


sponsored by

Gabe on EDA • EDAMarket



# Assembling the Future

A Newsletter About the Design  
and Production of Electronics

ISSUE 013 • MARCH 2012

**SUBSCRIBE**

PDF and Archives

twitter

## In this issue:

- **IP Is Core To System Design**  
by Gabe Moretti
- **The Business of IP**  
by Josh Lee
- **Verifying Antifuse NVM Hard Macro IP to Be Bulletproof**  
by David Hsu
- **Discovery Verification IP**  
by Neill Mullinger

## IP Is Core To System Design

**Gabe Moretti**

In January of 2011 I began the publication of a newsletter "Assembling the Future". With its March 2012 issue all readers of EDACafe can now access the newsletter for free. Each month I will write a blog that describes the contents of that particular edition. Just follow the link provided by EDACafe. If you loose it then use this one: <http://www.gabeoneda.com/newsletter>.

This month's topic is Intellectual Property (IP) and its necessary companion verification IP (VIP). The first article written by Josh Lee of Uniquify presents a short history of the IP business, from its wild west origins to the still developing present days. The second article, by David Hsu of Kilopass describes the relative large VIP package necessary to verify non-volatile memory (NVM) had macro IP. Finally Neill Mullinger of Synopsys introduces its company new VIP family based on the VIPER architecture.

Next month's topic is EDA standards. How they are developed, why we need them, and what is on the horizon. To contribute an article contact me at [gmoretti@gabeoneda.com](mailto:gmoretti@gabeoneda.com).

### **A few thoughts about IP**

The IP market needs not just IP cores, but also VIP functions and procedures, a way to find products, and standard procedures both commercial and technical.

The beginning of the IP market practically coincides with the commercial availability of logic synthesis. In the late eighties and early nineties, Synopsys introduced a package of synthesizable basic logic modules called Designware while at the same time HDL Systems also started selling synthesizable cores written in both Verilog and VHDL. HDL Systems later was assimilated by Philips semiconductors, victim of managerial incompetence. But Synopsys persisted and leveraging its large installed base of Design Compiler grew its IP business to a very lucrative and growing business. Its latest offering is based on the VIPER architecture described in this month's issue of the newsletter.

Denali, now part of Cadence, was the pioneer of the VIP business. When Sanjay Srivastava started the company no one paid much attention: there was nothing "sexy" about memories. But no system can exist without memory and interface protocols continued, and continue, to evolve to more complex systems. Testing the correct functioning of the system became expensive, and Denali's business grew and grew creating credibility for other companies to find funding in the VIP market. One of such examples of memory VIP is described in the article by Kilopass in the newsletter.

But there are more parts of the IP market that are either not well covered in the press or are not thought of as being "IP" per se. I am referring to the embedded software market. At a time when we have reached the realization that software and hardware are to be considered equally in architecting a system, embedded software remains a market in its own. But it is really IP that must have its own VIP to be of commercial quality.

I find it peculiar that the IP market and the embedded market are not learning from each other and establish commercial protocols and engineering practices that gain from both experiences. It is high time to stop thinking about computers and programs, or APPS as application programs are now called and start to think of systems.

Although in addition to Cadence, Mentor, and Synopsys there are some large sources of IP like ARM, Xilinx, Altera, and MIPS for example, there are also many smaller vendors. A designer does not always have either the knowledge or the time to find the IP core that would best fit his requirements. A service like that offered by ChipEstimate, now part of Cadence, is a solution. Although ChipEstimate sells its own products that ease the evaluation of IP cores and the integration of cores in a design, they also provide a free catalogue of IP that is believed, of course they cannot guarantee it, to be of commercial grade.

