

Fusion Compiler Unified Physical Synthesis

Unified Optimization Technologies For Advanced Node Designs

Authors

Akira Nikaido

Arvind Narayanan

Directors of Product Marketing,
Synopsys

Introduction

The semiconductor industry has been pushing the design envelope over the past few decades and reached an unparalleled level of innovation and success. In this quest for excellence, designers have to overcome a wide range of technological challenges and time-to-market pressure. With continual shrinking of process geometries and the ever increasing pace of technological adoption, designers are facing a myriad of new and complex design- and process-related challenges that must be overcome to break new frontiers. These new era designs are aggressively pushing chip performance, power, and area (PPA) as well as time-to-market targets, making the traditional RTL-to-GDSII design flow ineffective.

Smaller device nodes and the introduction of FinFET with multi-patterning have a significant impact on the implementation flow starting with floorplanning and placement. The new design rules for these advanced process nodes must be considered earlier during the synthesis phase and should be applied throughout the rest of the design flow including post-route optimization and chip finishing. Designing for the highest frequency targets while maintaining the lowest power continues to be a difficult objective for the designer to meet. Technologies for driving the maximum performance and the lowest power need to be utilized from architecture to RTL physical synthesis and through the backend flow. Stringent advanced node design rules requiring higher design utilization while reducing silicon area have become critical to reduce mask costs. Leading-edge process technologies and designs require a unified implementation environment that breaks the barriers of design phases and sharing of technologies and engines throughout the design flow.

Fusion Compiler RTL-to-GDSII System

Fusion Compiler™ is the next-generation RTL-to-GDSII implementation system architected to address the complexities of advanced node designs and deliver up to 20% improved PPA while reducing time-to-results (TTR) by 2X. Fusion Compiler is built on a compact, scalable, single data model that allows seamless sharing of Fusion Technology across the RTL-to-GDSII flow to enable hyper-convergent design closure. Fusion Compiler has been built from the ground-up using best-in-class RTL synthesis, place-and-route, and industry-standard golden signoff technologies for designing system-on-chips (SoCs). The tool delivers unmatched quality-of-results (QoR) for PPA while delivering the fastest design convergence. Fusion Compiler provides a complete RTL-to-GDSII design system including RTL physical synthesis, design planning, placement, clock tree synthesis (CTS), advanced routing, unified physical synthesis-based optimization, chip finishing, signoff quality analysis and ECO optimization.

Unified Physical Synthesis Solution

Fusion Compiler offers a solution that spans both RTL synthesis and place-and-route domains including interleaved floorplanning and synthesis, incremental compile, physically-aware data path representation, logic re-synthesis during physical implementation and a common Unified Physical Synthesis (UPS) optimization engine that delivers unprecedented QoR and design convergence. UPS is the nerve center and fundamental backbone for all optimization capabilities within Fusion Compiler (Figure 1). UPS offers advanced RTL physical synthesis technologies such as logic re-synthesis, layer-aware optimization, macro skewing and timing-driven buffering and sizing. UPS also leverages key place-and-route technologies such as a next-generation placer, advanced 2D legalizer, concurrent clock and data optimization (CCD), multi-bit banking and de-banking, total power-centric optimization, automatic non-default rules (NDR), and via ladder support for advanced nodes. UPS also utilizes a common infrastructure for placer, legalizer, timing analysis, and CCD across the flow leading to improved correlation and convergence. UPS offers the ultimate flexibility to the designer to go from RTL to coarse placement or directly to CTS using a single command from the same cockpit.




Figure 1: UPS delivers best QoR by combining best technologies from synthesis and place-and-route

To address the needs of leading-edge designs, UPS offers a slew of technologies that target best QoR, hyper-convergent design flow, and design cycle times. Some of the unique technologies include:

- Fusion data model architecture provides a scalable single database utilized by all implementation and signoff engines
- Advanced placement and 2D legalization algorithm provides improved design rule check (DRC) count, better pin access, optimal white space utilization, and faster design closure
- RTL synthesis with CCD-aware clock estimation for accurate clock structure prediction early in the design flow minimizes area and leakage, and maximizes performance
- Solver-driven, path-based optimization engine utilizing the same costing infrastructure for pre- and post-route optimization for best correlation and concurrent worst negative slack (WNS), total negative slack (TNS), power, and area optimization
- Pin density-aware placement for advanced process nodes improves pin access and maximizes design utilization and routability
- Wire and interconnect synthesis with automatic layer assignment and NDR generation to predict multi-routing layer assignment from synthesis and forward annotation to place-and-route
- New congestion prediction algorithm during synthesis for accurate prediction and avoidance early in the design flow
- Signoff-quality common timer, delay calculation, and RC extraction engine integration for signoff closure throughout RTL-to-GDSII flow

Advanced Placement and Legalization

Complex advanced node design rules must be applied early in RTL physical synthesis and during placement, legalization, CTS, and optimization to ensure that lithography requirements are met and to enable a smooth convergent flow. For example, the new inbound cells (turbo cells) that help boost timing have a unique transistor structure with half row offset that should be managed during placement and legalization. Similarly, support for optimal selection of via pillars early in the flow is critical for meeting performance requirements and DRC convergence.

