

High Level Power Estimation and Reduction Techniques for Power Aware Hardware Design

Sumit Ahuja

Dissertation submitted to the Faculty of the
Virginia Polytechnic Institute and State University
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy
in
Computer Engineering

Sandeep K. Shukla, Chair
A. Lynn Abbott
Michael S. Hsiao
Patrick Schaumont
Anil Vullikanti

May 12, 2010
Blacksburg, Virginia

Keywords: Electronic System Level, Hardware Software Co-design, High Level Synthesis,
Model Checking, Power Estimation, Power Reduction, Clock-gating, Coprocessors.

©Copyright 2010, Sumit Ahuja

High Level Power Estimation and Reduction Techniques for Power Aware Hardware Design

Sumit Ahuja

(ABSTRACT)

The unabated continuation of the Moore's law has allowed the doubling of the number of transistors per unit area of a silicon die every 2 years or so. At the same time, an increasing demand on consumer electronics and computing equipments to run sophisticated applications has led to an unprecedented complexity of hardware designs. These factors have necessitated the abstraction level of design-entry of hardware systems to be raised beyond the Register-Transfer-Level (RTL) to Electronic System Level (ESL). However, power envelope on the designs due to packaging and other thermal limitations, and the energy envelope due to battery life-time considerations have also created a need for power/energy efficient design. The confluence of these two technological issues has created an urgent need for solving two problems: (i) How do we enable a power-aware design flow with a design entry point at the Electronic System Level? (ii) How do we enable power aware High Level Synthesis to automatically synthesize RTL implementation from ESL?

This dissertation distinguishes itself by addressing the following two issues: (i) Since power/energy consumption of electronic systems largely depends on implementation details, and high-level models abstract away from such details, power/energy estimation at such levels has not been addressed thoroughly. (ii) A lot of work has been done in applying various techniques on control-data-flow graphs (CDFG) to find power/area/latency pareto points during behavioral synthesis. However, high level C-based functional models of various compute-intensive components, which could be easily synthesized as co-processors, have many opportunities to reduce power. Some of these savings opportunities are traditional such as clock-gating, operand-isolation etc. The exploration of alternate granularities of these techniques with target applications in mind, opens the door for traditional power reduction opportunities at the high-level.

This work therefore concentrates on the aforementioned two areas of inadequacy of hardware design methodologies. Our proposed solutions include utilizing ESL simulation traces and mapping those to lower abstraction levels for power estimation, derivation of statistical power models using regression based learning for power estimation at early design stages, etc. On the HLS front, techniques that insert the power saving features during the synthesis process using exploration of granularity and scope of clock-gating, sequential clock-gating are proposed. Finally, this work shows how to marry two domains, that is estimation and reduction. In this regard, a power model is proposed, which helps in predicting power savings obtained using clock-gating and further guiding HLS to selectively insert clock-gating.

This project received partial financial support from NSF CAREER, an NSF CRCD grant, and industrial funding from Cebatech Inc.

Dedication

To my family - Papa, Mummy, Bhaiya, Bhabhi and Kashvi

the friends and relatives

and

all my fellow Hokies.

- Sumit Ahuja

April 25, 2010, Blacksburg, Virginia (USA).

Acknowledgments

I would like to acknowledge various people with whom I was directly/indirectly associated with during the course of my PhD.

First and foremost, I would like to acknowledge my advisor, Dr. Sandeep K. Shukla for his guidance, motivation, encouragement, support, and friendship throughout my graduate studies at Virginia Tech. He is a great mentor and his commitment to work has been a source of great admiration and inspiration for me. I would like to thank Dr. Patrick Schaumont, Dr. Michael S. Hsiao, Dr. Lynn Abott and Dr. Anil Vullikanti for serving on my PhD committee. I would also like to thank Deepak Mathaikutty, Wei Zhang, Ajit Dingankar and Ramana Jampala for their mentorship, guidance and help. I would like to acknowledge the support received from Cebatech Inc., NSF-CRCO, NSF CAREER grants, which provided funding for the work reported in this dissertation.

Special thanks are due to my friends Luv Kothari and Deepak Mathaikutty for motivating and guiding me during the course of PhD. To Wei Zhang for numerous hours of technical discussions and guidance during my PhD, and for being patient and supportive through those discussions. To Gaurav Singh and Avinash Lakshminarayana for helping and discussing many research topics related to this thesis. To all my roommates, friends and colleagues from Virginia Tech: Hiren Patel, Syed Suhaib, Nitin Shukla, Debayan Bhaduri, Bijoy Jose, Mainak Banga, Harini Jagadeesan, Mahesh Nanjundappa, Bin Xue, Hua Lin, Maheshwar Chandrasekar, Vipul Chawla, Anupam Srivastava and Ananya Ojha. To all the other friends that have come and gone - thanks to all. To all my friends and batchmates in Europe and India.

Finally, I would like to thank my parents Ram Swaroop Ahuja, Uma Ahuja, my brother Vishal Ahuja and sister Amita Ahuja for their continued love, support and encouragement in my endeavors, and for always being there for me.

Contents

1	Introduction	1
1.1	The Need for Low-Power Design Methodologies	1
1.2	Trends in System Design with System Level Modeling, and High Level Synthesis	2
1.2.1	The Need for High Level Synthesis	2
1.2.2	Low Power Design and High Level Modeling	3
1.2.3	Current Power Aware Design Methodology at the High Level	4
1.2.4	Our Power Aware Design Methodology at the High Level	6
1.3	Overview of the Proposed Solutions	7
1.3.1	Application of the proposed techniques	9
1.4	Dissertation Organization	10
1.5	Author's Publications	11
1.5.1	Related Publications as a Co-Author	12
2	Related Work	14
2.1	High Level Power Estimation	14
2.1.1	Spreadsheet Based Approaches	14
2.1.2	Power Estimation Approaches utilizing Power Models	15
2.1.3	Power Macro-model based Approach	16
2.1.4	Summary of Power Estimation Research	18
2.2	High Level Synthesis	19
2.2.1	High Level Synthesis from C/C++ Specifications	20

2.2.2	High Level Synthesis from behavioral specifications	22
2.2.3	Summary of HLS Research	24
2.3	Power Reduction at the RTL and High Level	24
2.3.1	Summary - Low-Power High-Level Synthesis Work	28
3	Background	30
3.1	Average Power components	30
3.2	PowerTheater	31
3.3	FSMD Modeling using GEZEL	32
3.3.1	An Example	32
3.4	ESTEREL	33
3.4.1	Esterel Studio	33
3.5	Multi-variate least squares linear regression model	34
3.5.1	Theoretical Model	34
3.6	Transaction-level Modeling	34
3.7	High Level Synthesis using C2R	37
3.8	Power Reduction using Clock-gating	39
3.8.1	Clock-gating	39
3.8.2	Sequential Clock-gating	40
4	Statistical Regression Based Power Models	44
4.1	Regression based Power Model for FSMDs	46
4.2	Steps for Power Modeling	47
4.2.1	Learning Phase	47
4.2.2	Utilization Phase	50
4.3	Results and Conclusions	50
4.3.1	Tool Flow	50
4.3.2	Experimental Results	50
4.3.3	Discussion	53

5	High Level Simulation Directed RTL Power Estimation	56
5.1	Introduction	56
5.2	Rationale for Our Approach	57
5.3	Our Methodology	59
5.3.1	Activity Extraction from the High-level VCD	61
5.3.2	High-level Variable to RTL Signal Mapping	62
5.4	Results	64
6	Applying Verification Collaterals for Accurate Power Estimation	67
6.1	Introduction	67
6.2	Our Methodology	68
6.2.1	Assertions for Finding out Particular Mode of Design	70
6.2.2	Extraction of Modes from the Simulation Dump	71
6.2.3	High-level Power Estimation	72
6.3	Case Study	73
6.3.1	Directed Testbench Creation from the Specification	74
6.3.2	Utilizing Assertions for creating Directed Testbench	75
6.4	Results	78
7	Power Reduction using High-Level Clock-gating	81
7.1	Introduction	81
7.2	Why is clock-gating needed at High-level?	83
7.3	How to enable Clock-gating at High-Level?	85
7.3.1	Application of clock-gating for various granularities at the source-code level	85
7.3.2	Priority for clock-gating decisions	86
7.3.3	Algorithm for enabling clock-gating in C based co-processor synthesis framework	88
7.4	Results	89
7.4.1	Clock-gating Exploration at High-Level	90

