

Special Technology Report

Mixing Formal and Dynamic Verification, Part 2

By Bill Murray

In part 1 of this STR, we:

- Quantified the broadening and deepening of formal verification deployment
- Quantified the fast ramp of recent adopters
- Quantified the advancement of formal verification into design
- Discussed the crumbling barriers to formal's adoption
- Chronicled the quest for quality
- Discussed productivity
- Outlined formal verification's sweet 'n' sour spots

And we published our Exhaustive Proof cartoon.

In part 2 of this STR, we:

- Outline the 17 use cases used in the survey
- Discuss the detailed survey results
- See how formal is currently being used with dynamic verification
- Look at the application of formal in the ESL space
- Hear from both technology users and technology providers how formal methods might look in 2012
- Wrap up with the most significant findings in both parts of this report

The SCDsource Survey

According to Gary Smith of [Gary Smith EDA](#), formal property checking is used by 65 percent of chip design companies. Smith said "The proliferation of SystemVerilog and its assertion language, SVA, has boosted the adoption of formal property checking."

We surveyed 19 engineering managers and engineers in 16 companies to quantify and characterize their use of formal property checking. Our survey population spans the design and verification of application-specific chips, general purpose processors, graphics processors and silicon intellectual property (IP) for applications in computing, consumer electronics, medical electronics, military/aerospace, networking and wireless communications.

The 16 companies are industry-leading IP, chip and systems design companies – Alcatel-Lucent, Analog Devices, ARM, Cisco, DE Shaw Research (DESRES), Fujitsu Microelectronics Europe, HP, IBM, Infineon, Intel, Nvidia, Qualcomm, Saab, Silicon Logic Engineering (Tundra), STMicroelectronics and Sun Microsystems.

Of the 19 respondents, 14 were employing formal verification in 2006 – this STR refers to them as established users. This STR refers to the other five respondents, who adopted formal after 2006, as recent adopters.

The 17 Use Cases

The survey covered the following use cases. These use cases were gleaned largely from discussions with our interviewees, as well as from published material.

1. Early RTL code analysis to detect common design errors. Uses formal code checks and/or assertions generated either manually or automatically [5].
2. Early block-level sanity checking and bring-up verification to check micro-architecture and implementation intent, using [white-box assertions](#).
3. Share block design intent and integration assumptions, using assertions.
4. Capture and share coverage targets using assertions that validate formal constraints and define functional coverage points.
5. Protocol compliance verification.
6. Exploration and bug hunting for block/subsystem features, functions, operations, and transactions, using [grey-box](#) or [black-box](#) assertions.
7. Exhaustive verification of block/subsystem features, functions, operations, and transactions using grey-box or black-box assertions.
8. Verification plan-driven systematic, thorough block/subsystem verification.
9. Analysis of assertion sets for missing assertions – detects verification gaps and validates/improves informal specification.
10. Analysis of coverage items not reached by simulation/testbench.
11. Subsystem bring-up and integration verification.
12. Signal connectivity verification.
13. Engineering change order (ECO) verification.
14. Root cause analysis of complex errors found by simulation.
15. Root cause analysis of complex errors for post-silicon debug [10, 11].
16. Legacy block/IP exploration and redesign using assertions.
17. Functional sign-off criterion.

Survey Results

We show some of the more important survey results below. Note that all survey data in the following graphs are rounded, while derived data (such as growth rates and ratios) are calculated from the original, non-rounded survey data. The resulting tiny inconsistencies in the maths yield considerably greater clarity in the analysis.

Use Case Employment in 2009

Use case employment in 2009 by our 19 respondents is shown in figure 7. The results are ordered from most widely used at the top to least widely used at the bottom.

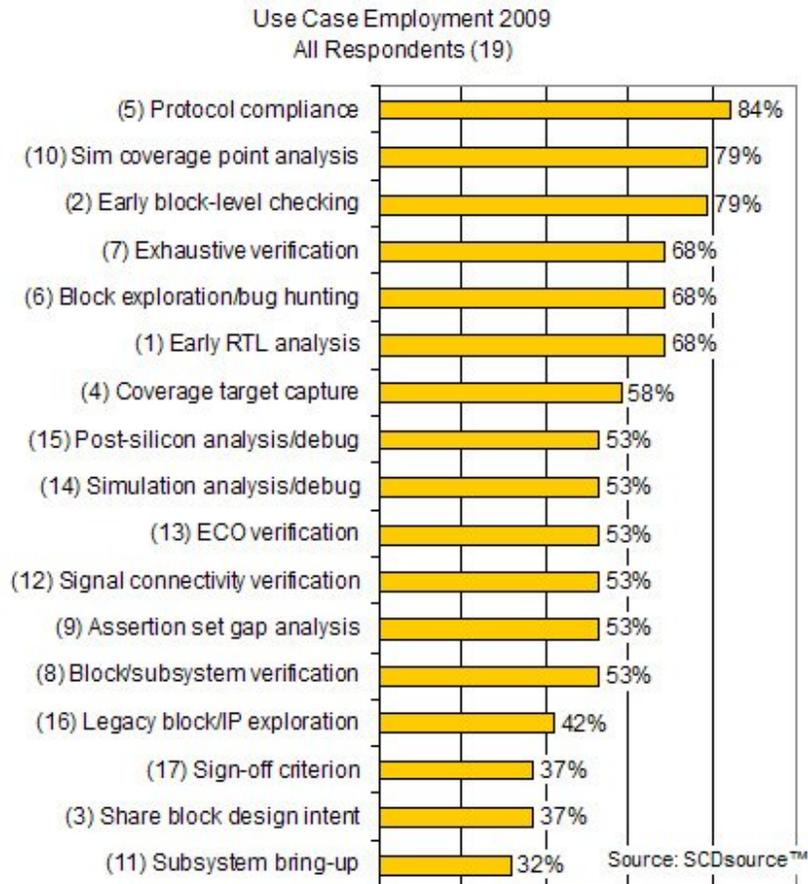


Figure 7: Use case employment 2009 (Source: SCDsource)

