

# FPGA SYNTHESIS: LOOKING BEYOND THE OBVIOUS

November 2008

### ABSTRACT

For successful FPGA implementation, getting synthesis right can often determine whether or not requirements are met and the product can ship. In many cases, design requirements refer to performance and area—the design needs to operate at a minimum frequency and naturally needs to fit into the selected device. Hence, designers or CAD managers looking to standardize on a synthesis flow tend to look for good out-of-the-box quality-of-results (QoR).

The criteria for selecting the right synthesis tool, however, should involve more than this. To meet aggressive QoR goals, constraints need to be refined, optimizations must be enabled, or small portions of RTL may need to be re-coded—but without help from the tool, it is difficult to identify when and where to make these changes. In other cases, QoR goals may have been met but design changes are constantly being introduced, and new changes either degrade previous QoR results or runtime for each change delays project schedule. FPGA project managers must take into account these scenarios and consider how their synthesis flow addresses them.

### Author:

Ehab Mohsen  
Technical Marketing Engineer

## INTRODUCTION

When it comes to successful FPGA implementation, getting synthesis right can often determine whether or not requirements are met and the product can ship. In many cases, design requirements refer to performance and area—the design needs to operate at a minimum frequency and naturally needs to fit into the selected device. Hence, designers or CAD managers looking to standardize on a synthesis flow tend to look for good out-of-the-box quality-of-results (QoR).

The criteria for selecting the right synthesis tool, however, can—and should—involve more than this. Adequate out-of-the-box QoR is a legitimate starting point, but it is not the whole story in terms of a successful FPGA design methodology. Often, to meet aggressive QoR goals, constraints need to be refined, optimizations must be enabled, or small portions of RTL may need to be re-coded—but without help from the tool it is difficult to identify when and where to make these changes. In other cases QoR goals may have been met but design changes are constantly being introduced, and new changes either degrade previous QoR results or runtime for each change delays project schedule. FPGA project managers must take into account these scenarios and consider how their synthesis flow addresses them.

## QUALITY-OF-RESULTS

Still, a tool's ability to produce adequate out-of-the-box quality-of-results is an important criterion and not to be downplayed. If the FPGA can not operate at the required frequency, the product simply cannot ship. And if it cannot fit into the selected device, having to upgrade to a larger device may be prohibitively expensive. It is no surprise that a common reason why users decide to move to 3rd party EDA synthesis tools is because they are not meeting the desired QoR requirements with their free vendor tools.

While there are many features contributing to QoR such as advanced technology inference, retiming, and resource sharing, one capability to be considered carefully is the tool's physical synthesis support. Physical synthesis, as opposed to standard RTL synthesis, is a method of synthesizing RTL based on the physical




Figure 1: Traditional wire load delay model treats these interconnect delays equally

characteristics of the target FPGA device.

Regular RTL synthesis only takes into account logic cell delays and wire load timing models. Figure 2 shows two interconnect delays that would be treated equally by regular RTL synthesis (given the same wire-load) but in fact would vary in final place-and-route due to placement.

Physical synthesis, in contrast, produces an optimized netlist by taking into account placement or potential placement of logic on the device and routing resources and performs . Different vendors have different approaches to physical synthesis, but the goal is the same—improve design performance—particularly for high-end devices or designs with aggressive performance goals.

Precision<sup>®</sup> RTL Plus applies a type of physical synthesis known as physically aware synthesis, a pre-place-and-route approach that estimates placement, thereby calculating more accurate routing delays, and performs optimizations based on this information to generate a timing-improved netlist, as shown in Figure 2. This approach has empirically been proven to be effective and scalable to many devices.