7.5	Summary	91
8	Model-Checking to exploit Sequential Clock-Gating	92
8.1	Introduction	92
8.2	Our Approach and sample properties	94
8.2.1	An Illustrative Example	94
8.2.2	Our Approach	95
8.3	Our Proposed Methodology	97
8.3.1	Changes applied on the RTL model	98
8.3.2	Automation of the flow	100
8.4	Results and Conclusions	100
8.4.1	Extending the framework for HLS	101
9	System Level Simulation Guided Approach for Clock-gating	103
9.1	Introduction	103
9.2	Power reduction model	104
9.2.1	Activity based power reduction model	104
9.2.2	Power reduction model with technology specific information	106
9.3	Rationale for Our Approach	107
9.4	Our Methodology	109
9.4.1	TLM interface for the generated CATL SystemC Model	112
9.5	Results	114
10	Conclusion and Future Work	118
	Bibliography	123

List of Figures

1.1	Power Savings at various abstraction levels [1].	4
1.2	Current design methodology used at the high-level	5
1.3	Proposed design flow	6
1.4	Dissertation Organization	13
3.1	System modeling discussed by Cai. et. al. [2]	35
3.2	System modeling discussed by Frank Ghenassia [3]	36
3.3	Typical register representation in hardware	39
3.4	Clock-gated register representation in hardware	40
3.5	An example of de-asserting the datapath registers to find out the sequential clock-gating opportunities [4]	40
3.6	After applying the power reduction technique on the example circuit [4]	41
3.7	An example of de-asserting the datapath registers to find the sequential clock gating opportunities [4]	42
3.8	An example of de-asserting the datapath registers to find the sequential clock gating opportunities [4]	43
4.1	Conceptual mapping of activity to the real design.	46
4.2	<i>SCoPE</i> Methodology for creating and utilizing power models of FSMDs	48
4.3	Various tools used in different phases of the <i>SCoPE</i> methodology	51
4.5	Comparison of measured power and predicted power for AES at 180nm	52
4.4	Comparison of measured power and predicted power for AES at 90nm	53
5.1	Detailed view of our power estimation methodology	60

6.1	Detailed view of system-level power estimation technique	68
6.2	Power estimation for each mode at RTL	69
6.3	Macro-state diagram of PSM	73
6.4	Specification of the different PSM ingredients	74
7.1	Flow chart for gating clocks of registers at high-level	87
8.1	Our proposed methodology for applying sequential clock-gating at RTL	98
9.1	Traditional clock-gating methodology for HLS based design flow	109
9.2	Pictorial Representation of Our Methodology	111
9.3	Detailed view of system-level interconnect for auto-generated CATL models with transaction level testbench	114

List of Tables

4.1	Comparison of measured power (90 nm) and predicted power for various designs .	52
4.2	Comparison of measured power (180 nm) and predicted power for various designs .	52
5.1	Results of our approach on different designs	65
5.2	VeSPA processor instruction wise power numbers	65
6.1	Percentage of simulation time spent in each state at system-level	79
6.2	Percentage of simulation time spent in each transition	79
6.3	power estimation of the states at RTL	79
6.4	Power estimation at RTL for each transition	80
6.5	power numbers at each state using our approach	80
6.6	Percentage of simulation time spent in each transition	80
6.7	Estimation accuracy as compared to RTL	80
7.1	Percentage power reduction as compare to the original design without clock-gating	89
7.2	Comparison of our approach with automated RTL clock-gating	90
8.1	Percentage power reduction summary on the tested Designs	101
8.2	Percentage power reduction summary at high-level	102
9.1	Percentage Power reduction summary on the tested Designs	115
9.2	Speedup with Our Approach	116
9.3	Number of Registers not to clock gate under various thresholds	117

Chapter 1

Introduction

1.1 The Need for Low-Power Design Methodologies

Managing power consumption in complex System-on-Chips (SoCs) and custom processors is emerging as the biggest concern for designers. These SoCs and processors are vital parts of embedded systems used in electronic equipments ranging from laptop computers, mobile phones to personal digital assistants (PDA), etc. Quality/Performance of such devices is not only measured by what kind of functionalities these devices are capable of performing, but also how long these gadgets survive without plugging into power outlets for recharging. There is an increased demand for reduced form factor or lower size of such devices, which is generally determined by the size of the battery used to run these devices. Another need for controlling power consumption of the design is to decrease power density of devices on a single die. A survey report [5] shows that power density of some designs can reach the power density of a rocket nozzle in future, if power aware design approaches are not used. The reasons cited above have motivated engineers to look into power aware design methodologies.

Traditionally the hardware design flow starts from the register transfer level (RTL) whereupon the design is synthesized to the gate-level, and after further processing at the lower abstraction levels a chip is manufactured. Methodologies for power estimation and optimization at the RTL or lower-level are well studied. In the recent past, modeling at higher level of abstraction was advocated for quick turn around. Increased demand and reduced time-to-market requirements for electronic equipments have forced designers to adopt high-level modeling in their design-flow. Design methodologies, where reduction of power consumption of such systems can be done at the high-level, have become the need of the hour. We know that any optimization requires good analysis of the design; and so is the case with managing the power consumption of the design. To reduce the power consumption in the final product, it is necessary to have an estimate of the power consumption early on. It is becoming essential to utilize design methodologies to reduce and estimate power consumption at the high-level.

1.2 Trends in System Design with System Level Modeling, and High Level Synthesis

Advances in semiconductor technology can be well represented by *Moore's law* [6], which states that every 18 months or so, almost double the number of transistors can be accommodated on the same die size¹. Because of such advances, every new generation of microelectronic systems can accommodate more computing elements than before. With the increased availability of computing elements, more and more applications are mapped onto the hardware. Increased demand for hand held electronic equipments capable of utilizing more applications require increased focus on system-on-chip (SoC) design approach. In an SoC design, a lot of reusable hardware blocks or intellectual properties (IPs) are integrated together. *Electronic System Level (ESL)* based design approaches are becoming necessary to design complex SoCs.

ESL methodologies are mainly targeted to improve the productivity of the designer. These methodologies help in reducing the design cycle time of the hardware design by adopting approaches that can be applied at the higher-levels. Some of the reasons why ESL is getting traction are the need for fast prototype development, quick functional verification, and easy architectural exploration. Most of these system models at such high-levels are very abstract and many details of the real hardware are not available. This level often lacks crucial information such as pipelining required to achieve appropriate throughput, selection of appropriate hardware blocks to perform the desired computation such as choice of various multipliers, finite impulse response (FIR) filter, etc.

High-level models are generally written by algorithm/system developers in C/C++. Traditionally, once these algorithms are fixed, hardware design starts from scratch using Verilog/VHDL. Hardware generation from the algorithmic description of the design is still more or less an unachieved goal. An algorithm developer generally thinks about the execution of a model using C/C++, which is sequential by nature; while hardware is inherently parallel. The gap between software implementation stage (algorithmic description) and synthesizable hardware implementation (RTL description) is often very large. In traditional design methodologies, hardware modeling at the RTL requires effort ranging from several months to a few years, depending upon the size and complexity of the hardware designs.

1.2.1 The Need for High Level Synthesis

High-level specifications allow the designer to ignore various RTL coding issues (such as synchronization and scheduling of operations) while capturing the functionalities of the target design. As stated above, High-level to RTL conversion is very tricky because each algorithmic functionality can have multiple implementations. Modeling of hardware designs at high-level has gained wide acceptance in functional verification, generation of quick simulation models and virtual prototyping. Various efforts to shorten the time consuming transformation of algorithmic description to

¹more recently this period has roughly become 2 years

RTL are addressed by the High-Level Synthesis (HLS) techniques [7, 8, 9]. Most of the high-level synthesis solutions convert the high-level specification model to the implementation model (such as RTL model). There are various levels of abstraction on which high-level specifications can be captured starting from the transaction level modeling (TLM) using cycle-accurate view, architecture view, programmer view, etc. Models at a higher-level (such as programmer view) contain lesser details of the hardware implementation, while at a lower-level (such as cycle-accurate view) contains more details [3].

High Level Synthesis (e.g. [10, 11]) allows the designer to write high-level (above RTL) descriptions of a design that can be automatically converted into synthesizable RTL code [1]. High-level code describes the operation on different data types of a hardware design without specifying the required resources and schedule cycle by cycle. Thus, high-level synthesis offers the advantage of automatic handling of the scheduling and synchronization issues of a design, which helps in faster architectural exploration leading to reduced design time and increased productivity.