The results show significant use of formal verification throughout the design and verification flow from initial RTL checks to verification sign-off, and even after first silicon availability. Even the least-used case is employed by nearly one-third of our respondents. Of particular note is:

- The evidence that formal verification is advancing backwards into design (see part 1 of the STR). Early RTL analysis is already among the most widely-employed use cases in 2009, used by 68 percent of the respondents; and about 37 percent use formal assertions to define and share design intent.
- Analysis of coverage items not reached by simulation is a second-ranked use case, with 79 percent of respondents employing it. Clearly, this complementary verification approach is a key value in mixing formal and dynamic verification.
- More than half of the respondents use formal analysis to identify the root cause of complex errors in post-silicon debug.

Top 3 Use Cases in 2009

Does “most widely used” necessarily mean “most important?” We asked our survey participants to identify their Top 3 use cases. The results are shown in figure 8, maintaining the use case order of figure 7.

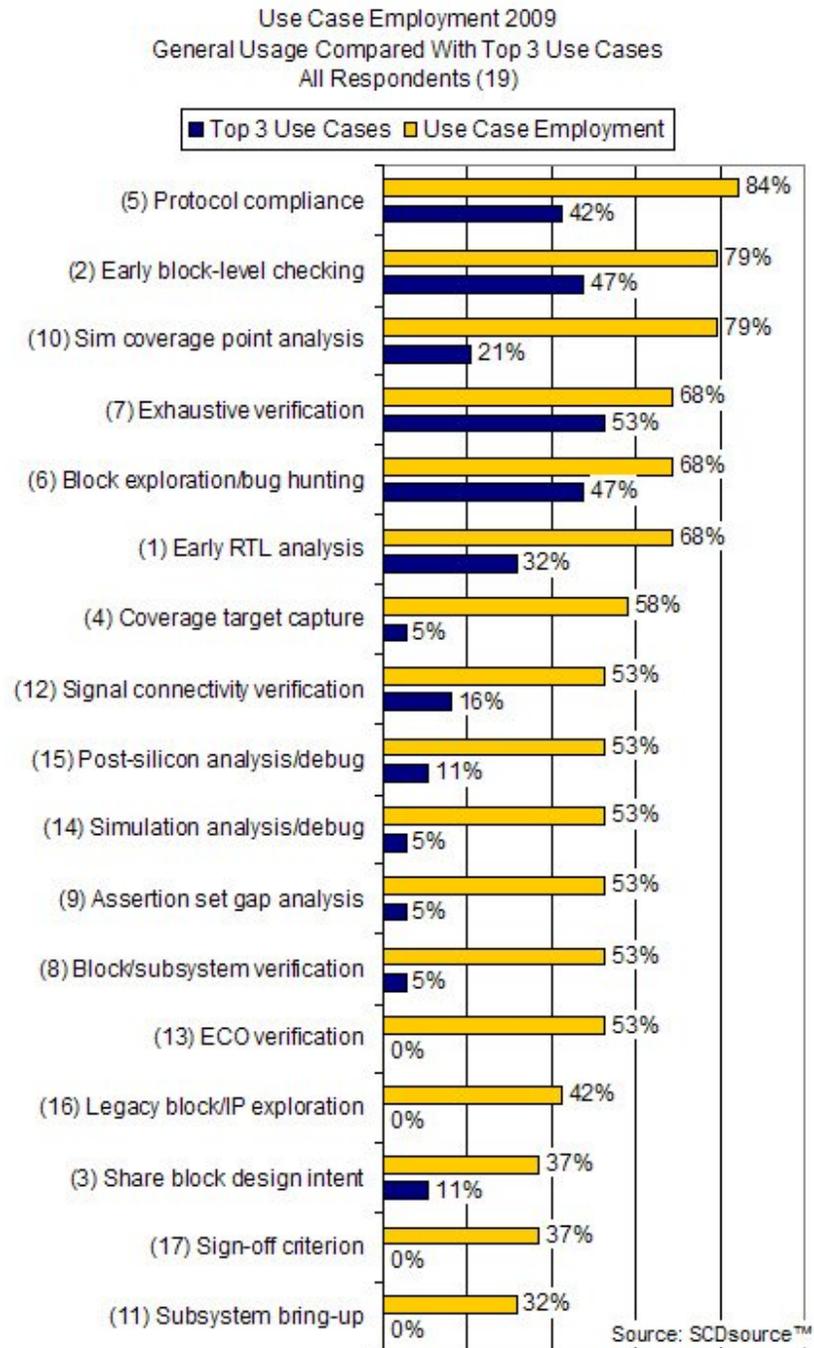


Figure 8: Aggregate Top 3 use cases vs. use case employment 2009 (Source: SCDsource)

The results illustrate the importance of the question. Protocol compliance verification is the most widely employed use case, employed by 84 percent of respondents, but only half of those respondents identified it as a Top 3 use case. By contrast, exhaustive verification is the fourth most widely employed use case, used by 68 percent of the respondents, but nearly four-fifths of them – 53 percent of respondents – identified it as a Top 3 use case. Result: protocol compliance verification ranks fourth in the Top 3 ranking, while exhaustive verification ranks first. Exhaustive verification – formal's major traditional perceived strength – remains the most important use case, despite formal's proliferation up and down the design and verification flow.

Two use cases tie for second rank, with 47 percent of respondents placing these use cases in their Top 3:

- Early block-level sanity checking and bring-up verification to check micro-architecture and implementation intent, using white-box assertions.
- Exploration and bug hunting for block/subsystem features, functions, operations, and transactions, using grey-box or black-box assertions.

