

IP quality lies beyond compliance testing

OF COURSE YOU WANT YOUR STANDARD-INTERFACE IP TO PASS COMPLIANCE TESTING. BUT THAT ACCOMPLISHMENT IS JUST THE BEGINNING. COMPLETE QUALITY ASSURANCE FOR IP CORES HAS FAR MORE CHALLENGES.

Compliance testing can reveal a great deal about whether an IP (intellectual-property) core functions according to the standard specifications, but does it really answer the most important question in the IP business, which is, Will passing compliance testing be enough to ensure that the IP will perform as expected in your chip? Long-term experience in providing IP to chip designers says that the answer to this question is “not quite.”

Compliance testing alone cannot guarantee the quality or performance of IP for a number of reasons. This article, using the PCIe (Peripheral Component Interconnect Express) interface as an example, discusses the verification steps that engineers use in ensuring that the IP is of the highest quality. This process includes ensuring functional correctness, testing across corners, validating in silicon, passing compliance, and selecting an IP with advanced, built-in capabilities. All of these steps require a great deal of effort, specifically from the IP vendor.

THE VALUE OF COMPLIANCE TESTING

Given the amount of standards-based IP on SOCs (systems on chips), the ability to comply with standard specifications has become crucial. IP that complies with standards dominates the market. These standards can be de facto, such as microprocessor architectures, including ARM (www.arm.com); interconnect standards, such as USB (Universal Serial Bus), PCIe, DDR (double-data rate), SATA (serial-advanced-technology attachment), and Ethernet; or standard IC functions, such as standard cells, memories, and I/Os. The ability to use standards is the cornerstone of the IP industry (Figure 1). Standards-based IP enables designers to use plug-and-play techniques to take full advantage of the millions of transistors that advanced process-technology nodes provide.

Compliance testing has thus become vitally important for determining whether the standards-based IP functions as the specifications define and to confirm its interoperability with other devices. To take its standards-based IP through a compliance program, the vendor must have a silicon implementation of the IP in an appropriate test system. This test system should create a series of tests based on the specifications,

including electrical specifications, such as a standard eye diagram (Figure 2).

Unfortunately, passing compliance and interoperability testing does not guarantee that the IP will work in your chip. The reasons range from basic functional issues to fabrication variations that the testing cannot verify. For example, compliance tests cannot check all of the functions the specification defines; they check only limited operating scenarios, and, specifically for IP, the tests do not include all configurations. Thus, compliance testing may confirm that a selected single-lane PCIe implementation passes the compliance and interoperability tests and works according to portions of the specification as defined by the standards organizations, but this test doesn't confirm that different configuration options built into the IP work.

Even when a version of the IP meets the standard, the tests do not capture its degree of robustness relative to the specification. This factor is especially important for mixed-signal IP—particularly for today's high-speed SERDES (serializer/deserializer) PHY (physical)-layer IP. For example, a PCIe PHY-IP core can pass the compliance tests by barely falling within the standard eye diagram and can have such a weak signal that it is less likely to work in your design than a PHY-IP core that provides plenty of margin relative to the specification.

That margin is necessary to accommodate the wide range of design-in, packaging, and board-level factors that tend to degrade signals. The ideal environment of compliance tests does