A typical high-level synthesis process starts by generating the intermediate representations in the form of various flow graphs such as Control Data Flow Graphs (CDFGs) [12], Hierarchical Task Graphs (HTGs) [13], etc. These representations help in capturing the data dependencies and control flow available in the high-level specifications. Various optimizations and transformations are then performed on such flow graphs before generating the RTL code. High Level Synthesis can be performed using a variety of high-level input specifications expressing the behavioral description in different model of computations such as term rewriting systems [8], communicating sequential process [14], synchronous data flow [15], etc. Moreover, different internal representations are used by different synthesis tools to represent these input high-level specifications before converting them into the RTL code. The input specification for HLS is generally provided as a variant of C/C++ such as SystemC [16], Handle-C [17], etc. Apart from the C/C++ based languages, other languages for which high level synthesis support exists include Bluespec System Verilog [8], Esterel [18], etc.

1.2.2 Low Power Design and High Level Modeling

Each stage of the design flow (shown in Figure 1.1) represents a different modeling level based on the need of the designer. For example a high-level stage generally entails the functional description of the design, whereas the RTL includes data flow between registers and logical operation performed on the signals. Gate-level design description include logic-gates (such as NAND, NOR, etc.) and flip flops connected to each other using wires. Layout-level view contains information on how different transistors are connected with each other in various layers of metals, silicon and other materials. This view is very close to real silicon implementation and contains a lot of design information such as geometric shapes of metal, oxide, semiconductor, etc. Every stage of the design represents the same functionality of the design but expressed in different format, style, etc.

At every stage, designer might have an opportunity to reduce power consumption of the design. However, any small change at the lower level requires more effort to explore possibilities and to

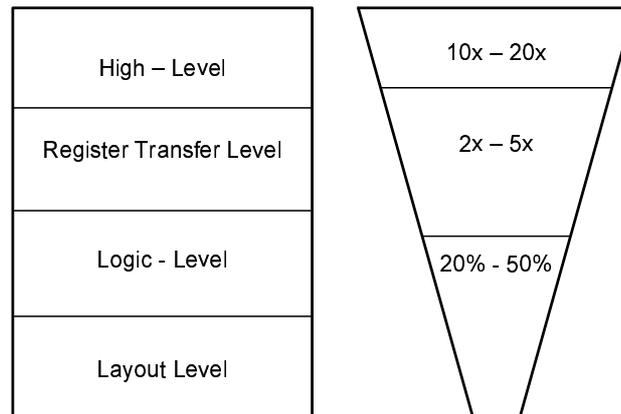


Figure 1.1: Power Savings at various abstraction levels [1].

validate the design (because of the increased amount of design information), thus increasing the complexity of optimization. Also, at the start of the hardware design process, the designer has more alternatives. At the RTL, micro-architecture of the design is fixed. Fixing micro-architecture implies fixing the data flow between registers, this leaves very few opportunities to optimize the design between/across registers. Based on the constraints on clock-frequency, area and power, RTL synthesis tools employ some techniques to reduce the power consumption of the design. At the logic and layout-level designers/tools are left with even lesser choices because structure (connection of various transistors) of the design is fixed. “However, no amount of skillful logic design can overcome a poor micro-architecture” [19]. Figure 1.1 presents a documented user experience [1] in reducing power consumption of designs at various stages of the design flow. As shown in the figure most power saving opportunities exist at the highest level of design abstraction.

1.2.3 Current Power Aware Design Methodology at the High Level

Current design methodology used at the high-level is shown in Figure 1.2 (for simplicity we show power aware design flow, other aspects of a typical design flow are omitted). A high-level design flow starts from a specification. Through rigorous simulations and other verification techniques, a system-level model that conforms to the design intent is constructed. As we know, the RTL design contains a lot more design information than a high-level design, hence different automated (high-level synthesis) or manual refinement techniques are applied to arrive at the desired RTL implementation (i.e., FSM encoding, micro-architecture variations, etc). Once the RTL implementation is stable, power estimation of the design is performed and various constraints are evaluated. If the power and other constraints are met, then the netlist is generated. Otherwise further refinements are applied at the RTL or even at higher levels in order to reduce the power consumption of the design.

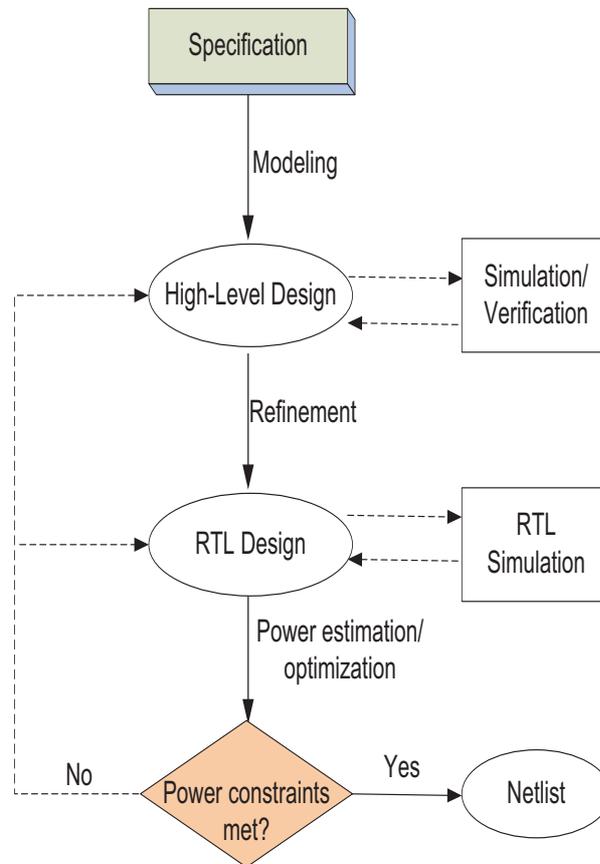


Figure 1.2: Current design methodology used at the high-level

Problems with the Traditional Methodology

As shown in Figure 1.2, power estimation is performed at the RTL stage, while refinements to reduce power are applied both at the high-level and the RTL. It may be too late to perform power estimation and optimization at the RTL as opportunities to tweak the design to meet the necessary design constraints are limited. High-level synthesis endowed with RTL or lower-level power optimization features (such as clock-gating, operand-isolation) can instead be utilized in traditional design methodology to exploit the maximum benefits. This will reduce the overhead of repeating the power optimizations at the RTL. This will also help in taking power aware design decisions early in the design flow. There are two main problems in the traditional design methodology:

- Lack of support for accurate power estimation at the high-level
- Lack of support to utilize the RTL power reduction techniques during high-level synthesis

1.2.4 Our Power Aware Design Methodology at the High Level

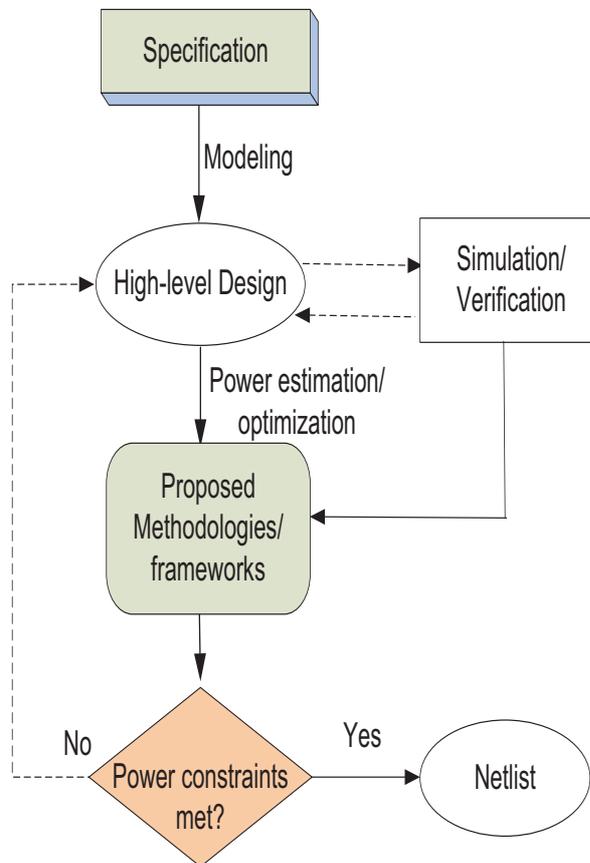


Figure 1.3: Proposed design flow

As solutions to the problems posed above, we propose a design methodology shown in Figure 1.3. Our proposed design methodology enables power estimation and reduction features at the high-level. Our solutions also include ways in which a high-level model can guide RTL techniques, in case complete migration to a high-level methodology is not possible. We propose the following solutions for both the problems:

- Enable Power Estimation at the High Level:
 - By creating the accurate power model for the hardware design above RTL
 - By utilizing the high-level verification collaterals
 - By utilizing the high-level simulation to guide the RTL power estimation
- Enable Power reduction during High-Level Synthesis:

- By utilizing the RTL power reduction techniques during HLS such as clock-gating, operand-isolation, sequential clock-gating etc.
- By enhancing and utilizing power models to guide the HLS to generate clock-gated RTL.