Note: protocol compliance verification is not ranked third because the second rank tie puts three use cases ahead of it – there is no third rank.

Is there a difference in use case employment between recent adopters and established users in 2009? The two are compared in figure 9, maintaining the use case order of figure 7.

The shapes of the two distributions are similar, indicating that the two groups may have similar assessments of the ROI of any given use case.

Interestingly, of the recent adopters, 100 percent employ formal methods to analyze coverage items not reached by simulation, compared with 71 percent of established users. The absolute percentages, however, should be used with caution – a single “vote” in the recent adopter population can change the use factor by 20 percentage points, while a single “vote” in the established user population can change it by only about 7 percentage points.

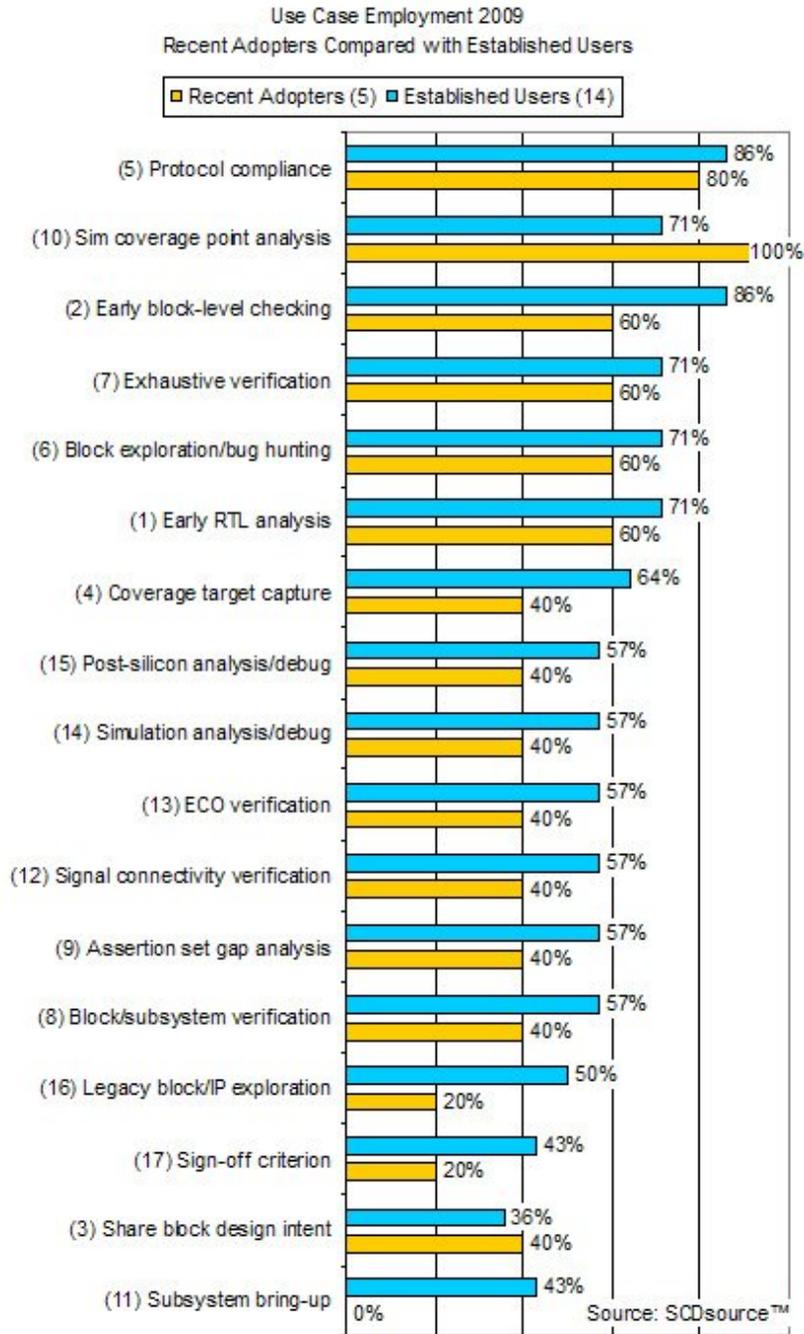


Figure 9: Use case employment 2009 recent adopters vs. established users (Source: SCDsource)

Adoption Speed and Profile

In part 1 of this STR, we compared the average use case employment of recent adopters and established users, and observed that the latter appear to be leading the former by about three years. Specifically, in 2009, recent adopters are employing about 7.8 use cases on average versus 6.6 by the established users in 2006; while recent adopters are projected to use 10.8 use cases in 2012 versus 10.4 by established users in 2009.

The obvious question is: how do the detailed employment profiles compare at similar stages of adoption? Are there any obvious similarities or dissimilarities? A comparison of the two groups' use case employment is shown in figure 10, maintaining the use case order of figure 7. It compares use case employment by established users in 2006 with that of recent adopters in 2009.