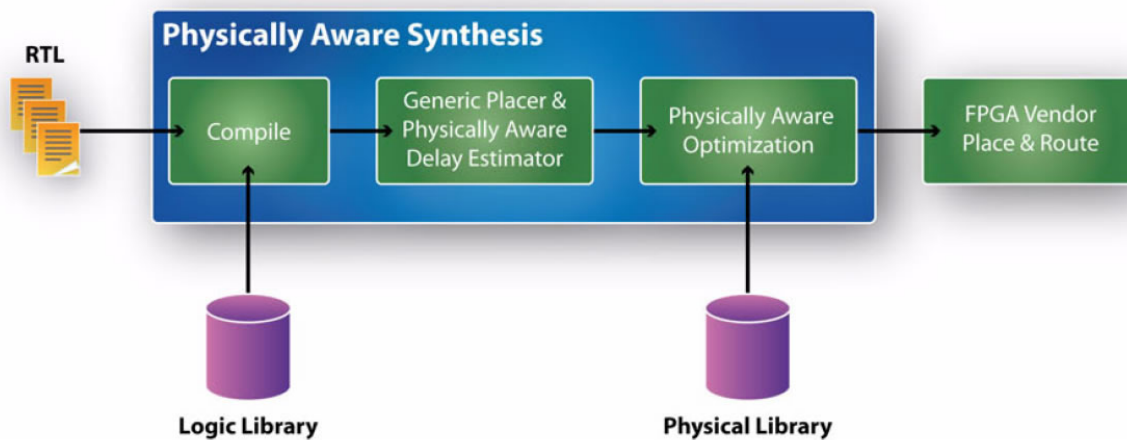


Figure 2: Physically Aware Synthesis in Precision RTL Plus

Scalability is important because an impressive optimization technology has limited use if it is only available for a few devices. Even if it supports the FPGA of the current project, a designer typically wants the freedom to switch devices or FPGA families for the next project and still have all the optimization technologies at his or her disposal. As of this writing, Precision’s physically aware synthesis supports 23 devices from all major FPGA vendors, three times the device support of the nearest competitor.

Ease-of-use is also another factor. Traditional methods of FPGA physical synthesis require expert knowledge of the chip, but many users may not want to invest time acquiring such device knowledge. Precision’s physically aware synthesis has been designed to maintain a push-button flow making it accessible to all users.

## DESIGN ANALYSIS

As design complexity continues to increase, design analysis remains a critical component of the synthesis flow. In an ideal scenario, a design is imported into the tool, a button is pushed, and a netlist is generated that meets all requirements—and no analysis is needed. Unfortunately, this scenario is neither common nor realistic. Few engineers will get the design and timing constraints right the first time. For aggressive QoR goals, synthesis trade-offs may need to be explored by the user. A good tool has the right analysis features to help analyze constraints, analyze the design itself, and guide the user through architectural trade-offs to meet timing or area goals. Descriptive timing reports, graphical views of critical paths, and the ability to cross-probe to relevant HDL are examples of the analysis features needed to understand QoR bottlenecks.

A surprising “gap” in many solutions is the ability to graphically analyze and control allocation of embedded memory or DSP resources on the device. Embedded resources are specialized areas on the FPGA that can map certain operators more efficiently for either performance, area, or both. Resource allocation has been available in most synthesis flows for many years but only via script commands or HDL attributes within the design code. Allocation results are only viewable after synthesis in the area report, and resource allocation changes require a complete re-synthesis.

The ability to graphically analyze how RTL functions will be mapped—before synthesis—can provide a desperately needed advantage to designers as it allows resource allocation changes to be made early and intuitively. Precision RTL Plus provides such a capability with Resource Manager. Figure 3 shows Resource

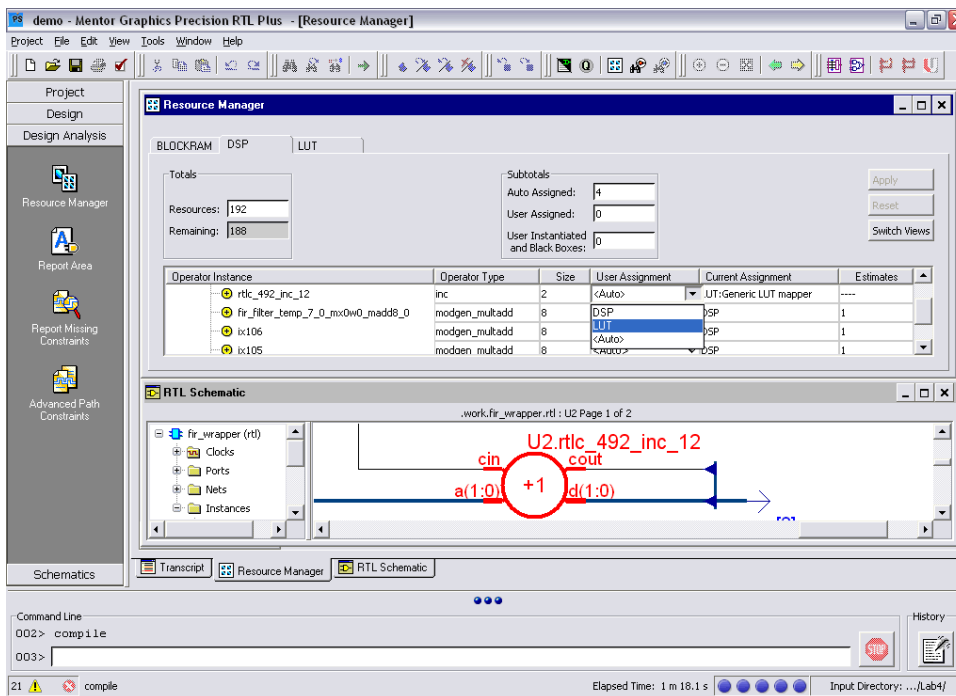


Figure 3: Graphical Resource Analysis

Manager displaying graphical view of all operators in a design that can be potentially mapped to embedded memory or DSP resources. Equally important is the ability to cross probe from such a view to the relevant design schematic or HDL source.