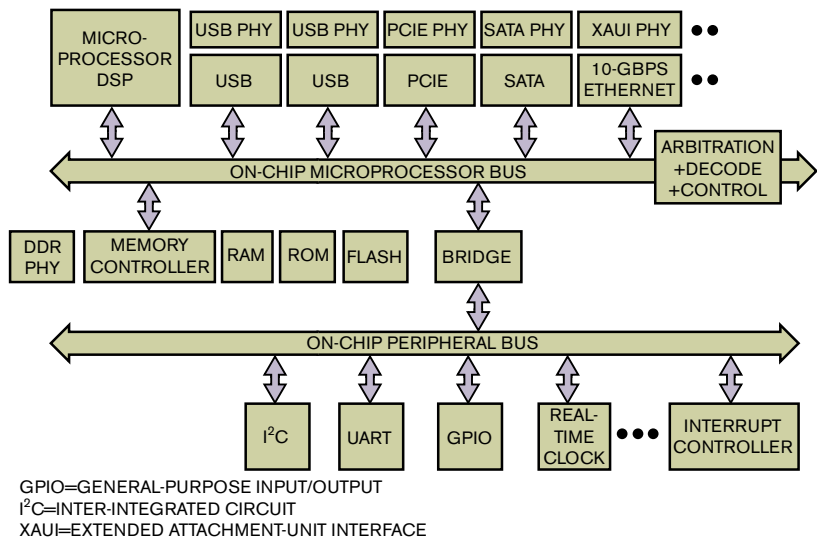


Figure 1 Standards-based IP may make up a significant portion of an SOC design.

not take these degradations into account. Moreover, the silicon for compliance tests may represent a beneficial combination of process corners, whereas your chips may tend toward corner combinations that can lead to problems. Manufacturing yield and reliability issues may result in the creation of many chips that you cannot use.

To address issues that may arise during its implementation into the overall system environment, the IP must go through much more than mere compliance testing. Ensuring high-quality IP involves a verification methodology that should consist of comprehensive simulation techniques, silicon validation, and proven interoperability. All of these factors should be a part of the IP vendor's development process.




Figure 2 An appropriate test fixture is necessary for IP-compliance testing.

START WITH GOOD IP DESIGN

Experienced designers often say that you must design quality into a product from the beginning. So it is worth taking a look at how a vendor develops its IP and how experience says that the IP should be developed. Using PCIe as an example, a front-end-development flow may include many verification steps, starting with a verification plan early in the design (Figure 3). Rather than a separate subflow, verification is an integral part of the flow that influences the IP design at every step.

The design practices the flow uses are also critically important, and the best practices of design for reuse should be readily available. Reference 1 details one resource, which lists best practices in reuse and SOC design based on the collective

insights of Alcatel (www.alcatel-lucent.com), ARM, Atmel (www.atmel.com), Cadence (www.cadence.com), IBM (www.ibm.com), Infineon (www.infineon.com), LSI (www.lsi.com), Mentor Graphics (www.mentor.com), Philips (www.philips.com), STMicroelectronics (www.st.com), Synopsys (www.synopsys.com), Texas Instruments (www.ti.com), and several universities. The book encourages IP developers to use these best practices in their designs so that the IP is more likely to perform as they expect.

DIGITAL-IP VERIFICATION

Verification and validation are interesting to examine for both synthesizable digital IP and mixed-signal PHY IP. A good