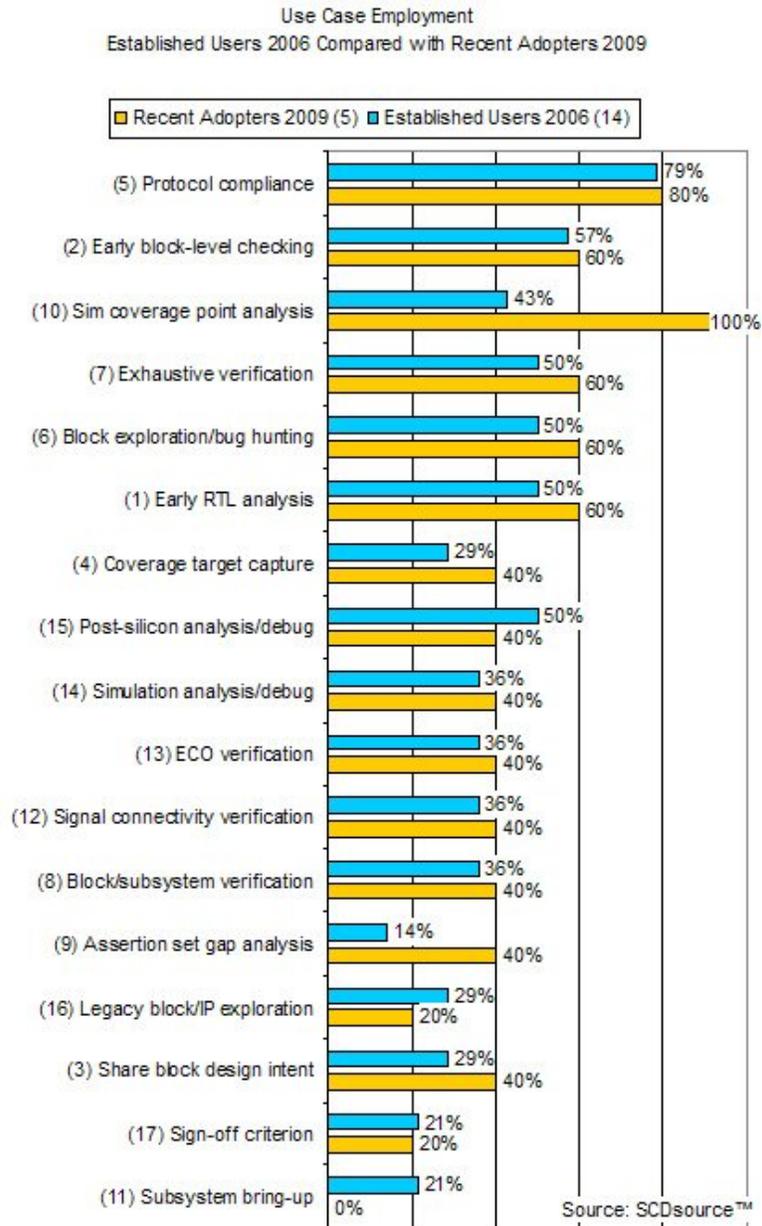


Figure 10: Established users 2006 compared with recent adopters 2009 (Source: SCDsource)

With the exception of the analysis of coverage items not reached by simulation, the shapes of the use profiles are similar. Moreover, the percentage of respondents in each group employing any given use case is similar – often nearly identical – with the exception of subsystem bring-up and the analysis of coverage items not reached by simulation.

In summary, then, the recent adopters in 2009 are in roughly the same state of adoption as the established users were in 2006. Could improved methodologies and technologies narrow the gap or, more accurately put, could they significantly accelerate the adoption of formal verification? If so, how? That is the subject of the next section.

Mixed Formal and Dynamic – How?

In a DAC 2005 paper [12], ST's Umberto Rossi observed that "From the user point of view, verification products still lack integration between dynamic analysis and formal assertion analysis. A closer cooperation or a better merger of the respective results may constitute a significant plus for the users."

This is the key to greatly increased formal verification use in the mainstream verification flow. Knowing where best to apply formal and dynamic verification methods and what the coverage/productivity payback is in any given scenario will help design and verification teams to make dependable, highest-ROI decisions with predictable results. Such consistent, comprehensive, objective data could identify and help proliferate best practices – and narrow the gap between recent adopters and established users.

Cross-Leveraging Assertions

So, how does the industry achieve that integration between formal and dynamic verification? Of course, both techniques leverage assertions. What kind of assertions, and where? In our survey, we partitioned assertions into two basic classes [13]:

1. Structural assertions
 - Low level assertions that verify detailed RTL behavior, such as that of a finite state machine (FSM).
 - Adequate for localized sanity checking and verification.
 - Must often be modified in response to RTL implementation modifications.
 - Difficult-to-impossible to verify a design exhaustively.
2. Operation/transaction assertions
 - Assertions that verify RTL behavior at a higher level, such as block reads and writes.
 - Correspond more closely to functionality as defined in the specification – a more intuitive approach to verification.
 - Must often be modified in response to specification changes, but generally not to RTL implementation changes.
 - Can deliver significantly higher verification productivity than structural assertions.
 - Suitable for exhaustive verification.

Our respondents' use of these assertions in simulation/testbenches and formal verification is shown in figure 11.

Verification Method	Assertion Type	2009	2012
Simulation	Structural	100%	100%
	Operation	88%	100%
Formal <small>Source: SCDsource</small>	Structural	88%	94%
	Operation	100%	100%

Figure 11: The use of different assertion types in simulation and formal verification (Source: SCDsource)

All of our respondents to this part of the survey (in this case, 16, not 19) use structural assertions in simulation, and all use operation assertions in formal. Eighty-eight percent of respondents use operation assertions in simulation, projected to increase to 100 percent by 2012; and 88 percent use structural assertions in formal, projected to increase to 94 percent by 2012.

This large degree of overlap would indicate the potential for a great deal of cross-leverage. So, how are they being cross-leveraged? Some respondents employ the same assertions in both environments, but many don't. Moreover, our survey shows that only 31 percent of respondents use formal to verify testbench assertions, both structural and operational. However, this is projected to increase to 63 percent by 2012.

Integrating Coverage Metrics

Establishing a tighter coupling between simulation and formal verification is clearly a significant need for most of the respondents in the survey, but they need a methodology to do it. Right now, most of our respondents mix formal and dynamic verification at only the verification plan level. Thereafter, the two approaches largely go their separate ways. Each has its own coverage goals; progress is measured or estimated separately; and total coverage is determined by eye-balling overall progress to the verification plan.

ARM wants to make eye-balling easier. Bryan Dickman said "We're moving towards database-driven planning. The idea is to upload results to a mySQL database and use the verification plan as a database query. We would then have a live view of results versus plan. Of course, both plan and database would contain both simulation and formal targets. That's our destination."

