

Trends in Functional Verification: A 2014 Industry Study

Harry D. Foster
Mentor Graphics Corporation
Wilsonville, Or
Harry_Foster@mentor.com

ABSTRACT

Technical publications often make either subjective or unsubstantiated claims about today's functional verification process—such as, 70 percent of a project's overall effort is spent in verification. Yet, there are very few credible industry studies that quantitatively provide insight into the functional verification process in terms of verification technology adoption, effort, and effectiveness. To address this dearth of knowledge, a recent world-wide, double-blind, functional verification study was conducted, covering all electronic industry market segments. To our knowledge, this is the largest independent functional verification study ever conducted. This paper presents the findings from our 2014 study and provides invaluable insight into the state of the electronic industry today in terms of both design and verification trends.

Categories and Subject Descriptors

B.6.3 [Logic Design]: Design Aids—Verification

General Terms

Design, Standardization, Languages, Verification

Keywords

Functional Verification Trends, Verification Effort

1. INTRODUCTION

<http://dx.doi.org/10.1145/2745404.2751548> Technical publications often make either subjective or unsubstantiated claims about today's functional verification process—such as, 70 percent of a project's overall effort is spent in verification. [2][5] Yet, there are very few credible industry studies that quantitatively provide insight into the functional verification process in terms of verification technology adoption, effort, and effectiveness. In 2002 and 2004, Collett International Research, Inc. conducted its well-known ASIC/IC functional verification studies, which provided invaluable insight into the state of the electronic industry and its trends in design and verification at that point in time. [3][4] However, after the 2004 study, no additional Collett studies were conducted, which left a void in identifying industry trends. Three private studies were commissioned in 2007, 2010, and 2012, which focused on functional verification. Although the data from these private studies has been referenced in various

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the Owner/Author.

Copyright is held by the owner/author(s).

DAC '15: DAC.com, Jun 07-11, 2015, San Francisco, CA, USA
ACM /15/06.

<http://dx.doi.org/10.1145/2745404.2744921>

publications and blogs, these studies were never officially published.

To address this dearth of knowledge, a recent world-wide, double-blind, functional verification study was conducted, covering all electronic industry market segments. To our knowledge, this is the largest independent functional verification study ever conducted. This paper presents the findings from our 2014 study and provides invaluable insight into the state of the electronic industry today in terms of both design and verification trends.

1.1 Study Background

Our study was modeled after the original 2002 and 2004 Collett International Research, Inc. studies. In other words, we endeavored to preserve the original wording of the Collett questions whenever possible to facilitate trend analysis. To ensure anonymity, we commissioned Wilson Research Group to execute our study. The purpose of preserving anonymity was to prevent biasing the participants' responses. Furthermore, to ensure that our study would be executed as a double-blind study, the compilation and analysis of the results did not take into account the identity of the participants.

For the purpose of our study we used a multiple sampling frame¹ approach that was constructed from eight independent lists that we acquired. This enabled us to cover all regions of the world—as well as cover all relevant electronic industry market segments. It is important to note that we decided not to include our own account team's customer list in the sampling frame. This was done in a deliberate attempt to prevent biasing the final results. In Section 6 we discuss other potential bias concerns when conducting a large industry study and describe what we did to address these concerns.

After data cleaning the results to remove inconsistent or random responses (e.g., someone who only answered “a” on all questions), the final sample size consisted of 1886 eligible participants² (i.e., $n=1886$). To put this figure in perspective, the 2004 Collett study sample size consisted of 201 eligible participants.

Unlike the 2002 and 2004 Collett IC/ASIC functional verification studies, which were conducted only in North America (US and Canada), our study covered all regions of the world. The final study results have been compiled both globally and regionally, but for the purpose of this paper (and due to space) we are presenting only the globally compiled results.

¹ A sampling frame is the list of population elements from which the sample is drawn.

² An example of eligible participant would be a self-identified design or verification engineer, or engineering manager, who is actively working within the electronics industry.

1.2 Confidence Interval

All surveys are subject to sampling errors. To quantify this error in probabilistic terms, we calculate a confidence interval. For example, we determined the overall margin of error for our study to be $\pm 2.19\%$ at a 95% confidence interval. In other words, this confidence interval tells us that if we were to take repeated samples of size $n=1886$ from a population, 95% of the samples would fall inside our margin of error $\pm 2.19\%$, and only 5% of the samples would fall outside.