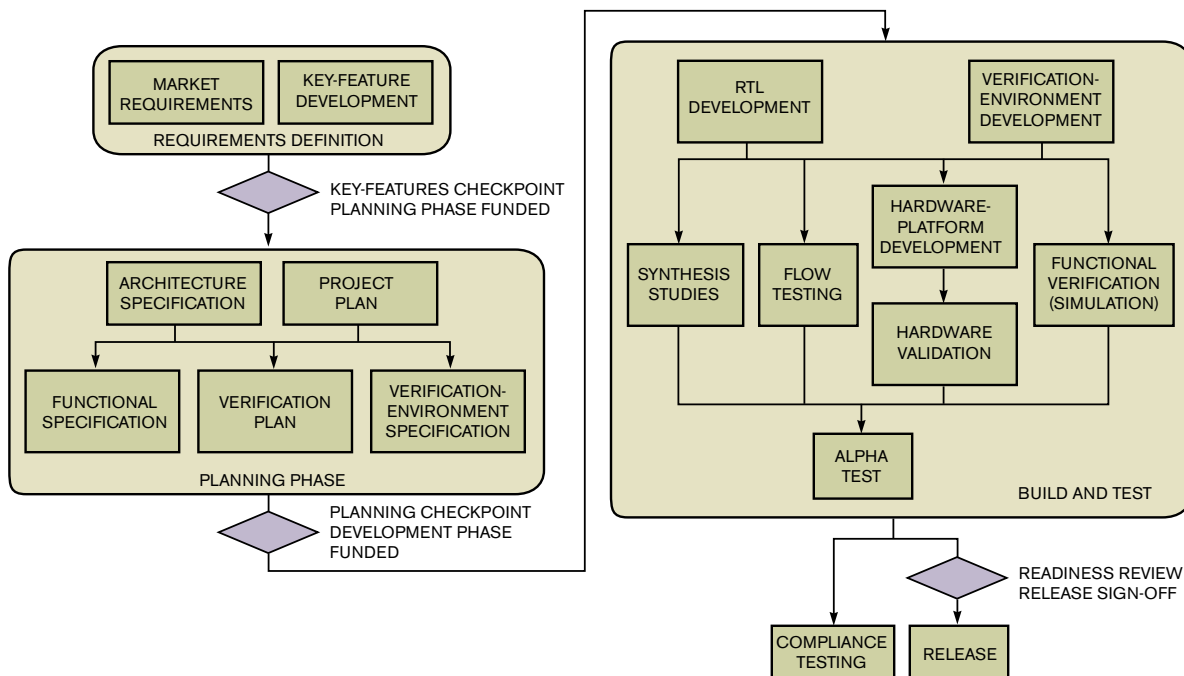


Figure 3 Verification may be complex, and it may start early in the design flow.

process for digital IP should begin with the generation of a functional-coverage check list (Figure 4). Verification engineers derive the check list from an interpretation of the protocol, design, and system-interface specifications, and this check list, in turn, drives the verification process.

Ensuring functional correctness typically requires a combination of techniques, such as simulation, writing assertions, emulation, and hardware prototyping. During simulation, verification engineers should conduct a series of random and directed tests to ensure maximum code coverage. The directed tests target functions within the blocks, including compliance-coverage tests, but often do not hit all of the functions of the device. Using constrained random verification techniques and functional-coverage points the verification team derives from the specifications, the team can methodically drive the IP in random simulations until it is fully verified, at almost all the corner conditions. With this approach, you can trace the verification-coverage points back to the functional-coverage check list and the primary specifications, and you can generate reports of progress against the specifications. The functional-coverage check list and report are key indicators of functional correctness, and they help ensure that the IP will work in your design.

Using directed and random verification testing, Synopsys routinely runs 35 billion to 40 billion cycles of simulations per day on IP-core regressions. This work requires the company to have approximately 500 CPUs running regressions all the time. That process goes a long way toward ensuring correct functioning of the IP. It also ensures that customers, no matter what tool chain they use, will get the same results the IP provider does. Consequently, the full validation process includes thorough flow-testing of the IP with individual EDA tools, as well as testing with major simulators from multiple vendors.

Even with billions of cycles of simulation, the design still needs to work in real-world applications. For example, even though each PCIe interface originates from the same specifica-

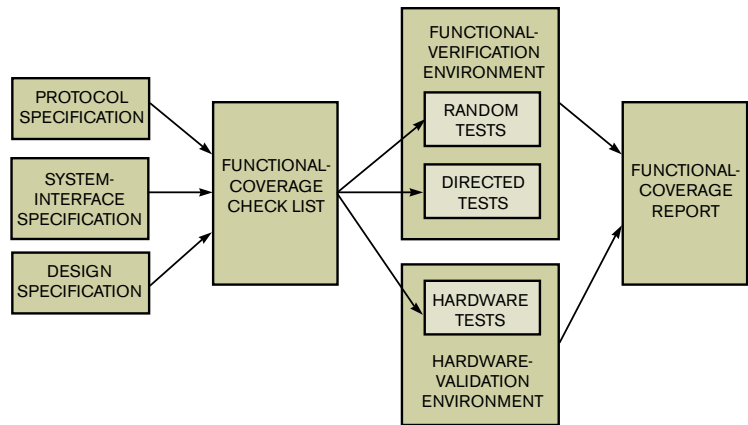


Figure 4 A functional-coverage check list informs the rest of the verification flow.

tion, some areas of the specification are not completely defined, or designers can interpret them differently. This problem is generally what leads to interoperability issues. An IP provider must test the IP by placing it into hardware, usually through an FPGA-prototyping board. Synopsys generally uses the same FPGA prototype for both hardware verification and compliance testing, so the FPGA implementation enables additional testing and running of interoperability tests, such as “plugfests,” with other PCIe systems and available test equipment.