Our proposed methodologies integrate high-level design languages & frameworks. These methodologies are independent of the:

- High-level modeling languages (C, SystemC [16], ESTEREL [20], etc.)
- High-level synthesis tools (Catapult-C [7], Cynthesizer [21], C2R [22], Esterel Studio [18], etc.)
- RTL power analysis tools (PowerTheater [23], Prime-Power [24], etc)

However, to experiment with our solutions, we have used various frameworks available in academia or industry such as GEZEL [25], C2R [22], Esterel Studio [18], PowerTheater [23], etc.

1.3 Overview of the Proposed Solutions

1. High Level Power Estimation Techniques

- (a) **Power Estimation using Statistical Power Models:** A power model is a model that captures the dependence of power consumption of a design block on certain parameters such as switching activity, capacitance, etc. Accuracy of the power model is very much dependent on the model of computation, input/output activity, capacitance, etc. In order to invent a suitable power model for designs, which are modeled as finite state machines with datapath (FSMD). We first analyze its models of computation and propose a power model based on the activity of each state. We capture the relationship between average power and activity of different states of FSMD as a linear regression, to model the power consumption of FSMD based designs [26, 27].
- (b) **Power Estimation using High-level Simulation:** Dynamic power is a very important component of total power consumption of a design. It is proportional to the activity/toggles of the design. For power estimation purposes, this activity is generally measured from the simulation of the design. RTL/gate-level simulation of big designs takes a lot of time. These simulation dumps for big designs are typically very large (with the value reaching upto 200 GBytes in some of our experiments). RTL/Lower-level simulation is one of the biggest bottleneck for power estimation because it is very time consuming. To solve this problem, we propose algorithm and methodology to utilize high-level simulation to guide the RTL power estimation techniques. Our methodology helps in significantly reducing the turn around time to obtain the power numbers with reasonable accuracy [28].

- (c) **Power Estimation using High-level Verification Collaterals:** Verification of a design is very time consuming process. Various verification collaterals such as tests, checker, assertions are used at the early design stages to check the functionality of the design. These collaterals help in finding the appropriate stimulus to check various design states, modes, etc. Measurement of power consumption is very much stimulus dependent. However, there are not many ways in which high-level verification collaterals can be targeted for accurate power estimation tasks. We propose a methodology to apply high-level verification collaterals for accurate power estimation [29, 30].

2. High Level Synthesis endowed with Power-saving Features

- (a) **Power Reduction using Clock-gating:** Clock-gating is a well known technique for reducing power consumption of the design. This technique is applied to to reduce the clock toggles of a register. Clock power is a major component of total power consumption, which makes clock-gating a very important power saving technique. RTL synthesis tools are capable of finding the clock-gating opportunities for a design, if the RTL model conforms to certain modeling guidelines. Since HLS is becoming increasingly important, there is a need to push such lower-level reduction techniques to a higher-level. Here, clock-gating opportunities that can be enabled in a high-level model are investigated. Various granularities of clock-gating at a behavioral level are examined such as function, scope or variable level. An algorithm for clock-gating, to automatically generate clock-gated RTL after HLS is also presented [31].
- (b) **Power Reduction using Sequential Clock-gating:** There are some techniques such as sequential clock-gating, which require a capability to provide power saving control information across clock boundaries. Verification of optimizations, which are applied across clock-cycles is very difficult. We use model checking to find out the relationship between two registers, and then find out the possibility of sequential clock-gating. Such an approach helps in discovering the relationship between two arbitrary registers. This relationship can then be utilized for sequential clock-gating. Existing solutions are limited to pipe-lined RTL implementations, while our approach is not limited to any particular design-style [32]. Model checking based approach helps in reducing the complexity of verification task, which otherwise is very difficult and limits application of such techniques to particular design style such as pipelined designs. Our approach is further extended to examine such an opportunity to determine sequential optimization opportunities at high-level.
- (c) **System level simulation guided HLS approach to generate clock-gated RTL:** Clock-gating in the current tools/methodologies is statically done i.e. HLS tool is informed to insert clock-gating for specific part of the design or whole design. Clock-gating may not save dynamic power all the time especially for finer granularities. This savings/wastage is also dependent on the stimulus. In current state of the art methods it is very difficult because power estimation is very time consuming and clock-gating is performed at netlist level, making a decision making task very difficult. Here, a power

model capable of checking the efficacy of clock-gating is proposed. This power model indicates how much power savings are possible with clock-gating. Once the indication from the power model is available, an HLS tool is guided to selectively apply clock-gating and generate the new RTL [33].

1.3.1 Application of the proposed techniques

In this dissertation, we propose various techniques that can be utilized with HLS or high-level design flows. Here we briefly discuss the application of the proposed techniques.

Early Power Estimation:

The approaches discussed in Chapters 4 and 6 are applicable in scenarios where designers want to measure power consumption of the target hardware using models that can be utilized at higher abstraction levels than the RTL. Given that RTL or gate-level power estimation is extremely time consuming, such techniques can save time if proven to be accurate enough for crucial design decisions. Chapter 4 presents a technique to characterize power model that is parameterized by the switching activity of each state of Finite State Machine with Datapath (FSMD) model of the design. Such a power model can then be simulated with high-level models. Chapter 6 shows utilization of directed test vectors and assertions. One could use such vectors or assertions to bring a high level simulation to a target state to measure state specific power consumption.

Since power consumption is very much design and technology dependent, it might be difficult to create power models that characterize power consumption of a target hardware at high-level. To perform power estimation at the RTL, the biggest bottlenecks are simulating, processing and extraction of the activity information. Chapter 5 presents a technique where power estimation is performed using an RTL power estimation tool but the simulation is done at a high-level. The resulting dump is interpolated to create required activity profile.

Power Reduction from high-level:

We consider clock-gating as a primary candidate for power reduction. We investigate how this can be enabled from high-level. At high-level, a designer is working on a behavioral description of the design. In this case he/she would need to insert clock-gating from the behavioral description before synthesis. This approach is presented in Chapter 7. Similarly, sequential clock-gating requires optimizations across the register boundaries. Such optimizations may cause a lot of verification problems. In chapter 8, we propose model-checking based approach utilizing invariants for sequential clock-gating specific optimization. In chapter 9, we show how system-level simulation can guide HLS to select the appropriate clock-gating candidates among all the registers. This approach is also useful to find out if aggressive application of clock-gating will lead to wastage.

1.4 Dissertation Organization

1. Chapter 2 discusses related research work done in the domain of high-level power estimation, high-level synthesis and power aware high-level synthesis.
2. Chapter 3 provides a detailed explanation of various background topics (such as components of average power, GEZEL, ESTEREL, transaction level modeling, C2R based synthesis flow, etc.) relevant to this dissertation.
3. Chapter 4 presents a technique to obtain early power estimation using vectorless RTL power estimation technique guided by high-level simulation.
4. Chapter 5 presents a technique for power estimation using power models of FSMD based design. This chapter details the analysis based on new power model which utilizes datapath activity of each state for power estimation.
5. Chapter 6 presents a technique which utilizes verification collaterals for accurate power estimation. This chapter presents a case study which helps in understanding verification collaterals needed for accurate power estimation. These collaterals include directed testbenches, assertions; which helps in obtaining state specific power numbers.
6. Chapter 7 presents power reduction using clock-gating at the high-level. This chapter provides detail on how clock-gating can be enabled in high-level specification. This chapter also presents an algorithm for clock-gating used in high-level synthesis.
7. Chapter 8 provides an overview of sequential clock-gating based power optimization. We show how model-checking can be utilized to find the opportunities for sequential clock-gating from the behavioral description and further utilizing the indications after checking few properties of the design to generate sequentially clock-gated RTL.
8. Chapter 9 presents an approach where power model guides power reduction. We present a power reduction model for clock-gating based power saving technique. This power-model is embedded in cycle-accurate SystemC view of the behavioral model. This model gives an indication on the registers that can save power and further guides the HLS to generate the clock-gated RTL. This chapter presents the complete methodology and theory for creating the power reduction model.
9. Chapter 10 presents conclusions and the possible extension for the approaches presented in this thesis.