1.3 Paper Organization

The remainder of this paper is organized as follows. In Section 2, we discuss the study results specifically related to various aspects of design to illustrate growing complexity. In Section 3, we discuss trends in terms of project resources. In Section 4, we examine verification technology adoption trends. Section 5 focuses on verification effectiveness. In Section 6, we describe various bias concerns when conducting a large study such as ours, and then discuss what we did to address these concerns. Finally, in Section 7, we draw some conclusions from our study and discuss other aspects of the verification process that we believe need to be studied in the future.

2. DESIGN TRENDS

In this section, we present trends related to various aspects of design to illustrate growing design complexity. Figure 1 shows the trends from the 2007, 2012, and 2014 studies in terms of participants by design sizes (gates of logic and datapath, excluding memories). The 2014 study added more resolution in identifying larger design sizes (up to 500M or more gates), while the 2012 studies' upper bound was limited to 60M or more gates.




Figure 1. Design Sizes

The key takeaway here is that the electronic industry continues to move to larger designs. In fact, 31 percent of today's designs are over 80M gates, while 17 percent of today's designs are over 500M gates.

But increased design size is only one dimension of the growing complexity challenge. What has changed significantly in design since the original Collett studies is the dramatic movement to SoC class of designs. In 2004, Collett found that 52 percent of designs included one or more embedded processors. Our 2014 study found that the number of designs with embedded processors had increased to 71 percent. Furthermore, 45 percent of all designs today contain two or more embedded processors, while 12 percent of today's designs include eight or more embedded processors. SoC class designs add a new layer of verification complexity to the verification process that did not exist with traditional non-SoC class designs due to hardware and software interactions, new coherency architectures, and the emergence of complex network-on-a-chip interconnect. [1]

In addition to the increasing number of embedded processors contained within an SoC, it is not uncommon to find in the order of 120 integrated IP blocks within today's more advanced SoCs. Many of these IP blocks have their own clocking requirements, which often present new verification challenges due to metastability issues involving signals that cross between multiple asynchronous clock domains. [6]

In Figure 2, we see that 93 percent of all designs today have two or more asynchronous clock domains.




Figure 2. Number of Asynchronous Clock Domains

One of the challenges with verifying clock domain crossing issues is that there is a class of metastability bugs that cannot be demonstrated in simulation on an RTL model. To simulate these issues requires a gate-level model with timing, which is often not available until later stages in the design flow. However, emerging static clock-domain crossing (CDC) verification tools can identify clock domain issues directly on an RTL model at earlier stages in the design flow.

3. RESOURCE TRENDS

In this section we discuss the growing resource trends due to rising design complexity. Figure 3 shows the percentage of total project time spent in verification. As you would expect, the results are all over the spectrum; whereas, some projects spend less time in verification, other projects spend more. The average total project time spent in verification in 2014 was 57 percent, which did not change significantly from 2012. However, notice the increase in the percentage of projects that spend more than 80 percent of their time in verification.




Figure 3. Percentage of Project Time Spent in Verification

Perhaps one of the biggest challenges in design and verification today is identifying solutions to increase productivity and control engineering headcount. To illustrate the need for productivity improvement, we discuss the trend in terms of increasing engineering headcount. Figure 4 shows the mean peak number of engineers working on a project. Again, this is an industry average since some projects have many engineers while other projects have few. You can see that the mean peak number of verification

engineers today is greater than the mean peak number of design engineers. In other words, there are, on average, more verification engineers working on a project than design engineers. This situation has changed significantly since 2007.




Figure 4. Mean Number of Peak Engineers per Project

Another way to comprehend the impact of today's project headcount trends is to calculate the compounded annual growth rate (CAGR) for both design and verification engineers. Between 2007 and 2014 the industry experienced a 3.7 percent CAGR for design engineers and a 12.5 percent CAGR for verification engineers. Clearly, the double-digit increase in required verification engineers has become a major project cost-management concern, and is one indicator of growing verification effort.

But verification engineers are not the only project stakeholders involved in the verification process. Design engineers spend a significant amount of their time in verification too, as shown in Figure 5. In 2014, design engineers spent on average 53 percent of their time involved in design activities and 47 percent of their time in verification.