Finally the market needs standards. A few exist, like the IEEE 1685 IP-XACT, but more are needed. The march of semiconductor manufacturing technology continues and with it come more and more stringent design rules that are approximating the limit of second sourcing and porting a design to a different foundry. Foundry specific IP and VIP is approaching and with it a possible revolution in the market. yet I find very few indications that any one is seriously considering how to

enable and support this transition. The market seems in denial hoping that business as usual can continue in spite of technology's requirements. If we do not act now to prepare for the general use of IP at 20nm and 14 nm we will witness a significant slow down in the transition to those processes with grave financial consequences not just for the electronics industry but for the world economy. Let's not forget that without electronics products, nothing gets done anymore.

---

## The Business of IP

### Josh Lee

**President and CEO, Uniquify Inc.**

It wasn't all that long ago when starting an IP business was all the rage. As the concept of SoC design began to take off, it created demand for IP blocks that could deliver common functions needed in most of these chips. Savvy engineers quit their day jobs and, with powerful PCs, started designing pieces of important and differentiated intellectual property in the comforts of their own homes. They were in the IP business. Those were heady days, indeed, and profitable for some. It was much like panning for gold during the gold rush of the 1840s.

Then, reality set in and semiconductor companies started to understand that purchasing IP from a third party was a bit like the American Wild West — seemingly, lawlessness prevailed. Nothing was illegal, just a small matter of product quality and reliability and the recurring challenge of integrating the IP into each new design. The lack of standardization also became an issue. SoC designs in the 1980s and 1990s had few common IP blocks or protocols. It was apparent that some kind of law and order was in order.

While many companies needed to rely on third-party IP to get their projects completed on time and within budget, they couldn't afford to take the risk of working with small, unproven vendors with unproven IP. These buyers were often large fabless semiconductor companies, electronic system manufacturers, semiconductor foundries, IDMs, ASIC companies and even EDA vendors, all of whom had customers relying on their high standards and reputations to protect.

Today, the IP Wild West is far more civilized and structured, while chip complexity continues to grow and the use of IP increases. We are now witnessing the emergence of the SuperSoC (SSoC), essentially a conglomeration of multiple SoCs, as they used to be defined, all interconnected on the same die. The IP content of many of SSoCs can be 50-60% or more. The rapid growth and intense global competition for electronic consumer devices is driving this demand for IP content as it holds the promise of enabling new SSoC designs to get to market faster.

In response to this demand, the IP business is slowly maturing. The semiconductor industry is beginning to converge around protocols and standards for the consumer and enterprise markets with well-defined standards for USB, PCIE, SATA, DDR, networking protocol, crypto blocks, and even some analog functions. Over the past few years, there has also been a move toward

consolidation with many of the independent IP providers being acquired by larger players in the semiconductor or EDA markets.

But, even more is needed. IP providers have realized that in addition to designing IP, they must also ensure that the IP is fully verified and can be easily integrated. Without these, companies with the best IP in the world will find themselves having a tough time lining up customers.

Many of the remaining third-party IP providers have stepped in to help. They have realized that they need to expand their services to offer a more complete solution to address the needs of these sophisticated buyers. After all, IP is not the end goal. The end goal is a working chip that adds value to an electronic system.

Some have invested in developing a larger portfolio of high-performance, silicon-proven IP building blocks. They are offering more of a solutions-based approach that may include ASIC/SoC design and IP integration and customization, in addition to their core business of high-performance and production-tested IP.

Still others have built experienced teams with expertise in design, IP design, qualification and integration and manufacturing operations. These kinds of offerings give a semiconductor company the ability to purchase a complete turnkey solution that spans chip specification, design and IP integration through the volume delivery of packaged and tested parts.

The madness of the gold rush may be over and the Wild West of the early days of the IP segment long gone, but the creativity and forward thinking remain. What's changed is the business of IP has become a real business with a solutions-based approach encapsulating the ever-more important pieces of IP. Many more companies are emerging with sensible and solid business plans to support the business of IP, and are being successful at implementing effective business models.

