

PANEL: PHYSICAL DESIGN AND SYNTHESIS: MERGE OR DIE !

Organizer and Chair: Massoud Pedram, University of Southern California, Los Angeles CA

Abstract

As IC fabrication capabilities extend down to sub-half-micron, the significance of interconnect delay and power dissipation can no longer be ignored. Existing enhancements to synthesis and physical design tools (such as non-linear delay modeling, custom wire load models, back annotation of calculated delays, early floorplanning, post-layout re-mapping and resizing) have not been able to solve the problem. It thus remains that tradeoffs in logical and physical domains must be addressed in an integrated fashion. Huge business opportunities will be lost unless more revolutionary changes to design flow are made.

This panel of experts will address the current split between logic synthesis and physical design and its effect on the design flow. It will then discuss possibilities for merging the two, or at least bringing them closer together. In particular, issues such as consistent wire load and timing models and algorithms which must be employed across the design flow, EDA standards and common databases to support the integration of layout and synthesis tools, evolving structured design styles that offer lower wiring overhead, interconnect-driven logic synthesis, and timing-driven physical design will be discussed. Finally, the panel will seek to highlight challenges and potential pitfalls that lie ahead.

Position Statements

Richard Bushroe
SEMATECH, Austin TX

Physical Design and Synthesis absolutely must be merged. But, this also must be expanded to include a merge with System Architectural Design at the front-end and interconnect parasitics extraction, delay calculations, transistor level design and analysis, and interface with TCAD generated models at the back-end. Design technologies must be advanced for design flows, methodologies, tools, standards to be able to produce chip designs with 100 million or more transistors with the same number of designers in the same time it takes today for a 5 million transistor chip. Without these advances, the semiconductor and electronics industries will die economically as they fall off the productivity curve (25-30% improvement in \$ per function) that the world has come to expect year-after-year.

Synthesis for these complex high performance ASIC's and structured custom chips must optimize for the now-dominant interconnect delays. Where the number of nets that may be at a performance risk for 0.5 to 0.25 micron 5 million transistor chips may have been 10's of thousands, the number of critical nets for 0.18 to 0.13 micron chips may be in the millions. Therefore, the semiconductor and electronics industries will die unless the EDA industry provides for a merged design technology from system architecture design through interconnect parasitics and includes integration of all aspects of performance driven physical design at every level, not just synthesis.

The resulting design technologies must also address, across the spectrum of the design process, power estimation and optimization, performance-driven system architectural design/planning/synthesis, interconnect analysis, signal integrity verification and optimization for total interconnect across the entire chip/system continuum, design-for-test, verification and validation (with increased use of formal methods). Also, EDA standards that support plug-and-play tool integration defined by an integrated design information-model and high-performance database that supports large geographically distributed design teams must be addressed to achieve cost effective solutions for the semiconductor and electronics industry.

Raul Camposano
Synopsys, Inc., Mountain View CA

The prevalent ASIC (standard cell or gate array) design methodology is based on synthesis of an RTL description, placement and routing, extraction, back-annotation and physical verification. This methodology becomes less feasible as so called "deep-sub micron" effects grow stronger. The design flow will change including several new elements. For example, hierarchical floor planning integrates estimation (logic, coarse placement, congestion analysis, global routing), logic synthesis and layout.

A flexible, integrated chip-level multi-level static timing analysis tool allows consistent timing calculations and time budgeting. More accurate (often called "3D or 2.5D") extraction and back annotation allow more precise timing and power analysis. We can expect profound changes in design methodology. The silicon industry will need to re-tool to cope with advancing technology and design, resulting in new challenges and opportunities which are likely to change the EDA industry.

Giovanni De Micheli
Stanford University, Stanford CA

The move to deep sub-micron technologies, where interconnect delays dominate, will foster a revolution in design styles, methods and tools. Library-based design is likely to decline, in favor of macro-cell approaches where module generators synthesize fairly complex combinational and sequential components. Such generators will control the layout topology, and therefore accurate delay/power predictor will be able to steer the simultaneous optimization of transistor configuration, placement and wiring. As a result, new paradigms for logic synthesis integrated with layout will become dominant.

Antun Domic
Cadence Design Systems, San Jose CA

The introduction of smaller geometry CMOS technologies is forcing us to re-evaluate the established design flows. In particular, the increasing dominance of interconnect delay (versus gate delay) in the total path delay is bringing the traditional separation of logic design (“logic synthesis”) and physical design (“place and route”) into scrutiny. The mismatches between predicted delay after synthesis and actual delays after placement and routing are causing excessive design iterations, and this situation is expected to worsen.

Can we then merge some of these steps? If not into routing, can we at least consider logic synthesis and placement in one step? Unfortunately, the algorithms in use in these tools are very different, including the basic data they use. Synthesis radically changes the circuit’s gate level implementation as it proceeds from boolean minimization to technology mapping, using a relatively simple model for wire delays. Placement, on the other hand, does not change the gate implementation, and is estimating net lengths by tree approximations based on physical location (plus other extremely important metrics even though less mentioned, such as routability).

So, unless new algorithms are developed soon, our best hope in the near future is to use “minor” re-synthesis based on placement data, and other very specific optimizations. In many cases, using only steps such as gate re-sizing, buffer tree modifications, and small logic changes, the basic original placement can be preserved substantially. This makes it possible for the effect of the changes to actually take place (through good overlap removal techniques). Evidence is appearing that these techniques can be very successful in fixing the timing and load related violations left after timing driven placement. If this holds up, the perceived need for combining synthesis and placement will be postponed.

Clearly, the success of specific optimizations will require a very clean flow among many tools. Use of tools such as floorplanners should help here. The transparency of the

data transfer, including the appropriate constraints, will be crucial for the success of such tools in production.

Chi-Ping Hsu
Avant! Corp. Sunnyvale CA

With the deep submicron effects from complicated parasitics of physical design dominating design performance, it is clear that the synthesis has to be done within the physical design process.

The traditional term of “synthesis” consists of two stages, i.e. language translation, and technology optimization. The technology optimization needs to be deeply embedded within physical design process, not only working with actual placement but also global routing and detailed routing processes. The impact of timing at later stages of physical design has as much impact as the placement alone due to complex routing parasitics. The deep submicron design paradigm will require “deep submicron synthesis” inside physical design tools and RTL floorplanning for early architectural decision with area/power/timing feedbacks instead of the “traditional synthesis” paradigm that has been used for over 20 years.

Michael Jackson
Motorola Inc., Austin TX

Over the next five years, future CAD flows will need to change to address the growing demands of deep submicron design (DSM). Classical logic synthesis followed by place and route will not converge fast enough on tomorrow’s high frequency and aggressive product development cycle time designs. There will be the need to introduce new technology that essentially consolidates the logic synthesis with layout tasks. As the percentage of interconnect to total delay increases, the value of stand alone synthesis and place and route is undermined until it is of zero value.

Five key technologies will be central to the success of tomorrow’s CAD flows:

- floorplanning technology that better integrates the estimation and analysis of logical, physical, timing, and power representations of a design.
- a merged synthesis and place and route technology that performs concurrent placement and technology mapping.
- dynamic libraries that better utilize the potential of the silicon.
- faster and more accurate RLC extraction.
- consistent timing and power analysis across all tools and the support for fast and incremental chip-wide analysis.