However, what everyone really wants is a unified coverage approach. DESRES' Ho said "A combined formal and simulation verification method would be much more effective if users could know how much coverage was achieved by each technique, and thereby avoid redundant verification." He continued "We need the ability to import coverage data from different tools from different vendors. Currently, there is no good way of doing that without a great deal of user involvement in translation."

Ho is now the chair of [Accellera's](#) Unified Coverage Interoperability Standard ([UCIS](#)) committee. The stated objectives of the committee are to:

- Identify interoperability opportunities between various coverage sources
- Define standard coverage models for commonly used metrics
- Define an operability standard that allows coverage data to be exchanged among EDA vendors' tools and IC vendors environments

Ho, together with Accellera's chair, Shrenik Mehta, reported progress-to-date. Ho said "The committee seeks to establish standards for those coverage items upon which members can agree, but also to

define an extension mechanism for vendor-specific or user-specific coverage types.” Ho observed that “Standardizing coverage types will enable the automatic merging of data, while standardizing an extension mechanism for new coverage types allows for innovation.”

A draft proposal [14] dated April 2009 identifies three use models for data merging, 16 verification coverage types targeted for standardization, and a conceptual API for the database.

The use models cover:

- Temporal merging: merging of coverage data across multiple runs of the same verification process.
- Spatial merging: merging of coverage data across different parts of a design.
- Heterogeneous data merging: merging of data from different verification processes.

The 16 verification coverage types are:

- Toggle coverage
- Line/statement coverage
- Branch coverage
- Condition/expression coverage
- Trigger coverage
- FSM state coverage
- FSM transition coverage
- Value-transition coverage
- Path coverage
- Cover point/coverage point
- Cross-coverage
- Cover group coverage
- Cover property
- Assertion coverage
- Assertion result
- Transaction coverage

Accellera expects to release a document to the general public in 2010.

System-Level Formal Verification

Formal Analysis

In his DAC 2005 paper [12], ST’s Umberto Rossi noted that, in order to verify the millions of complex configurations of a parametric IP “Formal methods have to extend their effectiveness by including so-called transactional level models (TLM) in their analysis, which abstract the mechanisms that SoC blocks use to communicate and synchronize among themselves. Two major problems have to be faced: the first is that TLM includes a range of abstraction levels and the second is that a mapping between TLM assertions and RTL assertions should be provided and formally proven.”

Three years later, not much had changed. In a DVCon 2008 presentation [15], Wolfgang Ecker, principal engineer in Infineon’s communications solution business group, observed “The barrier to SystemC TLM interconnect verification is that SystemC does not support native temporal assertions.

Previous efforts to rectify this deficiency in SystemC apply assertion-based verification (ABV) only at the RT-level.”

Mike Meredith, president of the Open SystemC Initiative (OSCI), said “There have been discussions in the OSCI verification working group about what temporal assertion checking at the TLM level really means. No firm consensus has been reached as to any standardization. However, there are OSCI members who are investigating formal property checking in SystemC, and they are making some progress in devising methods to do it. These would be proprietary, but are still valuable additions to the SystemC ecosystem.”

Does OSCI expect to see any progress by 2012? Meredith said “As far as development and deployment in the ecosystem is concerned, I certainly do. As far as standardization is concerned, it depends upon what makes sense to the developers of the proprietary approaches and the rest of the verification working group about what to try to standardize.”

System-level Equivalence Checking

However, there is one system level formal verification approach that *is* in use today. System level equivalence checking verifies equivalence between a C/C++/SystemC architecture model and the corresponding RTL micro-architecture implementation. The technology's sweet spot is assertion-based formal verification's sour spot – data transformation logic.

According to Prosenjit Chatterjee, hardware engineering manager at Nvidia, “Graphics processors leverage a broad range of data transform functions, such as arithmetic logic, address translation and, in general, complex mux logic. Many of these blocks are quite complex, mission-critical functions – not just simple blocks.”

Prior to adopting the technology, said Chatterjee “We relied upon directed and directed random simulation tests to compare the two models. It found a lot of bugs, but it could not exhaustively cover all possible input scenarios, so it was not possible to conclusively prove equivalence. Now, we can prove equivalence with system level equivalence checking, and use any equivalence mismatches to devise simulation tests.”

Adoption took a matter of a few days, according to Chatterjee. “There is not much training required. It is not necessary to write properties because, in essence, the C model is the property. The environment requires simple input constraints, such as min/max input values. Sometimes the equivalence check is push-button; sometimes we must decompose the input values into sub-cases or break up the two models at transaction matching internal interfaces, and then prove each one individually.”

The value of the approach grows with increasing expertise in writing C models that the technology can handle. According to Chatterjee, “A C model developer might describe data transform functions using dynamic memory allocation or loops with arbitrary bounds. However, hardware design needs fixed bounds, so C developers must replace these constructs with fixed-bound constructs – say, a fixed memory with fixed min/max bit count.” In other words, the developer has to use synthesizable code, regardless of whether the design team uses high level synthesis. Chatterjee added “As the technology is increasingly enhanced with further C constructs, developers have more freedom to use constructs that come more naturally to them.”

Nonetheless, the requirement that the code be synthesizable remains. Any disadvantages? “C is not really optimal for synthesis. The C model might actually be larger than the RTL model,” said Chatterjee. He also pointed to the inherent difficulty of identifying correspondence between the abstract C model and the RTL implementation model. “It's not the same as RTL/gate and gate/gate equivalence checking,” he said.

Chatterjee concluded "We plan to use system level equivalence checking on as many data transformation functions as possible. Looking ahead to 2012, I hope to see system level equivalence checking handle much bigger design sizes."