Figure 1.4 shows the organization in a pictorial form along with the separation of the work that is completed and the future work. The organization shown in Figure 1.4 presents the content dependencies to guide the reader in understanding the material that is of interest to them.

1.5 Author's Publications

Below is the list of publications on the work done by the author of this dissertation.

1. Sumit Ahuja, Wei Zhang, Sandeep K. Shukla. System Level Simulation Guided Approach to Improve the Efficacy of Clock-gating. *To appear in the proceedings of IEEE workshop on High Level Design Validation and Test (HLDVT), 2010, Anaheim, CA, June 2010.*
2. Sumit Ahuja, Wei Zhang, Avinash Lakshminarayana, Sandeep K. Shukla. A Methodology for Power-Aware High-Level Synthesis of Co-processors from Software Algorithms. *23rd International Conference on VLSI Design (VLSI 2010), Jan 2010, Bangalore, India, pp. 282-287.*
3. Sumit Ahuja, Deepak A. Mathaikutty, Avinash Lakshminarayana, Sandeep K. Shukla. SCoPE: Statistical Regression Based Power Models for Co-processors Power Estimation. *Journal of Low Power Electronics, Vol. 5, Number 4, December 2009, pp. 407-415.*
4. Sumit Ahuja, Sandeep K. Shukla. MCBCG: Model Checking Based Sequential Clock-Gating. *IEEE workshop on High Level Design Validation and Test (HLDVT), Nov 2009, San Francisco, CA, pp. 20-25.*
5. Sumit Ahuja, Deepak A. Mathaikutty, Avinash Lakshminarayana, S .K. Shukla. Accurate Power Estimation of Hardware Co-Processors using System Level Simulation. *22nd IEEE International SOC Conference, Sep 9-11, 2009, Belfast, Northern Ireland, pp. 399-402.*
6. Sumit Ahuja, Swathi T. Gurumani, Chad Spackman, Sandeep K. Shukla. Hardware Coprocessor Synthesis from an ANSI C Specification. *IEEE Design and Test of Computers, vol. 26, no. 4, July/Aug. 2009, pp. 58-67.*
7. Sumit Ahuja, Deepak A. Mathaikutty, Gaurav Singh, Joseph Stetzer, Sandeep K. Shukla, Ajit Dingankar. Power estimation methodology for a high-level synthesis framework. *10th International Symposium on Quality Electronic Design, Mar 16-18, 2009, San Jose, CA, pp. 541-546.*
8. Sumit Ahuja, Deepak A. Mathaikutty, Sandeep K. Shukla. Applying Verification Collaterals for Accurate Power Estimation. *9th International Workshop on Microprocessor Test and Verification, Dec, 2008, Austin, Texas, pp. 61-66.*
9. Sumit Ahuja, Deepak A. Mathaikutty, S. K. Shukla, A. Dingankar. Assertion-Based Modal Power Estimation. *Eighth International Workshop on Microprocessor Test and Verification, 2007, Austin, Texas, pp.3-7.*

All the work in terms of the research contributions, implementation and experimentation for these publications was done by the author under the guidance of Dr. Sandeep K. Shukla.

1.5.1 Related Publications as a Co-Author

Publications done with the authors other colleagues during the course of this dissertation work are enumerated below

1. J. Brandt, K. Schneider, Sumit Ahuja, Sandeep K. Shukla. The Model Checking View to Clock Gating and Operand Isolation. *to appear in the proceedings of Application of Concurrency to System Design (ACSD), June, 2010, Braga, Portugal .*
2. Gaurav Singh, S. S. Ravi, Sumit Ahuja, Sandeep K. Shukla. Complexity of Scheduling in Synthesizing Hardware from Concurrent Action Oriented Specifications. *under review, 2010.*
3. Avinash Lakshminarayana, Sumit Ahuja, Sandeep K. Shukla. Coprocessor Design Space Exploration Using High Level Synthesis. *11th International Symposium on Quality Electronic Design, Mar 16-18, 2010, San Jose, CA, pp. 879-884.*
4. Gaurav Singh, Jacob B. Schwartz, Sumit Ahuja, Sandeep K. Shukla. Techniques for Power-aware Hardware Synthesis from Concurrent Action Oriented Specifications. *Journal of Low Power Electronics (JOLPE), Vol. 3, No. 2, August 2007, pp. 156-166.*
5. Deepak A. Mathaikutty, Sumit Ahuja, Ajit Dingankar, Sandeep K. Shukla. Model-driven test generation for system level validation. *IEEE International High Level Design Validation and Test Workshop, Nov 7-9, 2007, Irvine, CA, pp. 83-90.*
6. Eric Simpson, P. Yu, Patrick Schaumont, Sumit Ahuja, S. K. Shukla. VT Matrix Multiply Design for MEMO-CODE 07. *International Conference on Formal Methods and Models for Co-Design, MEMOCODE 2007, Nice, France, pp. 95-96 .*

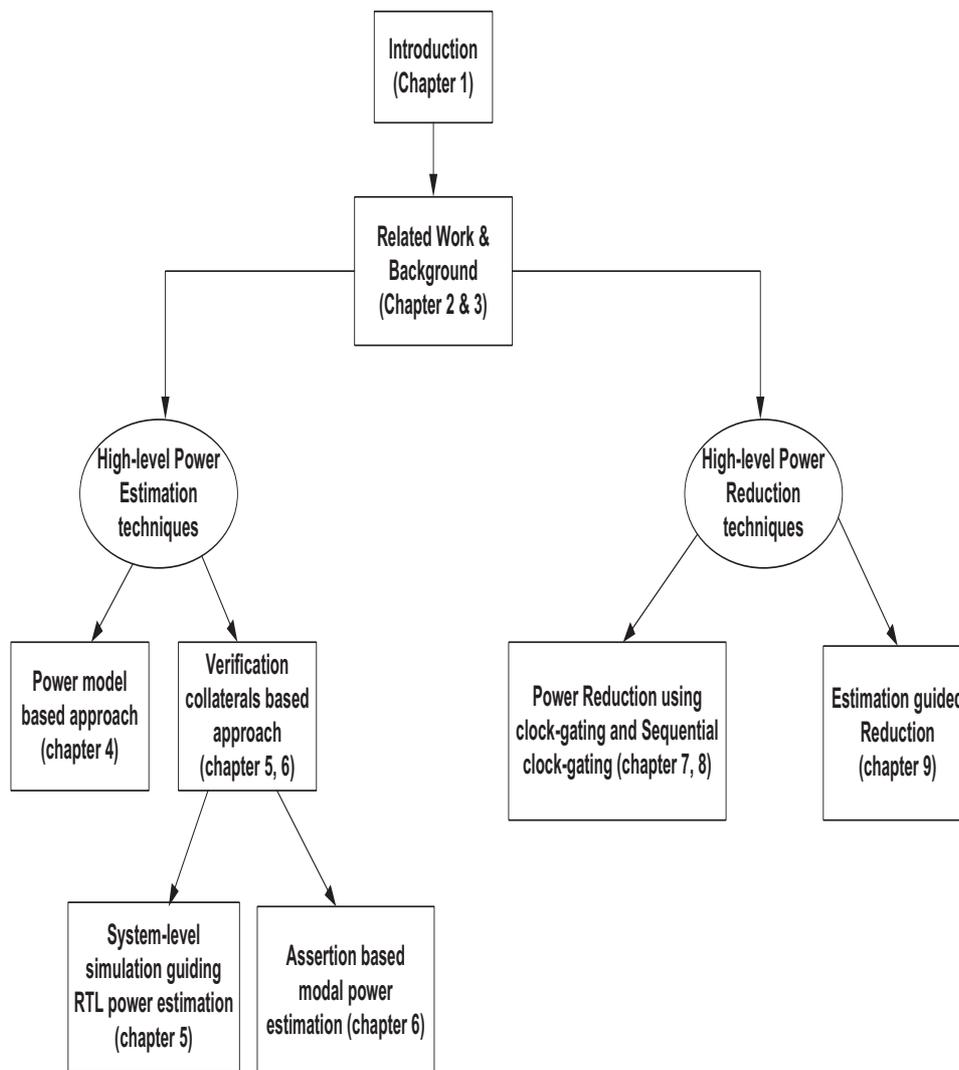


Figure 1.4: Dissertation Organization

Chapter 2

Related Work

In this chapter, we briefly discuss the existing techniques, methodologies in the area of power estimation, high level synthesis and power reduction.