Figure 5. Where Design Engineers Spend Their Time

However, this is a reversal in the trends observed from the 2010 and 2012 studies, which indicated that design engineers spent more time in verification activities than design activities. The data suggest that design effort has risen significantly in the last two years when you take into account that: (a) design engineers are spending more time in design, and (b) there was a nine percent CAGR in required design engineers between 2012 and 2014 (shown in Figure 4), which is a steeper increase than the overall 3.7 CAGR for design engineers spanning 2007 through 2014. We will discuss a few factors that might be contributing to this increased design effort in upcoming sections.

Figure 6 shows where verification engineers spend their time (on average). We do not show trends here since this aspect of project resources was not studied prior to 2012, and there were no significant changes in the results between 2012 and 2014.

Our study found that verification engineers spend more of their time in debugging than any other activity. This needs to be an important research area whose future solutions will be necessary for improving productivity and predictability within a project.




Figure 6. Where Verification Engineers Spend Their Time

4. VERIFICATION SOLUTION TRENDS

In this section we examine various dynamic verification technology adoption trends.

4.1 Dynamic Verification Techniques

Figure 7 shows the adoption trends for various simulation-based techniques from 2007 through 2014, which include code coverage, assertions, functional coverage, and constrained-random simulation.




Figure 7. Dynamic Verification Technology Adoption Trends

One observation from these adoption trends is that the electronic design industry is maturing its verification processes. This maturity is likely due to the growing complexity of designs as discussed in the previous section. Another observation is that constrained-random simulation adoption appears to be leveling off. This trend is likely due to the scaling limitations of constrained-random simulation. This technique generally works well at the IP block or subsystem level in simulation, but does not scale to the entire SoC integration level.




Figure 8. Formal Technology Adoption

4.2 Static Verification Techniques

Figure 8 shows the adoption trends for formal property checking (e.g., model checking), as well as automatic formal applications (e.g., SoC integration connectivity checking, deadlock detection, X semantic safety checks, coverage reachability analysis, and many other properties that can be automatically extracted and then

formally proven). Formal property checking traditionally has been a high-effort process requiring specialized skills and expertise. However, the recent emergence of automatic formal applications provides narrowly focused solutions and does not require specialized skills to adopt. While formal property checking adoption is experiencing incremental growth between 2012 and 2014, the adoption of automatic formal applications increased by 62 percent. In general, formal solutions (i.e., formal property checking combined with automatic formal applications) are one of the fastest growing segments in functional verification.

4.3 Emulation and FPGA Prototyping

Historically, the simulation market has depended on processor frequency scaling as one means of continual improvement in simulation performance. However, as processor frequency scaling levels off, simulation-based techniques are unable to keep up with today's growing complexity. This is particularly true when simulating large designs that include both software and embedded processor core models. Hence, acceleration techniques are now required to extend verification performance for very large designs. In fact, emulation and FPGA prototyping have become key platforms for SoC integration verification where both hardware and software are integrated into a system for the first time. In addition to SoC verification, emulation and FPGA prototyping are also used today as a platform for software development.

Today, 35 percent of the industry has adopted emulation, while 33 percent of the industry has adopted FPGA prototyping. Figure 9 describes various reasons why projects are using these techniques. You might note that the results do not sum to 100 percent since multiple answers were accepted from each study participant. Also, we are unable to show trend analysis here since previous studies did not examine this aspect of functional verification.




Figure 9. Why Was Emulation or FPGA Prototyping Used?

Figure 10 partitions the data for emulation and FPGA prototyping adoption by the design size as follows: less than 5M gates, 5M to 80M gates, and greater than 80M gates. Notice that the adoption of emulation continues to increase as design sizes increase. However, the adoption of FPGA prototyping rapidly drops off as design sizes increase beyond 80M gates. Actually, the drop-off point is more likely around 40M gates or so since this is the average capacity limit of many of today's FPGAs.

This graph illustrates one of the problems with adopting FPGA prototyping of very large designs. That is, there is an increased engineering effort required to partition designs across multiple FPGAs. In fact, the FPGA prototyping of very large designs is often a major engineering effort in itself, and one that many projects are trying to find alternative solutions to address this problem (e.g., virtual prototyping or actual silicon as a validation platform).




Figure 10. Emulation and FPGA Prototyping Adoption

4.4 Verification Languages and Libraries

In this section, we present the adoption trends for verification language and base-class libraries. As previously noted, the reason some of the results sum to more than 100 percent is that some projects are using multiple languages; thus, individual participants can have multiple answers.

