practical and impactful solutions. This guiding principle led to multiple successes in my career as an EDA researcher, and I illustrated some of them in this paper. I am pleased to have the opportunity to share this experience with the EDA research community and to participate in the celebration of Bryan's ISPD Lifetime Achievement Award.

6. REFERENCES

- [1] Preas, B. T., and C. W. Gwyn. "Architecture for contemporary computer aids to generate IC mask layouts," in *Record of 11th Asilomar Conf. on Circuits, Systems and Computers*, pp. 353-361, 1977.
- [2] Preas, B. T., and Charles W. Gwyn. "Methods for hierarchical automatic layout of custom LSI circuit masks," in *Proc. of the 15th Design Automation Conf.*, pp. 206-212, 1978.
- [3] Preas, B. T., and William M. vanCleave. "Placement algorithms for arbitrarily shaped blocks." in *Proc. of the 16th Design Automation Conf.*, pp. 474-480, 1979.
- [4] Preas, B. T. "*Placement and routing algorithms for hierarchical integrated circuit layout.*" Doctoral Dissertation, Department of Electrical Engineering, Stanford University, 1979.
- [5] Preas, B. T., and C. S. Chow, "Placement and routing algorithms for topological integrated circuit layout," in *Proc. Of the Intl. Symposium on Circuits and Systems*, pp. 17-20, 1985.
- [6] Preas, B. T., "An approach to placement for rectilinear cells," presented at Physical Design Workshop: Placement and Floorplanning, Hilton Head, South Carolina, April 1987.
- [7] J. Cong, D. F. Wong and C. L. Liu. "A New Approach to the Three Layer Channel Routing,". *Int'l Conf. Computer-Aided Design*, pp. 378-381, 1987.
- [8] Madden, Patrick H. "Partitioning by iterative deletion." *International Symposium on Physical Design*. ACM, 1999.
- [9] Sechn, C., and A. Sangiovanni-Vincentelli. "TimberWolf 3.2: A new standard cell placement and global routing package." *the 23rd Design Automation Conf.* 1986.
- [10] J. Cong and Bryan Preas. "A new algorithm for standard cell global routing." *IEEE International Conference on Computer-Aided Design*, 1988.
- [11] J. Cong and Yuzheng Ding. "FlowMap: An optimal technology mapping algorithm for delay optimization in lookup-table based FPGA designs." *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 13.1 (1994): 1-12.
- [12] J. Cong and Yean-Yow Hwang. "Boolean matching for complex PLBs in LUT-based FPGAs with application to architecture evaluation." *International Symposium on Field Programmable Gate Arrays*. ACM, 1998.
- [13] J. Cong and Yean-Yow Hwang. "Boolean matching for LUT-based logic blocks with applications to architecture evaluation and technology mapping." *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, (2001): 1077-1090.
- [14] Ling, Andrew, Deshanand P. Singh, and Stephen D. Brown. "FPGA technology mapping: a study of optimality." *the 42nd Annual Design Automation Conference*. ACM, 2005.
- [15] J. Cong and Kirill Minkovich. "Improved SAT-based Boolean matching using implicants for LUT-based FPGAs." *15th International Symposium on Field Programmable Gate Arrays*. ACM, 2007.
- [16] J. Cong, John Peck, and Yuzheng Ding. "RASP: A general logic synthesis system for SRAM-based FPGAs." *Fourth International Symposium on Field Programmable Gate Arrays*. ACM, 1996.
- [17] J. Cong, Jie Wei, and Yan Zhang. "A thermal-driven floorplanning algorithm for 3D ICs." *International Conference on Computer Aided Design*, 2004.
- [18] J. Cong, et al. "Thermal-aware 3D IC placement via transformation." *Asia and South Pacific Design Automation Conference*, 2007.
- [19] J. Cong and Yan Zhang. "Thermal-driven multilevel routing for 3-D ICs." *Asia and South Pacific Design Automation Conference*, 2005.
- [20] J. Cong and Yan Zhang. "Thermal via planning for 3-D ICs." *International Conference on Computer-Aided Design*, 2005.

- [21] J. Cong, Guojie Luo, and Yiyu Shi. "Thermal-aware cell and through-silicon-via co-placement for 3D ICs." *Design Automation Conference (DAC)*, 2011.
- [22] J. Cong and Guojie Luo. "A 3D physical design flow based on Open Access." *International Conference on Communications, Circuits and Systems*, 2009.
- [23] Thorolfsson, Thorlindur, et al. "Logic-on-logic 3d integration and placement." *IEEE International 3D Systems Integration Conference (3DIC)*, 2010.
- [24] J. Cong and Zhiru Zhang. "An efficient and versatile scheduling algorithm based on SDC formulation." *the 43rd Annual Design Automation Conference*. 2006.
- [25] J. Cong and Wei Jiang. "Pattern-based behavior synthesis for FPGA resource reduction." *International Symposium on Field Programmable Gate Arrays*. ACM, 2008.
- [26] J. Cong, Bin Liu, and Zhiru Zhang. "Scheduling with soft constraints." *International Conference on Computer-Aided Design*. ACM, 2009.
- [27] J. Cong, Bin Liu, and Zhiru Zhang. "Behavior-level observability don't-cares and application to low-power behavioral synthesis." *International Symposium on Low Power Electronics and Design*. ACM, 2009.
- [28] J. Cong, et al. "Behavior-level observability analysis for operation gating in low-power behavioral synthesis." *ACM Transactions on Design Automation of Electronic Systems (TODAES)* 16.1 (2010): 4.
- [29] J. Cong, et al. "Automatic memory partitioning and scheduling for throughput and power optimization." *ACM Transactions on Design Automation of Electronic Systems (TODAES)* 16.2 (2011): 15.
- [30] J. Cong, et al. "Platform-based behavior-level and system-level synthesis." *IEEE International SOC Conference*, 2006.
- [31] J. Cong and Yi Zou. "FPGA-based hardware acceleration of lithographic aerial image simulation." *ACM Transactions on Reconfigurable Technology and Systems (TRETS)* 2.3 (2009): 17.
- [32] Wang, Yuxin, et al. "Memory partitioning for multidimensional arrays in high-level synthesis." *Proceedings of the 50th Annual Design Automation Conference*. ACM, 2013