---

## Verifying Antifuse NVM Hard Macro IP to Be Bulletproof

**David Hsu**

**Senior Field Marketing and Applications Manager, Kilopass Technology, Inc.**

According to article in Design & Reuse from Cadence Design Systems entitled "The role of Verification IP in Complex core Design," the author Saverio Fazzari, writes, "general studies suggest that on average 60 to 70% of development schedules are taken up with verification." He calculates that, "in a hypothetical IC design case with a schedule of 1,000 man hours, 600 man hours, or 15 weeks, would be allocated to the verification task." It is no wonder that third-party IP verification needs to be bulletproof.

To ensure a hard core non-volatile memory (NVM) IP is bulletproof, a complete set of deliverables is provided to the design to cover every possible contingency that might arise. For an antifuse NVM IP

hard core, the deliverable will include the following files: GDSII module, GDSII Frame, LEF/Antenna LEF (Layout Extraction Format) and LVS (layout-versus-schematic) netlists, Verilog models, and .lib and .db models. The Verilog simulation model will verify the IP block's functionality; library model ".lib" file, a text file; ".db" binary file will verify timing, power, and connectivity for all the corner cases of a design: voltage, temperature, and process variations. Accompanying these data file is a set of documents including design integration guidelines, system design guidelines, data sheet, and test design methodology.

The antenna LEF file is a crucial verification model ensuring manufacturability of a chip's interconnect. Metal interconnect that's too long can build an unacceptable level of negative or positive static electricity. If the level is sufficient to produce an electrostatic discharge (ESD), it possible the transistor tied to the metal interconnect will be damaged or destroyed. Using via breaks that route the trace between layer 1, layer 2, layer 3, etc., the place and route tool reduces the continuous length of metal capable of producing a discharge able to damage a transistor. This length is specified in the process design kit (PDK) rules supplied from the foundry. The metal interconnect within the hard core must be measured with the external metal interconnect to ensure the maximum ESD length is not exceeded.

In addition to the files and documentation, there are three production tests that must be verified before the device tapes out. The first is "BLANK CHECK." The second is called "TESTDEC." The third is "WRITE TEST." These three test functions must be built into the chip during design, otherwise tests cannot be performed once the memory is in silicon. Built in self-test (BIST) and built in self-repair (BISR) are soft cores that are designed into the SoC around the IP vendor's memory block.

Blank check is an NVM IP incoming yield test, which determines that there are no defects in the transistors comprising the memory array—that all the transistors are in the "0" state. During verification, the SoC designer runs blank check simulation to verify the logic and state machine that interface the memory and the memory array itself properly execute a BLANK CHECK command. In the course of the simulation, an all zeros response from the memory array will be checked as well as the error condition of detecting a random faulty "1" bit.

No fab achieves a 100 percent yield and occasional a faulty "1" is detected during final test on actual die. With smaller process geometries, the cost of discarding die is more costly than at larger geometry processes. Instead of discarding the faulty die, it is far more cost effective to repair it using a small additional array of spare memory cells. If a memory cell fails the BLANK CHECK test, the defective cell can be mapped out and replaced with a functioning cell from the spare memory array. Documentation and a wrapper from the IP vendor detail how to perform the memory repair. The wrapper comprises two functions within the NVM IP hard core: built in self test (BIST) and built in self repair (BISR). The design team verifies these functions as part of the larger SoC verification.

While BLANK CHECK tests the memory's bit cells, TESTDEC tests the memory's wiring, the bit lines and word lines that select cells to be read and that carry the data to and from memory. Test checks for an open—a constant "0", a short—a constant "1" or a bridge—two lines shorted together. There is no repair function for a TESTDEC failure.

The last test is WRITE TEST, which tests read/write ability of the memory using an on-board charge pump or the SoC's supplied Vpp voltage. In addition to the spare memory used to repair a faulty bit

cell found during BLANK CHECK, a spare memory block of 256 bits of memory is set aside for WRITE TEST. During WRITE TEST mode, data is written and read to and from the spare memory block to ensure the memory is functional. During verification, this test function is also verified.