Figure 11 shows the adoption trends for languages used to create testbenches. Essentially, the adoption rates for all languages used to create testbenches are either declining or flat, with the exception of SystemVerilog. Nonetheless, the data suggest that SystemVerilog adoption is starting to saturate or level off at about 75 percent.




Figure 11. Languages Used for Verification (Testbenches)

Figure 12 shows the adoption trends for various testbench methodologies built using class libraries.




Figure 12. Methodologies and Testbench Base-Class Libraries

Here we see a decline in adoption of all methodologies and class libraries with the exception of Accellera's UVM³, whose adoption

³ Universal Verification Methodology (UVM) is a standard to enable efficient development and reuse of verification environments and verification IP (VIP) throughout the electronics industry. Accellera provides both an API standard for UVM and a reference implementation. That reference implementation is a class library defined using the syntax and semantics of SystemVerilog (IEEE 1800).

increased by 56 percent between 2012 and 2014. Furthermore, our study revealed that UVM is projected to grow an additional 13 percent within the next year.

Figure 13 shows the industry adoption trends for various assertion languages, and again, SystemVerilog Assertions seems to have saturated or leveled off.




Figure 13. Assertion Language Adoption

5. VERIFICATION RESULTS

In section 3, we provided data that suggest a significant amount of effort is being applied to functional verification. An important question the various studies have tried to answer is whether this increasing effort is paying off. In this section, we present verification results findings in terms of schedules, number of required spins, and classification of functional bugs.




Figure 14. Design Completion Compared to Original Schedule

Figure 14 presents the design completion time compared to the project's original schedule. The data suggest that in 2014 there was a slight improvement in projects meeting their original schedule, where in the 2007 and 2012 studies, 67 percent of the projects were behind scheduled, compared to 61 percent in 2014. It is unclear if this improvement is due to the industry becoming more conservative in project planning or simply better at scheduling. Regardless, meeting the originally planned schedule is still a challenge for most of the industry.




Figure 15. Required Number of Spins

Other results trends worth examining relate to the number of spins required between the start of a project and final production. Figure 15 shows this industry trend from 2007 through 2014. Even though designs have increased in complexity, the data suggest that projects are not getting any worse in terms of the number of required spins before production. Still, only about 30 percent of today's projects are able to achieve first silicon success.

Figure 16 shows various categories of flaws that are contributing to respins. Again, you might note that the sum is greater than 100 percent on this graph, which is because multiple flaws can trigger a respin.




Figure 16. Types of Flaws Resulting in Respins

Logic and functional flaws remain the leading causes of respins. However, the data suggest that there has been a slight improvement in this area over the past seven years.

Figure 17 examines the root cause of logical or functional flaws (previously identified in Figure 16) by various categories. The data suggest design errors are the leading cause of functional flaws, and the situation is worsening. In addition, problems associated with changing, incorrect, and incomplete specifications are a common theme often voiced by many verification engineers and project managers.




Figure 17. Root Cause of Functional Flaws

6. MINIMIZING STUDY BIAS

When architecting a study, three main concerns must be addressed to ensure valid results: *sample validity bias*, *non-response bias*, and *stakeholder bias*. Each of these concerns is discussed in the following sections, as well as the steps we took to minimize these bias concerns.

6.1 Sample Validity Bias

To ensure that a study is unbiased, it's critical that every member of a studied population have an equal chance of participating. An example of a biased study would be when a technical conference surveys its participants. The data might raise some interesting questions, but unfortunately, it does not represent members of the population that were unable to participate in the conference. The

same bias can occur if a journal or online publication limits its surveys to only its subscribers.

A classic example of sample validity bias is the famous Literary Digest poll in the 1936 United States presidential election, where the magazine surveyed over two million people. This was a huge study for this period in time. The sampling frame of the study was chosen from the magazine's subscriber list, phone books, and car registrations. However, the problem with this approach was that the study did not represent the actual voter population since it was a luxury to have a subscription to a magazine, or a phone, or a car during The Great Depression. As a result of this biased sample, the poll inaccurately predicted that Republican Alf Landon versus the Democrat Franklin Roosevelt would win the 1936 presidential election.