Asked about the progress towards an agreement on the synthesizable subset of SystemC, OSCI's Meredith said "We've made a lot of progress over the last twelve months. There is a significantly upgraded draft in the OSCI internal approval loop, and we expect to put it into public review this Spring.

The World in 2012

Use Case Employment in 2012

In part 1 of this STR, we showed that use case employment will grow by 32 percent from an average of 9.7 use cases per respondent in 2009 to 12.8 in 2012 – a significant deepening of formal use in this population of 19 respondents. What does the detailed 2012 projected use case employment look like? Results are shown in figure 12, maintaining the use case order of figure 7.

Use Case Employment
2012 Projection Compared with 2009

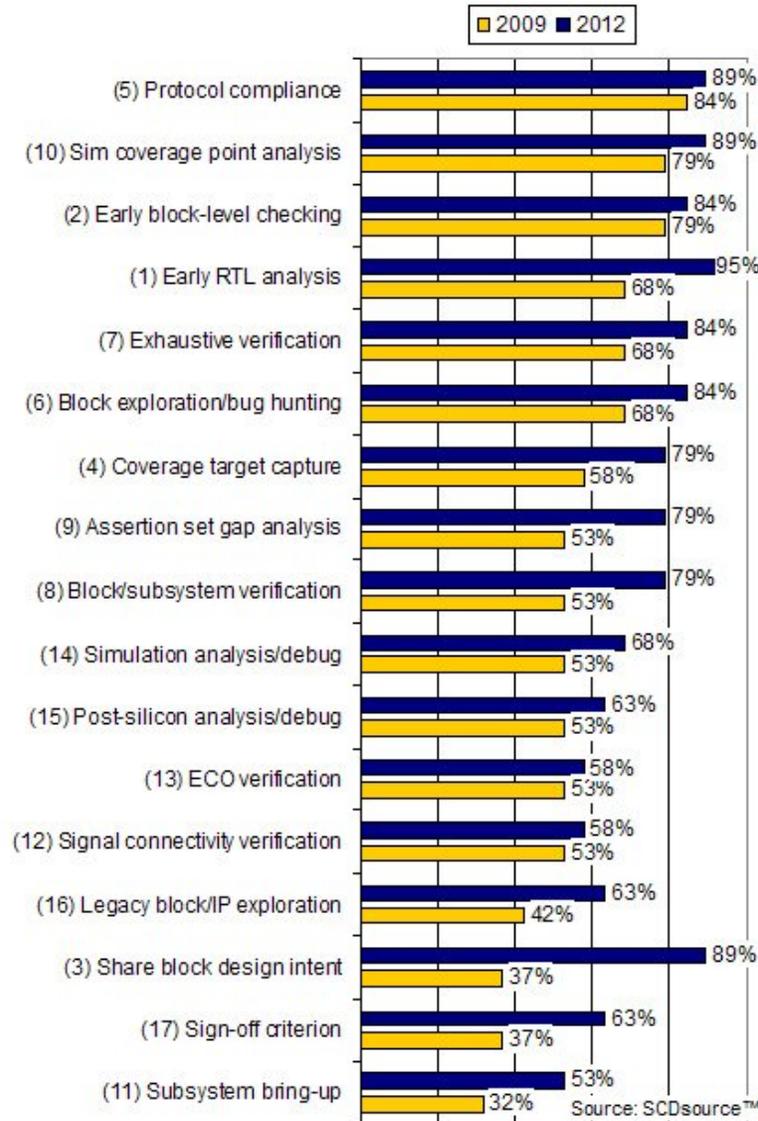


Figure 12: 2012 use case employment compared with 2009 (Source: SCDsource)

The most widely deployed use cases in 2012 are projected to be:

1. Early RTL analysis, identified by 95 percent of respondents. This is up from 68 percent of respondents in 2009. This proliferation and growth (38 percent over 2009) indicates widespread adoption of formal methods by design engineers.
2. Share block design intent and integration assumptions, identified by 89 percent of respondents. Again, this proliferation indicates widespread adoption of formal methods in design.

The top two growth cases from 2009 to 2012 are:

1. Share block design intent and integration assumptions, which is projected to grow by 143 percent from 37 percent of respondents in 2009 to 89 percent in 2012.
2. Sign-off criterion, which is projected to grow by 71 percent. Proliferation of this use case from 37 percent of respondents to 63 percent of respondents is consistent with (a) a growing confidence in formal verification and (b) a consequent reliance upon formal verification at a very critical point in the mainstream verification flow.

Users look to 2012

We asked interviewees to look out three years and visualize what formal verification may look like – or what they would like it to look like. We were not looking for a forecast, just a vision of how the methodology and the technology might look in 2012. Most foresee “bigger, better, faster” technology – more capacity, improved ease of use, faster proofs and debug. And most assume that these improvements will occur incrementally.

However, closer integration of formal and simulation methodologies and technologies was at or near the top of everyone’s list. Users pointed to the need for unified coverage metrics, a modular verification infrastructure, common constraints, and a correct-by-construction design flow. And most are keeping a close eye on the Accellera UCIS efforts.

Modular Infrastructure

For HP’s David Lacey, the importance of the UCIS effort is more than simply melding coverage results – it is critical to the establishment of a modular verification infrastructure. He said “Such an infrastructure would give us more of a toolbox approach to verification. For example, we would be able to integrate verification IP into formal and simulation in much the same way. So, formal needs to go down that path.”

Common Constraints

D. E. Shaw Research’s Richard Ho would really like constrained random testbenches to have the same constraints as formal. “This would enable simulation to be really smart. You could run simulation up to a coverage point of interest, then invoke formal analysis. This requires that the constraints for random testbenches be understandable by formal tools. The problem is that testbench writers use biasing of random variables and layered scenarios instead of using assertions to constrain the input sequences. A common language would solve a big interoperability problem and yield very smart verification.”