2.1 High Level Power Estimation

Accurate power estimation at high level is very important for any successful design methodology. This area has been extensively researched. In this section, we capture some research advancements relevant to this thesis, which include spread sheet based approach, power model and macro-model based approach, commercial tools available for RTL and gate-level power estimation etc.

2.1.1 Spreadsheet Based Approaches

Spreadsheets are very useful in the early stage of design process, when initial planning is going on and a lot of important decisions are being taken [34]. One of the biggest advantages of spreadsheet based analysis is that the user does not really need to learn any complex/sophisticated tool for taking design decisions. One of the basic application of spreadsheet is area estimation. Designers generally have a fair idea of the building blocks for a big design. He/She can easily get an estimate on area by using data sheets from intellectual property (IP) provider, library cell estimates, etc. Spreadsheet provides a capability to capture such information, which can be utilized for quick area estimation. Similarly, some decisions to control power can also be taken using spreadsheet based approach. Power budgeting approaches using spreadsheets are very helpful for printed circuit board (PCB), power supplies, voltage regulators, heat sink, and cooling systems.

Spreadsheet tools vary from utilizing excel sheets, word processors to Unified Modeling Language (UML) [35] etc. In industry, spreadsheet is being advocated by Field Programmable Gate Array (FPGA) vendors such as Xilinx [36], Altera [37]. Power analysis needs to be done very efficiently

especially for FPGAs, where basic building blocks are fixed (for example fixed size lookup tables and switch matrices). These power estimation tools provide current (I) and power (P) estimation for the different family of FPGAs. Power number associated to each block in these spreadsheets are very much dependent on architectural features such as the clock network, memory blocks, Digital Signal Processing (DSP) blocks. User can enter operating frequency, toggle rates, and other parameters in the early power estimator spreadsheet to estimate the design's power consumption. Spreadsheet based approach is useful for project planning but may not be able to provide accurate guidance for block-level hardware power estimation and reduction. This motivates a need to provide a power model, which can perform accurate yet efficient power analysis early in the design flow. Next subsection provides an overview of some of model based approaches for power estimation purposes.

2.1.2 Power Estimation Approaches utilizing Power Models

One of the first works in the area of model based power estimation was proposed by Tiwari et al. [38] for the architecture level power estimation of microprocessors. They provide a method to estimate power/energy number of a given program on a given processor. The way they have approached the problem was interesting, they have used a hypothesis for their work. This hypothesis states “By measuring the current drawn by the processor as it repeatedly executes certain instructions or short instruction sequences, it is possible to obtain most of the information that is needed to evaluate the power cost of a program for that processor” [38].

Power consumption of a microprocessor can be represented as $P = V_{CC} \times I$, where V_{CC} is operating voltage and I is the current drawn by the microprocessor. In their approach, they measured the current drawn by the processor and then utilize it for power measurement purpose. They assumed that during the computation, the operating voltage for the processor will not change. For average power estimation purpose they had first estimated energy over different cycles and then averaged it. To conclude with, they have proposed a way in which one can gauge the impact on average current from the execution of an instruction. Their approach also proposed a way to measure inter-instruction effects on current values. To calculate the average power consumption of a processor while executing a software program, they utilized these values.

Power modeling for the instructions of VLIW (Very Large Instruction Word) processors is discussed in [39]. Approach discussed by Tiwari et al. in [38] might not be applicable to the processors where number of instructions are reaching to a few thousand. The valid sequences for such processors can reach thousands. Characterization of power instruction based power model can be very time consuming and may take several months. Tiwari et al. in [38] propose a technique for VLIW processors for reducing the complexity of power models by using the clustering approach. They try to make cluster of instructions having energy number in the same range (individual as well as sequence of instructions). They propose an approach to reduce the complexity of characterization from exponential to quadratic. There are many other approaches that have been discussed in the literature for doing power modeling of processor instructions. More details on power estima-

tion utilizing instruction level power models and its variants are available in [40, 41, 42].

Most of the approaches for power models are proposed for CPU or micro architecture of processors. In an ASIC design flow where mainly design stage starts at RTL and mostly design is represented as a Finite State Machine with Datapath (FSMD), similar approach may not be useful. In the literature this is often termed as macro-model based approach. In the sub-section 2.1.3, we present an overview of existing macro-model based power estimation approaches. We have used quite a few concepts from such approaches in this thesis. A quick read of the subsection 2.1.3 will help the reader in understanding the approach discussed in the Chapter 4.

2.1.3 Power Macro-model based Approach

Power macro-models [43] generally consist of n-dimensional table to estimate the power consumed in a circuit for a given statistics, where n represent different variables/components capturing the relationship of power and dependent variables such as input probability, transition density etc. [43] present an automation methodology, in which, such a table for is automatically generated. Three variants in their model are average input signal probability, average input transition density, and average output zero-delay transition density. Power macro-model is a function of these variants, as shown in equation 2.1.

$$P_z = f(P_{in}, D_{in}, D_{out}) \quad (2.1)$$

Where P_z represents the entry for the average power corresponding to the average input signal probability P_{in} , average input transition density D_{in} and average output zero delay transition density D_{out} .

Signal probability P_i at input node x_i is defined as the average fraction of the clock cycle in which final value of x_i is high. Similarly, transition density D_i at an input node is defined as the average fraction of the cycles in which the node makes a logic transition (final and initial value should not be the same for a logical transition). Equation 2.2 represents the relationship between P_i and D_i .

$$D_i/2 \leq P_i \leq 1 - D_i/2 \quad (2.2)$$

Using the relationship discussed in equation 2.2 and equation 2.1 characterization of models is being done and the accuracy of these macro-models is evaluated. For characterization purposes they assume that input nodes are output of latches or flip-flops and make at most one transition per clock cycle. Also, sequential design is single clock system and clock skew is ignored in their analysis hence all the inputs switch simultaneously.

Bogliolo et al. [44] have proposed a methodology for creating power macro-models based on linear regressions but their flow is specific to the structural RTL macros and power estimation is done at the gate-level. Their analysis is restricted to structural RTL representation whose leaf components

are combinational logic blocks. This approach is based on a-) offline characterization in which they compute the power of the RTL macro based on certain tests, b-) online characterization, in which they do it adaptively for error minimization. Their approach utilizes all the inputs, outputs, transition functions of inputs and outputs on the successive cycles, and then they interpolate the relationship with energy consumption.

Potlapally et al. [45] present a technique in which they do cycle-accurate power macro modeling of the RTL component. This technique is based on the fact that RTL components exhibit different “power behavior” for different input scenarios. They create power macro model for each of these behaviors also known as power modes. Their framework chooses the appropriate power mode from the input trace in each cycle and then apply power macro-modeling technique discussed by Bogliolo et al. to get an estimate on power numbers. The technique discussed in [44] is limited to the typical average power estimation scenario while the technique in [45] covers non-trivial scenarios as well.

Negri et al. [46] have proposed a power simulation framework of communication protocols (Bluetooth and 802.11) using StateC. StateC is used to model the hierarchical state machines. Their flow is mainly targeted for simulator generation in SystemC. This flow is good for power exploration of protocol modeling but not presented on ASIC/FPGA design flow such as ours. They have mainly targeted wireless protocols in which relevant contribution to the power consumption of a node is due to the communication and not due to the datapath (computation) activity. Their learning phase requires execution of the real chip and can not easily be integrated to any ASIC/FPGA design flow. Also for practical purposes, it is difficult to create power model of the partial design or smaller part of the whole chip because current measured includes a lot of contribution from the other parts of the chip and isolation of the desired unit for power model purposes requires quite a lot of effort.

In [47], authors attempt to lift power estimation to higher levels than the RTL, and their choice for high-level modeling was Cycle-Accurate Functional Description (CAFD) of the design. They create virtual components for each design block and attach them to the CAFD model of the design block, and compute the power consumption dynamically as the CAFD is simulated. Since this additional overhead to the CAFD simulation causes inefficiency, they also allow periodic turning off of some of the virtual components during some cycles of the simulation. During those cycles, they estimate power based on the history of the power consumption for the turned off components. So even though, the abstraction at which they estimate power is cycle accurate modeling level as ours, their power estimation is not based on regression based technique, and the simulation of CAFD is slower due to the overhead of virtual components.