Unfortunately, it is difficult for an IP vendor to include every possible configuration of a flexible IP product into the FPGA prototype. So it is best to ask the IP provider not only what configurations it has tested in FPGAs but also what configurations its customers have placed into silicon and taken into volume production. Another customer may have already taken the configuration you need through compliance testing.

MIXED-SIGNAL-IP VERIFICATION

Although synthesizable digital IP offers plenty of verification challenges, challenges also arise when verifying analog/mixed-signal PHY IP. A look at a PCIe PHY at process nodes of 65

nm or smaller shows some fascinating trends, many of which involve physical-design dependencies.

One trend is the occurrence of variations due to changes in device size. Predicting these variations requires Monte Carlo simulations of threshold voltage and drain-to-source saturation current mismatches, and these simulations, in turn, demand large amounts of computation time and many tool licenses. Many other variations, including systematic variations due to physical effects, such as STI (shallow-trench-isolation) stress, WPE (well-proximity effect), and contact stress, also occur. The IP may also suffer from time-dependent variations due to NBTI (negative-bias temperature instability), HCI (hot-carrier injection), and EM (electromigra-

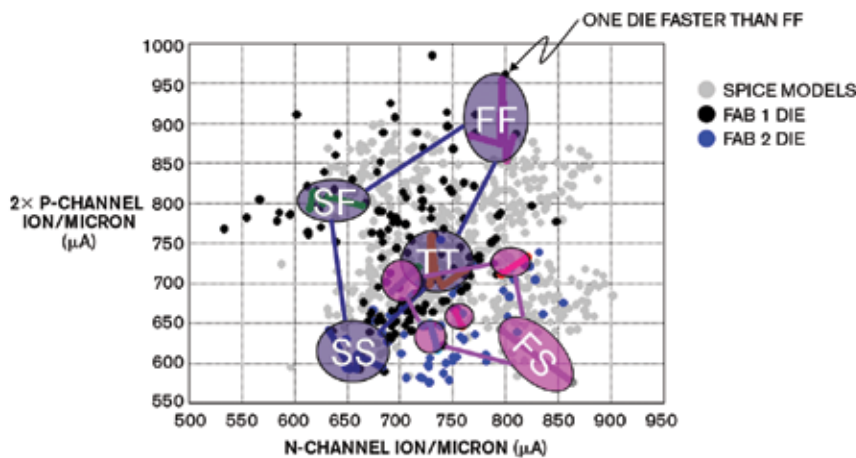


Figure 5 An on-chip process monitor can allow you to pick dice representative of each process corner.

tion). In the face of all of these variations, simulations that use prelayout parasitic-extraction parameters are becoming poor predictors of performance, simply because performance increasingly depends on physical factors that are subject to large layout-dependent variations. Postlayout parasitic-extraction effort is therefore increasing.

All of these variations in effect create process corners. But just as Monte Carlo simulations absorb computing time, simulating mixed-signal IP across process corners is an excellent way to consume huge amounts of computational resources. The simulation to verify the DesignWare PHY for PCIe, for example, covers 14 basic analog blocks. On average, the simulations require 18 test benches per block for a total of 250 test benches. Vendors must run each test bench across process, voltage, and temperature corners, including IR (current/resistance)-drop effects. Each test bench also includes independent corners for thin- and thick-gate devices, polysilicon resistors, and capacitance, as well as other variables. When you add up all of these considerations, each test bench runs across an average of 90 corners. This part of the PCIe PHY-layer verification process thus requires approximately 23,000 simulations.

Manufacturers may also do Monte Carlo simulation in addition to traditional corner analysis to assess yield and design centering. Altogether, the typical total simulation time for a DesignWare PHY layer for PCIe is approximately 3.25 CPU years. Following simulation is reliability analysis. Some companies analyze EM for both power and signal nets in postlayout parasitic-extraction simulation.