In Summary, IP verification is much more than testing the ability to read and write to an NVM IP memory block. It entails verifying all of the additional circuits incorporated in the IP block to facilitate test and provide repair. An IP block that provides bulletproof verification will contain all of the elements to ensure first time successful silicon.

---

## **Discovery Verification IP**

### **A New Generation of VIP to Address the Growing Challenge of Complex Protocol and SoC Verification**

**Neill Mullinger**

**Synopsys, Inc**

The role of Verification intellectual property (VIP) has become increasingly important over recent years as a vital component in achieving SoC verification productivity and is now established as an essential part of any verification solution. As the number of protocols and their complexity has increased, the demands on verification engineers to achieve productive verification and debug has also increased. This white paper discusses the limitations of protocol verification using first-generation verification IP and explains how Discovery VIP addresses these challenges.

#### **Device Convergence Driving Use of Standard Protocols**

Consumer demand for electronic products with more features, performance, capacity, and battery life has consistently driven the development of more efficient communications interfaces. The USB3.0 protocol, for example, offers 10x the performance of its predecessors. A whole family of MIPI (Mobile Industry Processor Interfaces) protocols has emerged and is rapidly evolving to serve the needs of the mobile market. On-chip buses like AMBA are also evolving with the addition of cache coherency extensions in the fabric as a means to improve the performance of multicore SoCs.

As communications standards have changed so rapidly, designing and supporting in-house protocol-based IP requires a huge commitment of resources and expertise. Buying design IP from credible vendors, rather than designing in-house, has significant value for engineering teams and has become one of the fastest-growing market segments in electronic design. Digital IP providers offer pre-designed, pre-qualified, “off the shelf” interface IP, which leaves design teams to focus their own design resources on the value-added parts of their products.

It's been a very different story for verification teams. Like design IP, developing in-house verification

IP is time-consuming and requires expert knowledge of both the protocol and testbench methodology. For this reason, there is a strong trend for companies to use commercial verification IP because:

- a) it saves the time and expertise of building and supporting in-house VIP, and
- b) if purchased from a credible vendor it will be industry proven.

However; much more so than design IP, Verification IP still requires a great deal of protocol expertise; verification engineers need to know specifications in enough detail to develop test plans, interpret coverage data, write new tests and debug errors to successfully complete verification. With the increase in complexity and number of IPs used in a typical design, protocol expertise has become a significant challenge and is becoming a critical factor in meeting schedules.

In addition to the expertise gap, the number of VIPs occupies an increasingly large percentage of the testbench size: there can easily be more than 10 VIPs and 3 million lines of VIP code in a typical SoC testbench. Performance of the VIP has become a critical factor affecting overall performance.

### **The Limitations of 1st Generation Verification IP**

Most of today's commercial VIP is based on a jigsaw puzzle of interconnected languages and methodologies. There is typically a base model developed using languages like C, e, or OpenVera that provides the underlying protocol behavior. Then, translation wrappers enable design teams to use the VIP in different languages and methodologies of the testbench, like SV, UVM or OVM. Supporting multiple languages and methodologies can often result in multiple layers, and the need to call secondary simulation engines for the underlying VIP. This layered approach creates limitations in performance and methodology. For example, mapping between environments can result in the translation of one SV object at the top-level SystemVerilog testbench interface to 10 or more commands at the VIP interface below. With today's protocols running at 500MHz and sampling on both clock edges, translating the language and methodology for each VIP in the testbench becomes a significant drag on overall performance, which either delays the time to results or demands high computer resource and more tool licenses.

Wrappers will typically do a reasonable job of mapping APIs to support the chosen methodology. Theoretically, a wrapper can provide all the translation that objects, messaging, callbacks, notifications, coverage data and error injection needed to provide full methodology support based on an underlying VIP written in a different language or methodology. This assumes that the underlying model also has the ability to support all of the messaging, notifications, coverage data and error injection; however the more comprehensive the wrapper becomes, the bigger the drag on performance. VIP vendors have to make this trade-off between methodology support and performance. Often, more code leads to more bugs.

