

ACHIEVING SCALABLE HARDWARE VERIFICATION WITH
SYMBOLIC SIMULATION

A DISSERTATION

SUBMITTED TO THE DEPARTMENT OF ELECTRICAL ENGINEERING

AND THE COMMITTEE ON GRADUATE STUDIES

OF STANFORD UNIVERSITY

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS

FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY

Valeria Bertacco

August 2003

© Copyright by Valeria Bertacco 2003
All Rights Reserved

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

Kunle Olukotun
(Principal Adviser)

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

David Dill

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

Mark Horowitz

Approved for the University Committee on Graduate Studies.

Abstract

In the past 40 years, electronic systems have become pervasive in modern society. Digital integrated circuits (ICs) are at the heart of the large majority of these systems. These digital ICs are complex systems containing millions of interconnected transistors in a very small area. Moreover, the underlying semiconductor fabrication technology used to fabricate these ICs allows doubling the number of transistors in the same area approximately every 18 months.

The design of digital systems is a complex and time consuming process that progresses through various phases and levels of abstraction, and relies heavily on CAD (Computer-Aided Design) software tools. Within this context, ensuring the correctness of these digital systems is a major consideration, especially because the cost of failures is becoming increasingly high. One of the most famous recent examples of its importance is the Intel, Inc. Pentium's flaw in the floating point divide unit of 1994 that eventually forced Intel to replace all the Pentium chips that were already in the market. In many cases, the possibility of failure is even unacceptable, examples of these applications are: transportation systems, medical applications and financial systems. Even though guaranteeing the correctness of a design is such a central aspect in its development, current verification methodologies are still inadequate to tackle the complex systems that are being developed nowadays. Hardware design companies try to compensate for mediocre CAD tools by dedicating the majority of their resources involved in a design to verification, yet are still unable to guarantee correct functionality over the entire design space.

Logic simulation is the most widely accepted method for ensuring the correctness of digital ICs in industry because of its scalability, flexibility and predictable run-time behavior. This technique is

based on verifying a digital system by providing sequences of binary values for each of the inputs of the system and checking that the corresponding outputs are correct, based on what the design team expected or described in a specification document. However, because of its inherent approach, this validation technique usually can visit only a small fraction of all the possible configurations of a system - also called the state space - and thus the discovery of bugs heavily relies on the expertise of the designer of the test stimuli to select a few crucial configurations to verify. Symbolic simulation is another verification method that is attracting increasing interest because it allows the verification engineer to explore all, or a major portion, of a circuit's state space without the need to design time-consuming test stimuli. However, this approach poses a high demand on the resources of the simulating host, and in particular, on the memory system, because of the complexity of the algorithms involved and their unpredictable run-time behavior. Thus, the scalability of this approach has been the main limiting factor to its mainstream deployment and so far its scope has been limited to small systems.

This thesis presents new symbolic simulation based approaches to the verification problem that radically improve scalability. We present two new techniques that narrow the performance gap between the complexity of digital systems that are being developed and the limited ability to verify them. The first technique, Cycle-Based Symbolic Simulation, is a unique combination of formal methods and logic simulation that can stimulate a circuit with a very large number of input combinations and sequences in parallel. The key concept is the use of a parametric form to represent the set of states visited during simulation. This approach maintains a high degree of scalability, in line with current cycle-based logic simulation techniques, while achieving better efficiency. To better exploit the use of parameterization in improving the memory profile of simulation, the second technique, Disjoint Support Decomposition Based Symbolic Simulation, exploits the disjoint support decomposition properties of the state functions. We develop a new algorithm that exposes the disjoint decomposition properties of a Boolean function by restructuring its BDD representation. The new algorithm is very efficient in the sense that it has worst-case complexity that is only quadratic in the size of the initial BDD, while previous algorithms had exponential complexity in the size of

the function's support. We deployed this algorithm to find the disjoint support decomposition of the state functions in symbolic simulation. By restructuring the next-state functions using their disjoint support components, it is possible to gain better insight about the role of each input variable. Consequently, the next-state functions can be transformed into a simpler parametric form without sacrificing simulation accuracy. Both of these techniques have been tested on the ISCAS benchmark suite. The results show that the first technique can simulate very large trace sets in parallel, maintaining a simulation speed and memory profile that are much closer to logic simulation. The second technique is effective in reducing the memory requirements of symbolic simulation while maintaining exact state exploration.