Infineon’s Darren Galpin concurs. “When writing constraints for random dynamic verification and formal verification environments, you’re essentially doing the same thing, but in a different way. For the random dynamic environment, you’re writing a set of code to generate – for example – possible stimuli. You then take a random test path through the protocol space. In the formal environment, you have to constrain the proof to legal operations and test everything simultaneously. I don’t see any reason why those two sets of constraints ultimately couldn’t be used for both tools.”

Correct-by-Construction Design Flow

ARM’s Bryan Dickman said “We hope to see a correct-by-construction design flow that may also reduce the overall verification burden. We could get there by having designers use the tools from design onwards in the flow to systematically explore and formally prove their code. At the moment, we’re using formal rather retrospectively.”

Formal "a must"

The last word from the users goes to IBM's Jason Baumgartner: "It used to be that you must have sinned greatly to have the formal verification guy in your office. But, with greatly enhanced capacity and reusability, formal has become substantially more mainstream, and is now much easier for casual users to set up. Our post-design 'lessons learned' reviews generally find that we should do more in formal and mixed formal/dynamic verification. Many now view formal as a must."

Technology providers look to 2012

Cadence Design Systems

Sarah Lynne Cooper Lundell, senior product marketing manager, said "By 2012, the focus will be on analyzing and optimizing clear and measurable unified verification metrics. How the contributions to the metrics are generated will be relatively less important. They could come from any combination of formal, hybrid, testbench simulation, acceleration, emulation or other engines. But the discussion will be about what the metrics mean for reaching verification closure."

According to Lundell, this shift will increase formal's ROI because "unified metrics eliminate the need to duplicate efforts to verify any given functionality, and allow project teams to deploy the right solution for a given task and for the skills available. Traditionally, formal analysis has been used either early in the verification process prior to the availability of a testbench, or on really tough verification targets; but generally adjacent to mainstream verification. Unified metrics will allow project teams to track formal's progress alongside that of testbench simulation. This will break down walls that have hampered formal's adoption."

Calypto Design Systems

Anmol Mathur, CTO, expects system-level equivalence checking to become a standard step in high-level synthesis (HLS) verification flows by 2012. He said "The two main enablers will have been established. First, the synthesizable subsets of C/SystemC will have been standardized; and second, a standard exchange format to transfer information from HLS tools to system level equivalence checking technology will have been agreed."

Mathur continued "This will enable checking tool to leverage information about boundaries such as function hierarchies or loop bodies to decompose the complex formal equivalence checking problem in order to scale to the block sizes that HLS tools will target in 2012."

Moreover, he said, "Fundamental progress in formal solvers that work at the operator-level rather than manipulating bit-level Boolean functions has been critical to allowing system level equivalence checking technology to tackle arithmetic intensive designs. I expect better merging of this technology with formal analysis of control to further boost tool capacity. And further improvements in sequential proof engines will allow designs with larger divergence in sequential behaviors to be checked."

Jasper Design Automation

Jasper sees formal methods extending beyond verification. Rajeev Ranjan, CTO, said "Higher tool capacity and scalability, coupled with innovative methodology, will enable the application of formal throughout the design flow, from architectural analysis, through design and integration, to silicon debug."

Ranjan continued "And formal technology will bring breakthrough changes in IP delivery. Third party IP will be encapsulated in a formal technology-based 'executable shell' that will make it easy for

designers to adopt the IP. It will allow designers to easily comprehend the critical behaviors of the IP, explore dependencies, and understand bus protocols.”

The company also projects that formal technology will be applied at the system level. Ranjan said “Standard ways of architectural and system-level modeling will enable the application of formal to property verification of high-level models, such as SystemC models. We shall see formal methods used in architectural models, TLM, HW/SW co-design, and algorithms developed in MatLab.”

Mentor Graphics

Harry Foster, chief verification scientist, foresees new, automatic applications built on formal technology that are easy to use and require no special expertise in formal verification or languages. “Furthermore, we predict a tighter integration of formal and simulation algorithms to solve specific problems in coverage closure as well as pre- and post-silicon debugging,” he said.

Foster pointed out that about 19 percent of the industry has adopted formal property checking, and about 41 percent has adopted simulation-based, coverage-driven techniques that leverage functional coverage properties specified using modern assertion languages.

He said “Clearly, there are opportunities to improve verification processes by taking advantage of assertions in either simulation or formal property checking. However, part of the challenge of adopting assertion languages is the expertise required to grasp temporal logic and declarative forms of specification. These concepts, which form the foundation of today’s assertion languages, are foreign to many of engineers. Perhaps the recent emergence of assertion libraries and assertion-based verification IP will assist the adoption of assertion-based techniques in the future.”

OneSpin Solutions

Michael Siegel, OneSpin's director of product marketing, stated “Today, formal verification is often used in addition to dynamic verification to increase design quality. However, the main potential of combining formal and dynamic lies in mastering the verification productivity challenge. We foresee enhanced methodologies and tools that guide users to the optimal verification approach in order to predictably reduce overall verification effort. Knowing when, where and how to apply the most appropriate method – formal or dynamic verification – is key.”

Siegel foresees enhancements that will spawn more automated and integrated formal/dynamic verification flows. “Formal will be a standard, mainstream technology in the verification toolkit for design, verification and integration engineers,” he said, “and enhanced tool interoperability will reduce the manual effort required in today’s combined formal/dynamic verification flows.”

“Increased verification productivity also depends upon the ability to accurately assess overall verification coverage and tool support to close coverage gaps,” Siegel continued. “So, it is critical that the industry establishes suitable standard coverage metrics shared between formal and dynamic.”

Real Intent

Carol Hallett, Real Intent's vice president of sales and marketing, said that automatic formal verification is an essential component in all functional verification flows today, and that its use will increase substantially by 2012.