Another productivity issue is that mixed-language VIP often requires multi-step compilation. The verification team has to compile the underlying VIP in a different step than the testbench and wrapper.

### **New Requirements for Protocol-Centric Verification**

Today's standard protocols have the flexibility to meet the price, performance and functional needs of many applications and products. For example, designers can often specify protocol speeds, number of lanes, interfaces definition, quality of service (QoS) features and so on. Consequently, there are many possible design configurations and verification engineers need to be able to configure VIP to precisely match the design specification. There can be literally hundreds of configurations and parameters present in an advanced VIP that enable this precise configuration, and, for the verification engineer, understanding which ones need to be set to which values can be another very time-consuming and error-prone step that delays time to first test.

Another task that is related to protocol configuration is defining a test plan. What needs to be tested? When is verification complete? This is another arduous task that requires analysis of the specification to extract the key coverage points that relate to the configuration used in the design in order to determine what needs to be tested during the verification phase. The test plan will change depending on whether the verification team is verifying a core or verifying the integration of a core into the design.

When it comes to the need for protocol expertise, there can be few more demanding areas than debug. Debug has become increasingly difficult with the multi-layered, highly interleaved protocols that are in use today. Most protocol debug is based around analyzing trace signals from simulator's log files, messages and transcripts. These help engineers to understand the lower-level events taking place in the traffic, but it can be extremely time consuming to cross-reference the various sources of information to understand what has been happening at the higher, more human-readable protocol-level. It requires a great deal of time and deep knowledge of the protocol.

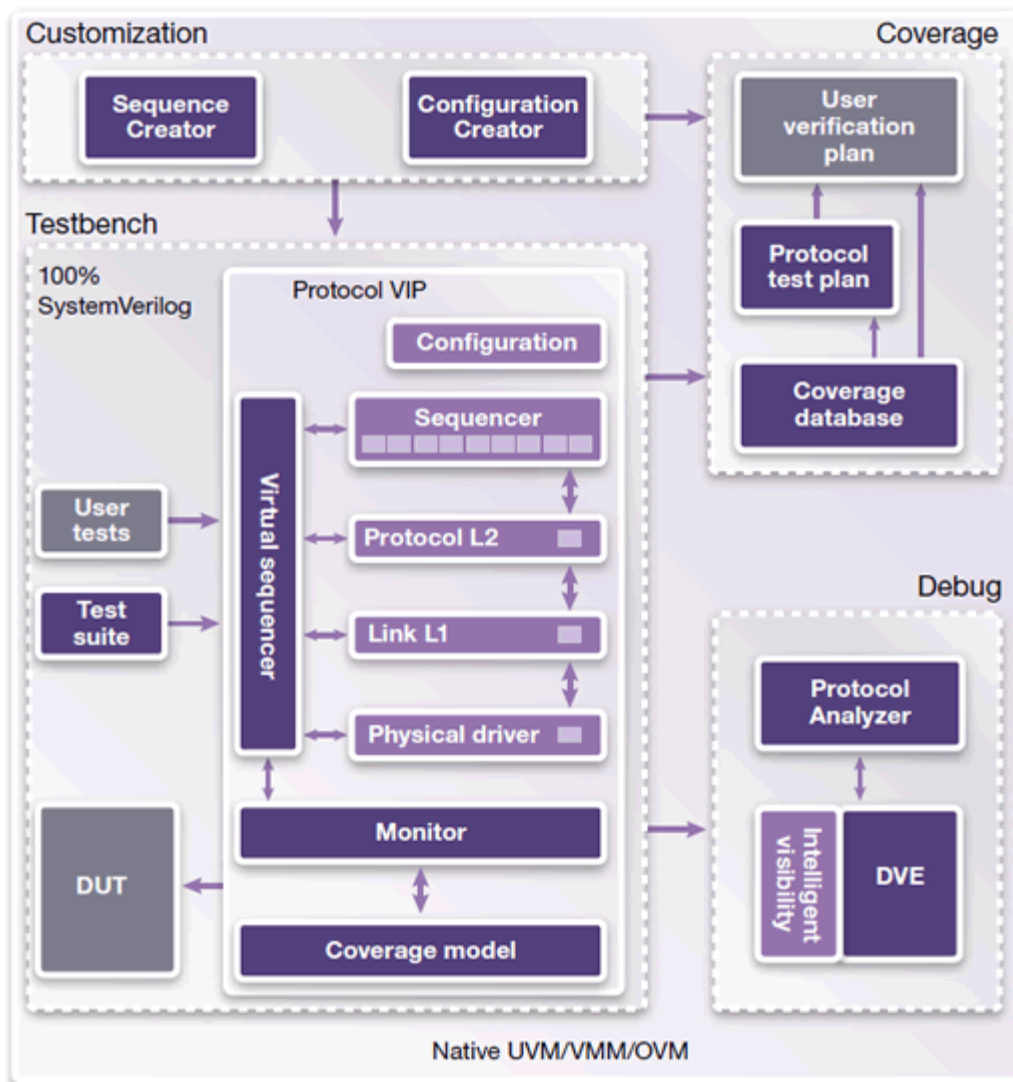
In summary, today's first-generation VIP does a reasonable job of modeling the protocol and some of the verification features. It generally does a poor job of accelerating the verification of today's highly complex protocols used in SoCS.