BUILT-IN INSURANCE

Because the performance of a high-speed PHY-layer IP core relies on a complex mix of difficult-to-predict factors, success may require more than just exhaustive analysis. You may need some new techniques for mixed-signal IP. For example, you can use an on-chip process monitor and an on-chip sampling scope in each die. The scope shows the on-chip signal from the PHY-layer chips, so you can determine its performance independently of packaging and board effects. The on-chip process monitor allows you to find out what corner combination you have on a die—information that is otherwise difficult to obtain. If you compare data from Spice models with on-chip monitor data for chips from two fabs, you can see where a die lies in the process (Figure 5). When you have this information, you can select dice representative of each process corner. You can predict yields and see where the center of the process is relative to Spice.

This ability enables more meaningful jitter analysis (Figure 6). Because jitter limits yield, chips must be better than the specification defines. If all your dice end up at the slow/slow corner, will your design still be OK? Without the on-chip monitor, it's virtually impossible to tell.

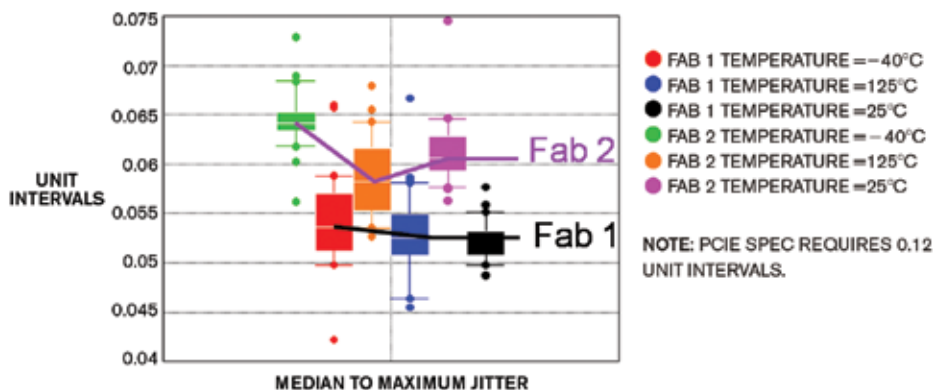


Figure 6 A process monitor can help visualize the range of jitter performance for dice.

Even before good Spice models are available, you can use the information from the on-chip monitor to take better advantage of programmability in the analog blocks. In a PHY-layer chip, for instance, you may be able to adjust the dc level, the crossover point, and the rise and fall characteristic to make the block faster or slower. Pre-emphasis is useful, but only if you can implement it without causing ringing. Knowing each die's corner combination from the on-chip process monitor enables you to accurately and efficiently set the programmable PHY-IP characteristics.

PLAYING SAFE WITH IP

You are highly likely to achieve working silicon using IP that was designed and verified using the best design practices and methods. In addition, incorporating advanced functions, such as the on-chip abilities for mixed-signal IP, provides a clear means of differentiating IP from various vendors in the market. Although passing compliance certification is a requirement, it is not sufficient to guarantee IP quality. It is also mandatory to examine a vendor's verification/validation and design practices. For digital IP, a functional verification plan and check list are keys. For analog- and mixed-signal IP, a development method for robust designs and margins, with a clear path to high yields, is essential. **EDN**

REFERENCE

1 Keating, Michael, and Pierre Bricaud, *Reuse Methodology Manual for System-on-a-Chip Designs*, Third Edition, Kluwer Academic Publishers, 2002.

AUTHOR'S BIOGRAPHY



Navraj Nandra joined Synopsys in February 2005 as director of product marketing for the mixed-signal products that include SERDES, USB, and DDR2. He has worked in the semiconductor industry since the mid-'80s as an analog/mixed-signal-IC designer for Philips Semiconductors, austriamicrosystems, and EM-Marin. Before joining Synopsys, Nandra was director of application engineering at Barcelona Design. He holds a master's degree in microelectronics, with a focus on analog-IC design, from Brunel University (London) and a postgraduate diploma in process technology from Middlesex University (London). For information on Synopsys' DesignWare IP, visit www.synopsys.com/designware.