Acknowledgements

I would like to first thank my graduate advisor Kunle Olukotun. Throughout these years, he has always been prompt and available in supporting whatever direction of research and of life I decided to pursue. In our technical interactions, he would always go straight to the results of my work and challenge me on their practical contribution to the quality of verification for industrial scale digital designs. David Dill has been the person I could always go to for bouncing ideas and have illuminating technical discussions. When my ideas could survive his dissecting analysis, I knew I could publish them. On a personal level, I always admired his bluntness that would eliminate a lot of useless conversation in our interactions. Thanks also to Mark Horowitz for being part of my defense committee and reviewing this thesis even though it is not central to his research area.

My years in Synopsys have played a central role in shaping my understanding of design verification as an industrial challenge first and a research area later. My colleagues have been crucial in providing me with invaluable opportunities: Ghulam Nurie, Swami Venkat and the marketing team of the Vera Group allowed me to interact with customers in meetings that have always been enlightening in my quest towards an understanding of the needs of the hardware designers. Pei-Hsin Ho, my manager in the Advanced Technology Group of Synopsys, gave me the chance to be part of a high-profile technical team and take part in seminars and technical conferences, all while never losing sight of the objective of providing solutions for the design industry. Most of all, he showed me how to efficiently achieve technology transfer, taking academic research and deploying it in software solution for the hardware design community. I would like to thank my other colleagues in the Advanced Technology Group, in particular: Stephen Edwards, Thomas Shiple, James Kukula,

David Cyrluk, Tony Ma, Kevin Harer, Jerry Taylor, Randy Harr and Robert Damiano.

I spent the past year at Stanford completing my PhD work. During this time I shared my office with John Davis. John has created a very positive work environment for me, he has always provided good advice and been very helpful. He has been crucial especially during the preparation of my oral defense talk for which he provided countless bits of advice, asked me all the most difficult questions and forced me to rehearse it until it would flow seamlessly, by which point he could give the talk himself. I would also like to thank all the people that supported me by making available all those resources that are involved in putting together a thesis. My thanks go especially to: Charlie Orgish, Darlene Hadding, Lance Hammond and Azita Emami-Neyestanak. My undergraduate advisor, Maurizio Damiani, first introduced me to research and to the area of Computer Aided Design for integrated circuits. I would like to thank him for the numerous interactions and collaborations that lasted long after my undergraduate studies and spurred many of the publications that led to this research work. The material presented in Chapters 4 and 5 has been shaped by many months of intense discussions with him.

On a personal level, I would like to thank my parents for teaching me the first concepts of mathematics and logic and for introducing me early in my life to pursuing both education and industry experience, contrary to the Italian tradition of completing all the studies before gaining any work experience. I also want to thank my family for supporting my choices in my path through life. My brother Livio provided all sort of technical support and advice and solved many system crashes, most often connecting from some remote location around Europe. Finally, I thank all my friends in the Stanford community who provided me with enthusiastic social entertainment during my years at Stanford.

As this work comes closer to completion, I look forward to new research in the years to come that will hopefully both have practical use and be intellectually stimulating. Thus, I see this dissertation more as a stepping stone in my research work than as the end of my efforts.

Contents

Abstract	v
Acknowledgements	ix
1 Introduction	1
1.1 Functional validation	2
1.2 Formal verification	3
1.2.1 Symbolic simulation	4
1.3 Contributions of the thesis	5
1.4 Organization of the thesis	6
2 Design and verification of digital systems	9
2.1 The design flow	10
2.2 RTL verification	14
2.3 Boolean variables and functions and their representation	17
2.3.1 Binary Decision Diagrams	18
2.4 Models for design verification	21
2.4.1 Structural network model	21
2.4.2 State diagrams	23
2.4.3 Mathematical model of Finite State Machines	25
2.5 Functional validation	26