### **Synopsys Discovery Verification IP Meets New Challenges**

In order to maximize productivity across the whole design and verification cycle, verification IP needs to change its focus to help reduce the burden of tasks that require deep protocol knowledge; enabling engineers to rapidly get to first test, run verifications, fill coverage holes, and debug unexpected events. VIP needs to provide greater performance, debug, coverage management features, ease of use, ease of integration into test environments and reuse.

Synopsys' Discovery verification IP is based on Synopsys' VIPER architecture, which has been engineered from the ground up for enhanced VIP performance, configurability, portability, debug, coverage management, and extensibility. VIPER uses a layered architecture implemented in SystemVerilog using best practices for all methodologies, including UVM, VMM and OVM. It offers significant advantages in increasing protocol-based verification productivity:





[ Figure 1: Synopsys VIPER architecture ]

## Performance

Verification performance of VIP has become increasingly important for SoC verification where the number of VIPs and instances of VIPs can easily become the bulk of the testbench.

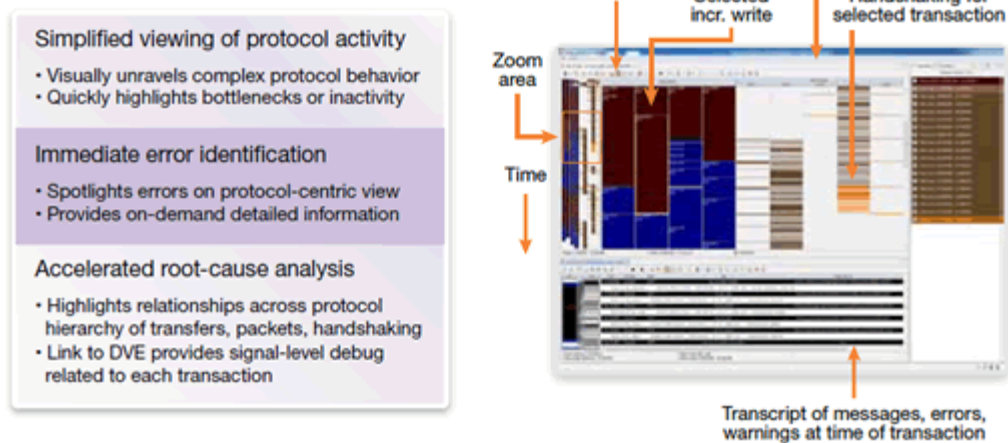
- Synopsys Discovery VIP is 100% SystemVerilog and runs natively in any simulator for optimum performance.
- It includes support for UVM, VMM, and OVM without using translation wrappers that degrade performance, as already discussed. A switch at compile time causes Discovery VIP to compile with whichever base-class library is being used.
- Discovery VIP enables verification teams to complete tasks in a fraction of the time compared to their previous approach. It boosts simulation performance by up to 4x compared to previous generations of VIP.
- Native support for UVM, VMM and OVM without the need for methodology-level interoperability wrappers or language translations. All layers are visible, providing complete controllability of protocol

