sponsored by

Gabe on EDA • EDAmarket • EDA Confidential

Assembling the Future

A Newsletter About the Design and Production of Electronics

ISSUE 003 • MARCH 2011

In this issue:

RTL Sign-off: Is It Realistic?

Gabe Moretti

A couple of years after the commercial introduction of logic synthesis the industry began to address the possibility of RTL signoff. Practically twenty years later we are still talking about it. It is time to ask: is RTL signoff ever going to be a reality? If not are the tools inadequate? Does the RTL netlist contain enough information to complete the design process? There is a common answer to these questions: it depends.

Just like no circuit is the same, whether or not a RTL description can be processed into a gate level netlist suitable for place and route by a third party depends on the nature of the design. Whether it is solely digital or mixed/signal, the product execution environment, issues like power consumption and thermal limits, and the manufacturing process characteristics are determining factors in whether or not a RTL description of the design is sufficient. In general, designers must consider all these factors before deciding on the design method and flow to follow.

Feasibility Of RTL Sign-off

There are certainly many designs that can be signed off at the RTL level. Designs that are fully digital and that do not target the latest manufacturing process, or designs that target FPGA devices fall in this category. Logic synthesis and physical synthesis tools are advanced enough to be reliably counted on to generate accurate gate level netlist descriptions. Place and route tools, whether part of the foundry's own tools suite or from a trusted third party, are also capable of producing accurate input to the manufacturing process.

Yet, there are practically no designs that are signed-off at the RTL. The reason is that the description of a circuit done with a Hardware Description Language (HDL) does not contain enough information for a foundry to guide synthesis and place and route tools correctly. More information is required. The missing information is in some cases tool dependant, and in others introduced by designers later in the development cycle.

For example in the case of power consumption there are tools that allow architects and designers to evaluate the total power required by a specific circuit architecture. The modular nature of HDLs would allow the assignment of specific power characteristics to various portions of the design and given the proper characteristics of a logic simulator provide designers with the information required to describe a power distribution network without having to wait for a gate level topology of the circuit. Yet, such capabilities seem to be missing.

The process of designing and manufacturing a circuit proceeds from one description to the next by eliminating abstractions and replacing them with specific parameters and implementation descriptions. The greater the amount of information about the physical and behavioral characteristics of a circuit that can be defined at a higher level of abstraction, the less feedback work that will be required if errors or limitations are found later on. As complexity of designs increases, so must the level of abstraction at which the design can be signed-off. Not to do so implies increasing the cost of development to the point that many custom IC designs are, or soon will be, so expensive as to make the resulting product unprofitable.

It also helps if designers are not required to push the manufacturing process to its limits in order to implement a circuit that meets specifications. Although this is not always possible, it is very prudent to allow margins between what the circuit requires and what physical characteristics a process can implement. Designing to the limit is an engineering characteristic but, especially when implementing products destined for the consumer market, it is prudent to allow for some margin of execution reliability and lower fabrication costs.

I believe that every design that can be implemented in an FPGA device is a potential candidate for RTL signoff. But changes are needed in the design methods.

Improved Design Methodology And Verification Tools

Designers are increasingly using algorithmic languages, like C/C++, SystemC, and MATLAB to describe the circuit functionality and relying on High Level Synthesis (HLS) tools to derive either HDL descriptions or, in some cases, gate level circuit netlists. Although the acceptance of HLS tools was slower than expected, the quality of results from their use is now acceptable, and it continues to improve, as would be expected, with use. More can and will be done in this area now that vendor are assured of a growing market. Tools like Forte's Cynthesizer that give users the option of producing either an HDL or a gate level netlist provide the required flexibility to designers.

The use of algorithmic languages allows engineers to verify the correctness and accuracy of their implementation of a specific algorithm, thus diminishing the need for logic simulation later on.

The growth in IP blocks availability is also simplifying design verification. Unfortunately the quality of commercial IP is not uniform. This is a major problem not only for engineers that must integrate the

IP block in the design, but also for verification engineers that must insure that the block functions as expected in a specific design. True, tools like Cadence's ChipEstimate InCyte (www.chipestimate.com), try to provide a way to choose and qualify IP blocks, but there is really no standard way to evaluate IP, in spite of the IEEE-1685 , also known as the IP-XACT, standard.

Two significant standards, TLM and UVM are providing a much needed structure to architectural design and verification respectively. TLM not only allows the integration of hardware IP blocks but also supports hardware/software co-design. Accellera just released UVM, and there was significant interest in the methodology at the just ended DVCon conference. The overt, and even enthusiastic, support for the method by all three major EDA vendors bodes well for the standard.

Formal Verification is another growing market segment. Here again the adoption of the tools was slow. But improvements in both user interfaces and size of design capabilities have widened the use of formal verification, both for equivalence checking as well as for formal verification through the use of assertions.

Equivalence checking tools can allow designers to verify that a gate level circuit is functionally equivalent to its HDL description for example, and assertions not only simplify design verification, but can be used by designers to turn an implementation into a functional specification of the design.

Assertion based verification has the potential of increasing the level of abstraction at which analog circuits can be verified. Two papers delivered at the just concluded DVCon conference, one by Texas Instruments and the other by Cadence, show the successful use of assertion to verify analog circuits.

There Will Always Be Exceptions

In electronic design, as in most things in life, there is never absolute uniformity of methods. Thus there will always be designs that must be signed-off at abstraction levels lower than RTL. Custom analog circuits, for example, and mixed/signal circuits with a high percentage of analog components, could continue to require sign-off at gate level, or maybe even GDS-II level. But analog IP, once certified and properly documented, can be used in a design as a black box component. It can then be instantiated at RTL and should not produce any problem to the sign-off process.

There are also products that for competitive reasons will require their developers to push the limit of fabrication technology in order to achieve the highest possible performance level for a given process. Graphics chips, CPUs, and network controllers are examples of such devices. These devices are practically handcrafted at the cell level and today even require the intervention of developers in the process of producing the diffusion masks required for manufacturing.

Foundries Are Increasing The Level Of Hand-off Abstraction

On the other hand, the message from semiconductor manufacturers is that they foresee at 20 nm and below the need to exercise more control in the fabrication process. During the Common Platform event held in Santa Clara earlier this year, Dr. Gary Patton of IBM said that when using the 20 nm process, his company would only take a GDS-II description from a customer and would reserve any further processing for itself. The reason: the complexity of the process. When asked to speculate about the next node, 14 nm, he declined to go into details but pointed out that there will

be new material used in the process, and thus the level of abstraction of the description of the circuit required by the foundry may be higher still.

With every new processing node designers face an increasing number of design rules that must be met for the design to be fabricated with acceptable yields. In spite of the fact that EDA vendors are releasing more powerful tools with every process node to manage design rules, the set is so large and complicated that verification is becoming too expensive for many products. To decrease the number of rules and their complexity requires increasing the level of abstraction at which the design is turned over to the foundry. This allows the foundry to take advantage of methods that implement the circuit with less probability of breaking a rule.

The way a custom design is placed and routed, then, becomes more and more similar to how a design is implemented in an FPGA device. Perfectionists would observe that the final layout is not as efficient as it could have been. But if the product meets the specifications and can be fabricated for the right price, then the company has met its goal: going to market with a product capable of generating a profit.