For our study, we carefully chose a broad set of independent lists that, when combined, represented all regions of the world and all electronic design market segments. We reviewed the participant results in terms of market segments to ensure no segment or region representation was inadvertently excluded or under-represented.

6.2 Non-Response Bias

Non-response bias in a study occurs when a randomly sampled individual cannot be contacted or refuses to participate in a survey. For example, spam and unsolicited mail filters remove an individual from the possibility of receiving an invitation to participate in a study, which can bias results. It is important to validate sufficient responses occurred across all lists that make up the sample frame. Hence, we reviewed the final results to ensure that no single list of respondents that made up the sample frame dominated the final results.

Another potential non-response bias is due to lack of language translation, which we learned during our 2012 study. The 2012 study generally had good representation from all regions of the world, with the exception of an initially very poor level of participation from Japan. To solve this problem, we took two actions:

1. We translated both the invitation and the survey into Japanese.
2. We acquired additional engineering lists directly from Japan to augment our existing survey invitation list.

This resulted in a balanced representation from Japan. Based on that experience, we took the same approach to solve the language problem for the 2014 study.

6.3 Stakeholder Bias

Stakeholder bias occurs when someone who has a vested interest in survey results can complete an online study survey multiple times and urge others to complete the survey in order to influence the results. To address this problem, a special code was generated for each study participation invitation that was sent out. The code could only be used once to fill out the survey questions, preventing someone from taking the study multiple times or sharing the invitation with someone else.

6.4 2010 Study Bias

While architecting the 2012 study, we did discover a non-response bias associated with the 2010 study. Although multiple lists across multiple market segments and across multiple regions of the world were used during the 2010 study, we discovered that a single list dominated the responses, which consisted of

participants who worked on more advanced projects and whose functional verification processes tend to be mature. Hence, for this paper we have decided not to publish any of the 2010 results as part of verification technology adoption trend analysis. However, we did include the 2010 data in mean number of peak engineers per project analysis, shown in Figure 4, and the percentage of time a design engineers spends in design versus verification, shown in Figure 5—and the reader can ignore the 2010 results in these two data points if they choose.

The 2007, 2012, and 2014 studies were well balanced and did not exhibit the non-response bias previously described for the 2010 data. Hence, we have confidence in talking about general industry trends presented in this paper.

7. CONCLUSION

This paper presents the findings from a recent world-wide, double-blind, functional verification study, covering all electronic industry market segments. To our knowledge, this is the largest functional verification study that quantitatively provides insight into today's functional verification process in terms of verification technology adoption, effort, and effectiveness.

The data our study reveals is certainly of value, but it does not represent all challenges associated with SoC design (such as SoC integration verification and system validation).⁴ In fact, many of the techniques used for block and subsystem verification that we studied do not scale well to the SoC integration and system-level validation space (e.g., constrained-random, functional coverage, and general formal property checking). In addition, our study does not encompass ESL or virtual prototyping. We believe that future studies should be expanded to include these emerging challenges.

Finally, we believe that the benefit from an industry study is not necessarily the quantitative values that the answers reveal but the new questions they raise and the healthy dialogue that ensues.

8. ACKNOWLEDGMENTS

The authors would like to thank Larry and Zachary Wilson (from Wilson Research Group), and Merlyn Brunken (from Mentor Graphics Corporation) for their expertise and guidance in conducting very large industry studies.

9. REFERENCES

- [1] ARM® AMBA® 5 CHI (Coherent Hub Interface) Specification, 2013, ARM Ltd.
- [2] J. Bergeron, *Writing Testbenches: Functional Verification of HDL Models*. Kluwer Academic Publishers, January 2000.
- [3] R. Collett, "2002 IC/ASIC functional verification study," Industry Report from Collett International Research, Inc. 2003.
- [4] R. Collett, "2004 IC/ASIC functional verification study," Industry Report from Collett International Research, Inc. 2005.
- [5] H. Foster, *Applied Assertion-Based Verification: An Industry Perspective*, Foundations and Trends® in Electron Design Automation, vol 3, no 1, pp 1-95, 2009.
- [6] R. Ginosar, "Metastability and Synchronizers: A Tutorial," IEEE Design & Test, Issue No.05 - Sept/Oct (2011 vol.28)

⁴ Due to space limitation, we have decided not to present power management verification trends from our study in this paper, and will publish them at some future point in time.