According to Hallett, “Automatic formal verification applies pre-determined rules to perform error checking automatically, with virtually no engineering effort or design style restrictions. It runs formal analysis automatically, and generates failure reports that minimize debug effort.” She continued “The methodology tackles two endemic problems. Its early bug detection mitigates the spiral of bug

detection, debugging and bug fixing; and its automation deals with the ever-growing problem of constrained resources.”

Hallett predicts that improved automatic error checking rules applied early in the design flow; faster convergence on larger designs; and much more advanced reporting and debugging support will increase the bug detection rate from about 50 percent today to about 70 percent by 2012.

Synopsis

Dan Benua, principal corporate application engineer, sees a continuing trend to use more formal and assertion-based verification to complement the dominant constrained-random coverage-driven simulation approach. “This trend is being driven by the need to improve verification cost-efficiency, and is also facilitated by progress in formal technology, language and library standardization, and easier-to-use tools,” he said

He continued, “There will still be experts doing hard-core model checking, but the average designer or simulation user will be leveraging more formal analysis also. In some cases, engineers may not even be aware that formal algorithms are running inside their simulator; providing benefits like debugging hints, coverage closure, or smarter stimulus generation. Multi-core hardware runs these activities in the background with minimal impact on simulation runtime.

“We are still in the early stages of understanding how to integrate and potentially merge these technologies, said Benua. “We see big opportunities ahead to improve the functional verification process and we are hard at work on the technologies, methodologies, and tools needed to achieve them.”

Wrap-Up

So, what are the major findings of this STR?

- In this population of 19 users, formal verification’s use nearly doubled from 2006 to 2009.
- The 5 recent adopters are closely tracking the adoption curve of the 14 established users.
- The results show significant use of formal verification throughout the design and verification flow from initial RTL checks to verification sign-off, and even after first silicon availability.
- Adoption barriers are crumbling – capacity increases and ease-of-use enhancements make adoption easier – but many respondents want to see a lot more progress.
- The primary reason for adopting formal was verification quality – but productivity is a key factor in the extent of formal’s deployment.
- Formal’s sweet spots are expanding.
- Formal is “a must.”
- System level property checking is in a barely embryonic stage, but system level equivalence checking has been established as a viable approach.
- Last, but not least, we still need a robust, systematic methodology for mixing formal and dynamic verification. Our interviewees view the work of Accellera’s UCIS as critical to this effort.

References

1. [The Art of Verification with SystemVerilog Assertions](#). Faisal Haque, Khizar Khan, Jonathan Michelson. [Verification Central](#) 2006.
2. [Practical Approaches to Deployment of SystemVerilog Assertions](#). Faisal Haque, Jon Michelson. EE Times 2007.
3. Structured Assertion Design Verification for Complex Safety-Critical Hardware. Kristoffer Karlsson, Håkan Forsberg. Military and Aerospace Programmable Logic Devices (MAPLD) Conference 2008. [Paper](#) and [Presentation](#).
4. [A comparison of three verification techniques: directed testing, pseudo-random testing and property checking](#). Mike Bartley, Darren Galpin and Tim Blackmore. Proceeding of the 39th Design Automation Conference, 2002.
5. [Early formal verification of conditional coverage points to identify intrinsically hard-to-verify logic](#). Richard Ho, Michael Theobald, Martin M. Deneroff, Ron O. Dror, Joseph Gagliardo, David E. Shaw. Proceedings of the 45th Design Automation Conference, 2008.
6. [Automated formal method verifies highly-configurable HW/SW interface](#). Joachim Knäblein and Hans Sahm. SCDsource 2009.
7. Guidelines for creating a formal verification testplan. Harry Foster, Lawrence Loh, Bahman Rabii, Vigyan Singhal. DVCon 2006. [Paper](#) and [Presentation](#).
8. [Automatic Formal Verification of Fused-Multiply-Add FPU's](#). Christian Jacobi, Kai Weber, Viresh Paruthi, Jason Baumgartner. Proceedings of DATE 2005.
9. [Achieving Certified IP Quality Efficiently](#). Lorenzo di Gregorio, Carlo del Giglio, Michael Siegel. EE Times 2007.
10. Zero Escape Plans: Tying Together Design, Simulation, and Formal Methods for Bulletproof Stepping Validation. Erik Seligman, Carl Dreyer, Ken Haren, Raman Nayyar. Proceedings of DVCon 2008. Reported in Formal verification expands its use model by Bill Murray, SCDsource 2008.
11. [Post-Silicon Debug Using Formal Verification Waypoints](#) by C. Richard Ho, Michael Theobald, Brannon Batson, J.P. Grossman, Stanley C. Wang, Joseph Gagliardo, Martin M. Deneroff, Ron O. Dror, David E. Shaw. Proceedings of the 43rd Design Automation Conference, 2006.
12. [Can We Really Do Without the Support of Formal Methods in the Verification of Large Designs?](#) Umberto Rossi. 42nd Design Automation Conference, 2005.
13. [Assertion-Based Design](#). Harry Foster, Adam Krolnik, David Lacey. Springer Verlag, 2004.
14. Unified Coverage Database Application Programming Interface (API) Design Specification Document, Draft 29, April 2, 2009. Accellera [Unified Coverage Interoperability Standard](#) (UCIS) committee.
15. Formal Techniques Speed Up Interconnect Verification of SystemC Virtual Platform Models. Wolfgang Ecker, Volkan Esen, Robert Schwencker, Thomas Steininger, Michael Velten. DVCon 2008. Reported in "Formal Verification Expands Its Use Model by Bill Murray, SCDsource 2008.

Further Reading

1. Formal Property Checking - What the Users Say, by Richard Goering, SCDsource 2008
2. Increasing Confidence of Complex Hardware in Safety-Critical Avionics Using Formal Methods. Kristoffer Karlsson and Håkan Forsberg. Military and Aerospace Programmable Logic Devices (MAPLD) Conference 2006.