Book Reviews

A physical-design picture book

Igor Markov
University of Michigan

Reviewed in this issue

Practical Problems in VLSI Physical Design Automation, Sung Kyu Lim, (Springer, 2008, ISBN: 978-1-4020-6626-9, 264 pp., $139.00).

FEW AUTHORITATIVE SOURCES on physical design algorithms can be found today. Monographs available in English are either hopelessly outdated or focus on narrow topics. Several edited volumes offer broad coverage by mustering chapters written by top researchers. These chapters show varying levels of detail but are overwhelmingly consistent in their lack of worked-out examples. With *Practical Problems in VLSI Physical Design Automation*, Sung Kyu Lim has addressed those problems and pioneered a new genre for an EDA book, built around a collection of exercises with solutions.

Physical design, of course, is unique in the way it affects almost every field of EDA, as well as microarchitecture and VLSI design. A computer-architect intern at a legendary US company mentioned recently that CPU designs in Verilog are sometimes returned by the back-end teams when the designs' performance fizzles during physical design (a word to the wise: watch out for multiplexer nests that lead to fatal routing congestion). System-level planning and performance estimation for ASICs involve partitioning and floorplanning, but manual floorplanning can be prohibitively labor-intensive for modern SoC designs. IP reuse in SoCs requires large-scale mixed-size placement, which was practically unknown in the 1990s. Transistor tuning by adaptive body-biasing is often supported with dedicated detail placement. In another example, industry colleagues working on FPGA logic synthesis recently obtained a 25% logic-level delay improvement with new algorithms, but saw it evaporate entirely when performance was evaluated after placement.

The ongoing breakdown of traditional design abstractions is striking, and commercial tools are only starting to address it in earnest. In the late 1990s, several industry luminaries and top executives famously proclaimed that VLSI placement "was solved," and some advocated hiding latency instead of optimizing it. Ten years later, physical synthesis tools still struggle with placement, while several leading US and Taiwan design houses indicate that available commercial EDA tool chains are lacking in placement quality, speed, and features. Top design companies have partnered with universities to build in-house tools or to articulate the need for improvement to EDA vendors. Smaller players continue waiting until commercial tools improve.

Here's an easy guess to venture. Commercial physical design tools will never provide encapsulation sufficient for competitive high-end design because of the vast number of difficult optimizations required at advanced technology nodes—most optimizations are moving targets. State-of-the-art chips will continue requiring manual optimization or script-based hacking, specific to individual chips. Whoever writes those scripts (presumably current graduate students) must understand basic algorithms for large-scale physical design. Researchers and practitioners in related fields can also benefit from understanding what makes physical design algorithms tick.

Enter *Practical Problems in VLSI Physical Design Automation,* a book that will help those who would like to get a general sense of foundational algorithms for physical design. Sung Kyu Lim's book grew out of a graduate-level course and is organized in six chapters: clustering, partitioning, floorplanning, placement,

0740-7475/09/$26.00 © 2009 IEEE Copublished by the IEEE CS and the IEEE CASS

single-net (Steiner) routing, and multi-net routing. It does not aim to be comprehensive in the selection of topics or in discussing key properties of specific algorithms. To this end, each chapter is structured around three-to-five seminal publications on a given topic—some of the papers covered go back as far as the early 1980s. Lim supplements them in four ways by providing a quick overview, solved exercises, additional exercises, and short summaries of more recent work. The focus is on helping the reader grasp the essence of each algorithm (by example) and develop relevant algorithmic skills. Numerous custom-designed illustrations—a hallmark of this book—make the examples fun and carry the reader through fairly sophisticated algorithms, which would otherwise make for some pretty dry reading. The writing is clear, although grammar and style could be improved in future editions.

To appreciate this book, it is important to keep a sense of perspective. Quick overviews of algorithms given in each chapter do not attempt to define technical terms used, and there is no discussion of the VLSI design process or typical back-end tool chains. From an algorithmic viewpoint, readers might wonder how to select an algorithm most appropriate for a given situation—a discussion of advantages and drawbacks of individual techniques could be useful here. Moreover, some algorithms covered in the book seem to be largely abandoned in practice today, despite their mathematical beauty. A case in point is spectral partitioning, which seems unable to support "fixed terminals" and is dominated in both speed and solution quality by multilevel variants of Fiduccia-Mattheyses (MLFM). Given that MLFM is only justified in practice for netlists with more than 100–200 nodes, any complete example illustrating it would look farfetched. The same applies to analytical placement algorithms, which only shine at the large scale. In fact, this is one of the reasons why competitive physical design algorithms are so hard to develop—small examples can be misleading.

I found chapters 3 and 5 most relevant: on floorplanning and Steiner routing respectively. On the other hand, chapter 6 ("Multi-net Routing") is sorely lacking negotiated-congestion routing (the Pathfinder algorithm by McMurchie and Ebeling), which is used in the majority of contemporary commercial and academic routers. Chapter 4 takes a surprising turn in its discussion on placement by listing, under a section on the GORDIAN algorithm, several entirely unrelated techniques, such as Mongrel and APlace. Chapter 1 could be reduced by limiting its focus on clustering algorithms used in multilevel partitioning and placement, and perhaps folded into chapter 2. But clock-network synthesis deserves a separate chapter—DME (deferred-merge embedding) algorithms are practically useful and lend themselves naturally to illustration by small examples.

THIS BOOK IS A SIGNIFICANT step forward in that it presents a new type of educational material for students of physical design and experts from related fields. Some readers will view it a glass half-full, and some a glass half-empty. Perhaps it is also a harbinger of a new approach to physical design education and an open invitation to authors in other EDA fields. ■

■ Direct questions and comments about this department to Igor Markov, 2260 Hayward St., Ann Arbor, MI 48109-2121; imarkov@eecs.umich.edu.