2.6	Formal verification	31
2.6.1	Symbolic Finite State Machine traversal	32
2.7	Symbolic Simulation	35
2.7.1	The algorithm	37
2.7.2	The challenge in symbolic simulation	41
3	Cycle-Based Symbolic Simulation	43
3.1	Parametric transformations	43
3.2	Parameterizations in symbolic simulation	46
3.3	The CBSS algorithm	47
3.4	The parameterization phase	49
3.4.1	Using functional dependencies	50
3.4.2	How to classify the components of the state vector	53
3.4.3	The <code>remap</code> function	57
3.5	Implementation and complexity	59
3.6	Experimental results	61
3.7	Conclusion	65
4	Disjoint Support Decompositions	67
4.1	Introduction	68
4.2	Related work on Disjoint Support Decompositions	69
4.3	Terminology	71
4.3.1	Decomposition trees.	73
4.4	The unique maximal Disjoint Support Decomposition	73
4.4.1	Decomposition by prime functions.	75
4.4.2	A characterization of F/K_F	78
4.4.3	The normal Decomposition Tree	86
4.5	On the decomposability of Boolean functions	92

5	A novel algorithm for Disjoint Support Decompositions	95
5.1	Building the decomposition bottom-up	96
5.2	Case 1. Neither A_{10} nor A_{11} is constant and $A_{10} \neq \overline{A_{11}}$	99
5.3	Case 2. Exactly one of A_{10}, A_{11} is constant	103
5.4	Case 3. $A_{10} = \overline{A_{11}}$ and A_{10} is not a constant	106
5.5	New decompositions	108
5.6	Putting it all together: The DSD procedure	117
5.6.1	Inherited decompositions	120
5.6.2	New decompositions	125
5.7	Complexity analysis and considerations	127
5.8	Experiments on the decomposability of industrial testbenches	129
5.9	Conclusion	136
6	Exact Parameterizations for Symbolic Simulation	137
6.1	Re-encoding the state function	138
6.2	Reduction at Free Points	140
6.3	Elimination of Prime functions	143
6.4	Removal of non-dominant variables	146
6.5	DSD-SS Implementation	150
6.6	Experimental results	151
6.7	Summary	155
7	Conclusion	157
7.1	Parameterized approaches in symbolic simulation	157
7.2	Disjoint support decompositions	158
7.3	The future of this work	159
	Bibliography	161

List of Tables

3.1	Cycle Based Symbolic Simulation results	62
5.1	Disjoint Support Decomposition results	131
6.1	Disjoint Support Decomposition-based simulation results	154

List of Figures

2.1	Design flow of a digital system	11
2.2	Approaches to verification	16
2.3	Binary Decision Diagrams	19
2.4	Graphic symbols for some basic logic gates	22
2.5	Structural network model	22
2.6	Network model of a 3-bits up/down counter with reset	23
2.7	State diagram of a 3-bit up/down counter	24
2.8	State diagram of a 1-hot encoded 3-bit counter	25
2.9	Compiled logic simulator	27
2.10	Logic simulation - pseudocode	29
2.11	FSM state traversal - pseudocode	35
2.12	Logic and symbolic simulation	36
2.13	Symbolic simulation algorithm - pseudocode	38
2.14	Symbolic simulation for Example 2.8 - Initialization phase	38
2.15	Symbolic simulation for Example 2.8 - Simulation Step 2	40
2.16	Iterative model of symbolic simulation	40
3.1	Parameterization of the state vector during symbolic simulation	44
3.2	Three steps of symbolic simulation for the counter of Example 2.2 and possible parameterizations of the reached state sets	46

3.3	Cycle-Based symbolic simulation flow	49
3.4	The CBSS algorithm - pseudocode	50
3.5	The parameterized frontier subset $\mathbf{PS}_{@k}$	51
3.6	<code>parameterize</code> function - pseudocode	53
3.7	Classifying simple and complex variables - pseudocode	55
3.8	Classifying shared variables - pseudocode	57
4.1	Decompositions for Example 4.1	68
4.2	A decomposition tree for Example 4.3.	74
4.3	Decomposition representation of the function of Example 4.7	88
4.4	Decomposition tree for Example 4.8.	89
5.1	PRIME decomposition.	111
5.2	Function for Example 5.6.	115
5.3	Two functions and the construction of their $Max(G,H)$ tree.	127
6.1	A decomposed state vector for a small design	139
6.2	The parameterized frontier set $\mathbf{PS}_{@k}$	140
6.3	A vector function and its free points	141
6.4	Free points elimination for Example 6.1	143
6.5	General case for prime function elimination: (a) before and (b) after	145
6.6	Prime elimination for Example 6.2.	146
6.7	Non-dominant variable removal for Example 6.4	150
7.1	Trade-offs of in the breadth vs. scalability plane	159