Caldari et al. presented a relative-power analysis methodology [48] for system-level models of the Advanced Micro-controller Bus Architecture (AMBA) and Advanced High-performance Bus (AHB) from ARM. It relies on creating macro-models from the knowledge of the possible implementations. Similarly, Bansal et al. presented a framework in [49], which uses the power-models of the components available at the system-level simulation stage by observing them at run time. It selects the most suitable power-model for each component by making efficiency and accuracy trade-offs. In [50], the presented framework employs co-simulation techniques for power esti-

mation with the capability of performing accuracy and efficiency trade-offs. They utilize multiple power estimation engines concurrently and have proposed several speed-up techniques to eliminate the overhead of co-simulation.

SEAS [51] presents a system and framework for analyzing SoCs in early design stage. Power Analysis in this system works at a granularity of a processor cores, where pre-characterized data for power is utilized based on the power state of the design. Power states of the cores are high level states based on the workload such as active, idle, sleep states in today's processors. By utilizing these high level states of the SoC an early power estimation can be performed which is more efficient and accurate than the traditional spread sheet based approach.

Shin et al. have proposed a methodology [52] for power estimation of operation modes but their analysis is done at logic-level and proposes a way to create power models based on the switching frequencies. Our approach is different from their approach as we are targeting the modal designs for power estimation at system-level keeping good accuracy in mind. [53, 54, 55, 56, 57, 58] utilize the similar approach for power estimation purposes and provide various accuracy and efficiency trade-offs based on the quality of inputs and power modeling. Power estimation accuracy can be significantly increased but generally it impacts the efficiency of power estimation procedure.

2.1.4 Summary of Power Estimation Research

Techniques discussed above show that power modeling of a hardware block can be very complex and application dependent. As we increase the levels of abstraction, power estimation of a hardware block becomes difficult. Various techniques at different abstraction levels exist to obtain the power consumption starting from spread-sheet, power model to macro-model based power estimation. The most popular approaches in industry are mainly power model based approach or by performing power estimation at RTL/gate-level description of a hardware design. While, at the lower level of abstraction, commercial tools provide good accuracy with respect to silicon, but as we go to higher levels of abstraction accuracy of the power estimation methodologies reduces.

In an ASIC design flow, most of the mentioned power-estimation solutions either rely on the creation of power macro-models from lower level descriptions of the designs or on the availability of power consumption information for different simulation scenarios. The methodologies we propose in this thesis are not dependent on any such infrastructure, thus making it suitable for easier and accurate power analysis at early stages of an ASIC design flow. Previous approaches in this area are focused on creating macro-models for the design based on input activity and relationship with the corresponding power. Such input and output patterns help in modeling power of the design but design is considered as *blackbox* in most of the cases. However, approaches used in this thesis provides explicit visibility on the toggles and corresponding control state and datapath, and can be viewed as a gray box approach.

In this thesis, we also present approaches, where high level fast simulation is guiding lower level tools. Most of the power estimation tools rely on simulation and simulation dump processing,

which takes majority of power analysis time. We investigated techniques that work with high-level simulation to provide the maximum benefits. We experimentally show very high accuracy can be achieved.

2.2 High Level Synthesis

In 1970's transistor based designs were prevalent, 80's gate-level, 90's RTL, after RTL there is no consensus. RTL based designs have increased designer's productivity as newer chips are utilizing more and more transistors in the same design. Managing such a design complexity requires a lot of effort on design and verification side. Engineering budget to include more and more engineers at lower level along with the increasing cost of mask is another compelling reason to automate the process from the earliest possible design stage. This brings a need for Electronic System Level (ESL) based design approach, where design, verification and prototyping is envisioned at a level of abstraction where designer's concern is to take care of the best possible algorithm to complete the task while other tasks such as creation of hardware, functional verification, validation, etc. will be seamlessly integrated to the design flow. One of the most important tasks for achieving this goal is to be able to produce hardware from C/C++ based specification or the specifications which are above register transfer level (RTL) abstraction of the design.

High-level synthesis is an emerging area. There are lot of tools and methodologies exist in the industry and academia. Most of these tools are in infancy stage. In hardware design industry, adoption of these tools has recently started. High-level synthesis of hardware designs has been shown to be possible from a variety of high-level languages. These languages are used to specify the behavior of a design at a level of abstraction above RTL. Hardware descriptions written using such HDLs are passed as inputs to high-level synthesis tools to generate the RTL code. There are various ways high-level synthesis can be enabled but most of the contributions are in the area of C/C++ based synthesis.

The other motivations to designing at high-level with respect to Verilog, VHDL for hardware design are:

- Most of the HDLs are foreign to algorithm developers (they generally prefer C/C++)
- Increasing level of abstraction using HDLs is very difficult. As most of the times, without the availability of FSM/FSMD, a good RTL cannot be developed for the design
- In practice, HDL based design flow is time consuming. Generally, the golden model (from which functional correctness is verified for the lower level models) comes as an executable, which is C/C++ based implementation of the design. Hence an automatic path from the C/C++ implementation would first reduce the time. Secondly, once flow is stable, it will reduce a lot of verification effort, because generated RTL needs to be validated every time with respect to functional specification for any small manual change done at RTL

2.2.1 High Level Synthesis from C/C++ Specifications

Here, we briefly touch upon the existing high-level synthesis tools and methodologies used in industry and academia generating hardware from C/C++ specifications. In this section, we cover in detail the methodology presented by CatapultC and PICO tool. We also provide a brief overview of some other tools.

CATAPULT: In Catapult synthesis [7] flow various manual methods are automated to reduce the time to produce the hardware. The flow is centered around Catapult HLS tool, where micro-architectural selection is done based on the constraints (provided by the designer/user). Tool creates the RTL based on these constraints. Other important point to note is that one can specify the target technology used, clock-period or clock frequency as constraints. Such information helps the scheduler of the HLS to schedule and allocate resources based on the constraints. Below, we briefly discuss various important points of Catapult based design methodology.

- *Verification of generated RTL against original C code:* Functional verification of the generated RTL is one of the most important stage of the design flow. Here, RTL is wrapped around a SystemC transactor, where it is called as foreign module. By doing this, original C++ testbenches can be compiled with this SystemC top module instantiating generated RTL module and finally comparator is used to compare the outputs. The wrapper code along with the makefiles is auto-generated to complete the verification flow.
- *Synthesis Constraints for Catapult Flow:* There are two types of constraints that can be provided in the Catapult flow. First set of constraints are related to target technology, clock-frequency, etc. Second set of constraints are utilized to control the architecture. These constraints can be inserted using GUI or using directives. These directives facilitate loop-unrolling, loop pipelining, hardware interface synthesis, etc. These constraints are not encoded in the source code; hence appropriate micro-architectures are created during synthesis stage.
- *C++ and optimization support:* Catapult supports the pure C++ based specification as input. Most of the C++ constructs are supported by the tool except the code, which requires dynamic memory allocation/de-allocation such as use of malloc, free, new, delete. In other words code should be statically deterministic so that all the properties, memory allocation can be performed during compile time. Catapult supports pointer synthesis, classes and templates, bit-accurate data-types, etc. On the optimization side Catapult C supports loop pipelining, loop-merging, loop unrolling, technology driven resource allocation and scheduling, etc.

Design flow for PICO tool:

Program In Chip Out (PICO) tool's [59] approach is mainly targeting SoC platforms. Before going to the discussion of tool it will be helpful to see the characterization of different IPs in the context

of this tool and their approach. These IPs are mainly characterized in four categories as discussed below:

- *Star IPs* : Star IPs are CPUs/DSP blocks and generally these blocks are fixed for many generation of SoCs. Most of the times these IPs are manually created and their instruction level characteristics are well defined. During the development of SoC various models of such IPs are utilized at different granularities, which includes Instruction level simulation model, RTL/gate-level model, etc. These IPs are not altered for design changes.
- *Complex Application Engines*: Especially for embedded systems, complex application IPs such as video-codec, wireless modem, etc. are differentiating factor for the end product. These codecs or modems generally change as new standards are introduced. Functionality of these complex IPs require small to significant changes for a new SoC development and in many cases direct utilization of these IPs may not be advisable.
- *Connectivity and Control IPs*: This set of IPs include connectivity and control IPs, examples of such IPs include DMA, USB, etc. These IPs are generally utilized in communication and it can be considered as system-level glue. Their functionality is not needed and requires very minimal tailoring.
- *Memory*: Memory is generally the biggest contributor to silicon area and generally at system-level their functionality is not defined. It's functionality generally does not differentiate the end product. Memory models are compiled and built bottom-up.