In fact, the ability to cross-probe to and from different graphical and textual views is foundational to a good analysis environment. A user should be able to cross-probe from generated reports, to the post-synthesis technology schematic, to the generic RTL schematic, to the relevant HDL code to understand all issues in their appropriate contexts.

### INCREMENTAL FLOWS

Even when QoR goals are met, they risk being affected by subsequent design iterations. Time spent meeting performance goals on critical blocks may be lost due to changes in other areas of the design. This and the time consumed from long synthesis and place-and-route run-times slow down productivity and risk project schedule. For this reason, a tool's incremental flow is critical.

There are a variety of incremental flows available, and some of these require pre-planning, such as the block-based partition flow. Perhaps the ideal solution to look for is a fully automatic incremental synthesis flow, where the user does not need to identify logic areas that may potentially be changed in future runs. Equally important is that the flow not prevent full optimization of the design to obtain best results. Precision RTL Plus provides a fully automatic synthesis flow where re-implementation is localized to the affected design areas. Coupled with an incremental place-and-route flow (when available), as shown in Figure 4, time savings can be achieved and QoR results of previous runs can be preserved.

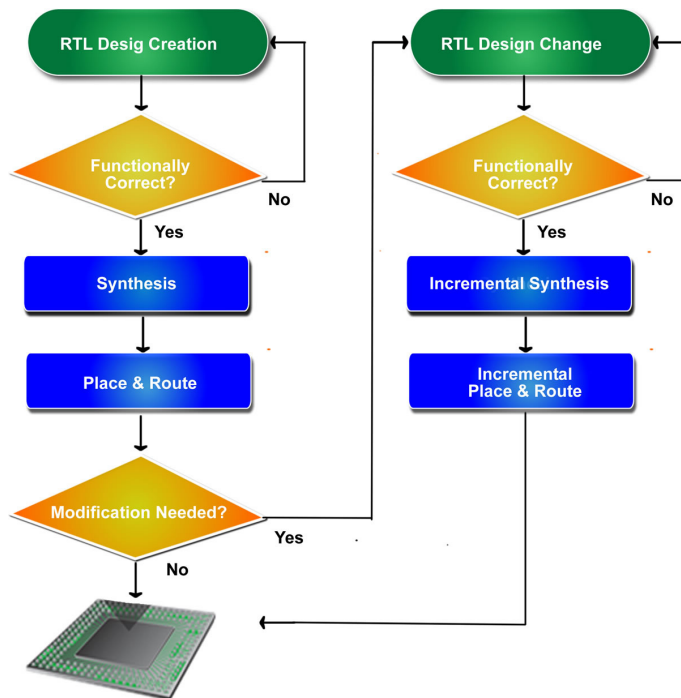


Figure 4: Incremental Synthesis & Place-and-Route Flow

---

## CONCLUSION

There are naturally other criteria to consider for FPGA synthesis, depending on the application. The point to remember is that selecting the right tool is not a one-dimensional exercise. Even when performance or area is the highest priority, capabilities such as analysis and incremental support can directly relate to quality-of-results. Designers and managers should realistically consider the various stages of their RTL-to-FPGA-implementation flow and select the tool that best supports each stage.

**Visit the Mentor Graphics website at [www.mentor.com/fpga](http://www.mentor.com/fpga) for the latest product information.**

©2008 Mentor Graphics Corporation. All rights reserved.

Mentor Graphics and Precision are registered trademarks of Mentor Graphics Corporation. All other trademarks are the property of their respective owners. This document contains information that is proprietary to Mentor Graphics Corporation and may be duplicated in whole or in part by the original recipient for internal business purposes only, provided that this entire notice appears in all copies. In accepting this document, the recipient agrees to make every reasonable effort to prevent unauthorized use of this information.

---

**Corporate Headquarters**  
Mentor Graphics Corporation  
8005 SW Boeckman Road  
Wilsonville, OR 97070-7777  
Phone: 503.685.7000  
Fax: 503.685.1204  
**Sales and Product Information**  
Phone: 800.547.3000

**Silicon Valley**  
Mentor Graphics Corporation  
1001 Ridder Park Drive  
San Jose, California 95131 USA  
Phone: 408.436.1500  
Fax: 408.436.1501  
**North American Support Center**  
Phone: 800.547.4303

**Europe**  
Mentor Graphics  
Deutschland GmbH  
Arnulfstrasse 201  
80634 Munich  
Germany  
Phone: +49.89.57096.0  
Fax: +49.89.57096.400

**Pacific Rim**  
Mentor Graphics (Taiwan)  
Room 1001, 10F  
International Trade Building  
No. 333, Section 1, Keelung Road  
Taipei, Taiwan, ROC  
Phone: 886.2.87252000  
Fax: 886.2.27576027

**Japan**  
Mentor Graphics Japan Co., Ltd.  
Gotenyama Garden  
7-35, Kita-Shinagawa 4-chome  
Shinagawa-Ku, Tokyo 140-0001  
Japan  
Phone: 81.3.5488.3033  
Fax: 81.3.5488.3004