UPS offers new advanced legalization and placement engines that are common between the synthesis and place-and-route domains. Sharing the same engines tightens correlation between the two domains and delivers significantly better design convergence (Figure 2). Congestion, area, and timing predictions made early in the design cycle are consistent through the flow, thereby minimizing unexpected surprises at the very end of the flow. During synthesis, the UPS advanced legalizer handles inbound cells and regular standard cells as well as multi-bit registers while considering the trade-offs to meet timing and area constraints. The advanced legalizer can handle multiple routing layers (>10), support via pillars for both boosting performance and providing robustness for electro-migration, and handle cross row dependencies.




Figure 2: Perfect congestion correlation between synthesis and place-and-route

With the introduction of FinFETs, standard cells have become very compact causing high pin density and pin access challenges. The synthesis engine makes optimal choices of cells to minimize pin density hotspots and reduce routing congestion downstream during physical implementation. The placement engine is architected to support the fin grid, spacing rules for implants, continuous diffusion, variable threshold cell spacing rules and cross row rules to minimize multi-patterning violations. The placer also considers pin density congestion and coloring during placement, avoids placing standard cells below power straps, and utilizes horizontal and vertical pin alignment to improve pin access and routability (Figure 3).




Figure 3: Pin grid (PG)-aware placement and pin alignment to improve routability

New Optimization Engine

To achieve lowest power, best performance, and smallest area, UPS has a completely revamped optimization engine. The solver-based global optimization engine enables path-based TNS targets for effective design closure. The same costing and infrastructure are used for both pre- and post-route optimization to enable consistent correlation through the flow. Solver-based optimization targets all endpoints with negative slack, then traces and identifies every cell on the critical path as an optimization candidate. It computes the timing, area, and power trade-offs for each candidate in parallel, selects the optimal set, and concurrently minimizes the objective functions. Path-based optimization (Figure 4) traces and identifies one cell on each critical path to optimize TNS, and handles all scenarios, path groups and violating endpoints for effective timing convergence. Both solver-based and path-based optimization runs in parallel by distributing the candidates on multiple threads for fast turnaround time. New pre-route RC extraction statistical modeling enables accurate RC modeling during synthesis and early placement stages for tight correlation with the post-route stage.




Figure 4: Exhaustive path-based optimization significantly reduces TNS

Total Power Optimization

UPS offers total power optimization that identifies the right trade-off between leakage and dynamic power reduction while maintaining other design metrics. The solver-based algorithm delivers fast convergence by optimizing across the entire multi-dimensional solution space and finding the global minima. The solution avoids getting stuck in local minima, by tunneling through to other parts of the solution space to get out of local minimum. The flow also performs knee-based power optimization where the algorithm does not implement small delay improvements at the expense of huge power spending.

UPS power optimization technologies include power-aware CCD, power-driven restructuring, multi-bit inference, IR drop-driven placement, UPF-based multi-voltage design, Dynamic voltage frequency scaling (DVFS) and advanced multi-Vth flows. Cell libraries for leading-edge processes have multiple Vth selection and variance of channel widths. The optimization engine should choose the right combination between the different Vth classes and channel widths to reduce total power. Figure 5 illustrates how Fusion Compiler considers both dynamic and leakage power concurrently to identify the best candidate.




Figure 5: Total power optimization provides best trade-off between dynamic and leakage power reduction

Concurrent Clock Data Throughout the Flow

CCD optimization is traditionally performed after the CTS stage as clock latency is essential for useful skew optimization. In a conventional approach early in the synthesis stage, clocks are treated as ideal, leading to all sinks having the same clock delay. Designers manually tweak critical clock delays by setting clock latency to solve this issue. For large SoC designs this technique can become very tedious. The UPS infrastructure enables CCD optimization to be fully integrated with synthesis, providing accurate clock structure estimation to improve PPA early in the design cycle. This technology builds a virtual clock tree and observes entire endpoint slacks and arrives at positive slack distribution to minimize area and power simultaneously. Figure 6 describes how CCD enabled during synthesis stage performs slack distribution and optimizes the clock structure to improve timing QoR. CCD during synthesis is completely automated and forward-annotates CTS offset information for the place-and-route CTS engine to build efficient clock trees. CCD during physical implementation is enabled throughout the flow including place opt, clock opt and post-route opt stages.




Figure 6: Concurrent clock data optimization during synthesis to improve PPA

Fusion Data Model

Fusion Compiler is built on a single data model and contains both logical and physical information to enable sharing of library, data, constraints, and design intent throughout the implementation flow. The compact and scalable advanced data model is architected to support ultra-high capacity designs with the smallest possible memory footprint. Fusion Compiler integrates all synthesis, place-and-route as well as signoff engines on the single data model, eliminating data transfer to deliver the fastest design closure with highest throughput.

Unified Physical Synthesis Results

Figure 7 illustrates the QoR and runtime advantage of utilizing UPS technologies over the traditional design flow for different design sizes and process nodes. The comparison is done after legalized placement.




Figure 7: UPS QoR and runtime advantage over traditional design flow

Summary

Fusion Compiler RTL-to-GDSII implementation system addresses the time-to-market pressure and delivers the QoR required for leading-edge designs. UPS in Fusion Compiler seamlessly shares technologies and common engines between synthesis and place-and-route domains to deliver the best PPA in the shortest time.