verification. Verification engineers are able to work at the highest layer required to meet their verification plans, yet are still able to inject errors at the lowest layers.

## Debug

Debug, as already mentioned, is a challenging and time-consuming task requiring deep, protocol knowledge. First-generation VIP provides error messages when illegal behavior is detected at the interface for one of a multitude of reasons. However, with the complexity of today's highly interleaved and hierarchical protocols, it can be extremely challenging to find out what caused the error using the data extracted from log files, transcripts and trace data. Synopsys has developed a protocol-aware debug environment that unravels all of this complexity, and enables users to see the protocol at a higher-level, while preserving the lower-level detail. The Discovery VIP Protocol Analyzer:

- Gives users a graphical, hierarchical view of the transfers, transaction, packets and handshaking of a protocol. It shows detailed information for all of these objects.
- Annotates errors, warnings and messages onto the protocol view to quickly find the cause of the errors and warnings.
- Is simulator independent.
- Displays the full VIP class hierarchy with built-in help, and includes full links to VIP user documentation.
- Can link to simulator trace views.



[ Figure 2 ]



[ Figure 3: Protocol Analyzer enables protocol-aware debug ]

## Coverage

Coverage is the measure of progress against the test plan and has historically been built by verification engineers; a very time consuming but critical part of the testbench. With directed test methodologies, this is not so key, but with the advent of mainstream constrained random verification, coverage is a key part of the test environment.

- Built-in coverage is included in the Discovery VIP to generate coverage reports that match the verification plan and enable users to quickly identify coverage holes. This is another place where references back to the specification are very helpful.
- The protocol-based sequence library provided with the Discovery VIP gives engineers a jump-start on achieving coverage for their designs.

## Configuration

Getting to first test is one of the major steps in any testbench development, and reducing this is clearly an area of focus for any team. Discovery VIP has a number of built-in features to address this requirement, including:

- HTML-based, task-focused documents walk engineers systematically through the key steps of integrating and using the VIP in testbenches of any methodology. They save hours of time sifting through traditional user documentation and include links to relevant code examples
- Configuration Creator provides all of the VIP configuration options in one environment that shows the default values and possible values for all of the parameters. Built-in help is there to help engineers understand the implications of all the possible parameters. Validity checking confirms there are no conflicts.

- Built-in test plans are provided for the different configurations of the VIP. The coverage points in the test plan reference back to the specification so engineers can quickly refer back to specified functionality covered by the test. These become the basis for coverage and are updated to match the parameters that are set in Configuration Creator. After all, there is no point getting coverage data for coverage points that can't be reached because of the configuration. Users can edit the plans to add or remove specific coverage points.

### **Summary**

The performance requirements and complexity of the latest interfaces demand a different approach to verification if design teams are to meet increasingly aggressive project timescales.

By delivering a new generation of verification IP, Synopsys is enabling non-specialist engineers to get up and running faster than ever, and, in the process, become verification experts. Together with other innovative features creating a VIP solution around SystemVerilog, Discovery VIP reduces the time to first test, makes debug easier, and reduces overall verification time and the number of engineer-hours it takes to achieve the coverage goals.