All the IPs discussed above are essential for SoC development but complex application engine generally require the bulk of effort for design and verification purposes.

Programming Model of PICO

PICO can accept the sequential C specification and tries to extract the parallelism from the sequential C code. In that sense their approach is very much specific to particular application domain such as signal processing applications. Because in such domains a lot of parallelism is available while application is processed by the hardware. Such a programming model is useful where a part of function has no dependency between the different tasks. If one task works on a block of data and other works on another block of data, then a lot of parallelism can be extracted in pipelined manner. The execution model of PICO is based on Kahn Process Network (KPN), where set of sequential processes communicate via streams with block-on-read and unbounded buffering. Here processes are the hardware blocks, which communicate with each other through streams. The restriction on unbounded buffering is a big restriction, which is basically solved by imposing additional constraints on the execution model.

GAUT: GAUT [60] is an academic high-level synthesis tool based on C as design language but applicable for digital signal processing applications. GAUT starts from bit-accurate specification written in C/C++ and extracts the possible parallelism before going through the conventional stages

such as binding, allocation and scheduling tasks. GAUT has mandatory synthesis constraints, which are throughput and clock-period. GAUT utilizes the algorithmic C libraries (from mentor graphics) for bit-accurate integer and fixed-point data-types using `ac_int` and `ac_fixed` data types. HLS in GAUT utilizes the DFG generated from gcc and library characterization for particular target library for area, speed, etc. to generate the RTL VHDL or cycle accurate SystemC. The synthesis process is not completely technology independent but can be useful for virtual platform development for micro-architectural analysis.

SPARK: SPARK [13] presents a high-level design methodology based on the high-level synthesis tool, which takes Behavioral C as input design language and capable of generating RTL VHDL. Main contributions of the methodology based on SPARK are: a-) inclusion of code motion and code transformation techniques at compiler level to include maximum parallelism for high-level synthesis, b-) proposal of high-level synthesis starting with behavioral C input. The approach presented in SPARK helps designer in understanding how and by what amount the quality of results can get affected by the language level transformations such as loop unrolling, loop-invariant code motion, etc. on generated circuit from HLS. Also the approach presented in SPARK based methodology suggests that no single code transformation approach is universally applicable but such techniques with heuristics applied for particular application domain leads to better quality or results. The transformations and techniques applied in this methodology include exploitation of instruction-level parallelism, such as speculative code motions, percolation scheduling and trail-blazing, etc.

C2R: C2R [22] is also a tool, which takes ANSI C based description as input and capable of producing synthesizable RTL Verilog as output. C2R benefits the designer in two folds : a-) facilitate algorithmic developer to remain in C/C++ native execution environment, b-) provide directives to include concurrency, implementing interfaces using C functions, exploiting micro-architectural parallelism, concurrency inside a function/process, facilities to exploit pipelining behavior, etc. Functional verification of the design is performed on the generated RTL output by comparing its output against the output generated from the pure ANSI C based implementation. Other advantage tool provides is instantiation of complex memories in behavioral style in pure ANSI C environment. Such an approach helps SoC development where parts of the designs are IPs coming from other vendors and some can be developed in-house quickly. Finally, Cebatech [22] provides a methodology to designer to exploit the multi-threaded C code using pthread and convert it into Verilog RTL.

2.2.2 High Level Synthesis from behavioral specifications

Bluespec [61] provides high-level synthesis solution based on bluespec system verilog (BSV), which is based on atomic transactions. The compiler is capable of generating synthesizable HDL from BSV. The reason for using BSV instead of C/C++ is because they are not considered appropriate for certain control dominated applications. Such applications include processors, caches, interconnects, bridges, etc. Secondly, atomic transaction based model of computation suits well to

some of the application mentioned above, which otherwise may be very complex if implemented using conventional HDL based approach. For verification BSV based methodology can either utilize the simulation based approach or formal verification based approach, which can be applied to the atomic model of computation, such verification methods are based on term rewriting systems, SPIN model checker, etc. However, designers need to learn a new language to develop the hardware using Bluespec.

Cynthesizer [62] provides the high-level synthesis capability around SystemC transaction level models. The main contributions of this tools comes in the synthesis of transaction level models. On the verification/simulation side, since the input language is SystemC, for the synthesizable subset of SystemC designer gets the flexibility to utilize the C++ execution environment. SystemC has support for processes using `sc_thread`, `sc_method`, `sc_thead`, etc. Also SystemC helps designer to write efficient interfaces for communication between computation blocks. Such a high-level synthesis approach is very helpful in system design where designer wants to utilize IPs/hardware blocks and utilize various communication interfaces to develop SoC. C++ execution helps in quick verification and generation of interface and hardware blocks to RTL quickly, which helps in improving the design cycle efficiency.

Apart from above mentioned contributions, various techniques for hardware synthesis from many C-like languages have been proposed [63, 64]. Most techniques either use a subset of C or propose extensions to C in order to synthesize the high-level code to RTL. Cones [65] synthesis system from AT&T Bell Laboratories allows the designer to describe the circuit behavior during each clock cycle of a sequential logic. It takes a restricted subset of C (with no unbounded loops or pointers) as input for hardware synthesis. Similarly, HardwareC [66] is a C-like language having a cycle-based semantics with support for hardware structure and hierarchy. It is used as an input to Olympus synthesis system [67]. Celoxica's Handel-C [17] is a superset of ANSI-C with support for parallel statements and rendezvous communication. SpecC [68, 69], which is a system-level design language based on ANSI-C, has an explicit support for hierarchy, concurrency, state transitions, timing and exception handling. SpecC's synthesizable subset is used to generate hardware designs. Other tools and languages, which are mainly used in high-level modeling and for which synthesis support exists include Celoxica's Agility Compiler [9], NEC's Cyber [70]. ESTEREL [71] is a synchronous language, allows the description of a design at high-level, which can then be synthesized using Esterel Studio [18]. ESTEREL supports implicit state machines and provides constructs for concurrent composition, preemption, and exceptions [72, 73].

There are various approaches proposed in software synthesis area by Jose et al. The proposed techniques in this thesis or other related works can be adopted for such synthesis frameworks as well. Synchronous languages such as Esterel, Lustre, etc. can also be used to generate embedded software [74]. Synthesis tools such as Esterel Studio and SCADE are used to generate control software in safety critical applications like avionics, power plants and so on. An alternate programming model called polychrony, which deals with multi-rate signals whose event occurrences are independent of the software execution rate is also used for generating embedded software. The programming language SIGNAL [75] based on relational models or the actor based network MRICDF [76, 77] can be used to capture polychronous specifications. These

programming formalisms are associated with their software synthesis tools Polychrony [78] and EmCodeSyn [79, 80] respectively. To obtain efficient embedded software implementations on distributed environment, the generation of multi-threaded code [81, 82] and the communication between distributed synchronous systems are also being researched [83].

2.2.3 Summary of HLS Research

A lot of initial work for HLS started in academia but unfortunately the quality of the hardware generated was not very good. However, in the recent past more and more design wins are shown in industry using HLS tools. HLS tools are maturing day-by-day. Initial research in this area was about finding good model of computation and language to represent the hardware. A lot of tools have shown success in producing hardware from the behavioral description in various different modeling languages. At the time of writing this thesis, I can think of atleast two viewpoints for different languages to represent the design at high-level:

- Use C/C++ as design language because most of the specifications in practice come in C/C++. A lot of companies have kept this viewpoint in mind and have introduced HLS tools, which require minimal code restructuring to generate synthesizable RTL. Such examples include tools like CatapultC, PICO, C2R etc.
- Use the behavioral description amenable to verification or particular model of computation. This helps in creating hardware, which can be easily verified in the later design stages. Also, it greatly helps in resolving issues, which exist in the behavioral description in C/C++ itself but it introduces an intermediate step of creating another description in different language. Examples include Bluespec compiler for Bluespec system verilog , Forte Cynthesizer for SystemC, Esterel Studio for Esterel language etc.

In this thesis, we have used tools from both the viewpoints. We have used Esterel Studio and C2R. We have shown power measurement and reduction approaches for these tools. Most of the proposed approaches in this thesis, can be seamlessly integrated with tools accepting C/C++ or some other behavioral description as the modeling language.

2.3 Power Reduction at the RTL and High Level

Substantial research has been done in the area of power reduction at the RTL and higher-levels of abstractions. Clock-gating is implemented at the RTL by tools such as PowerTheater [84], Power compiler [85], etc. PowerTheater suggests opportunities to clock-gate registers for which there is no multiplexer in the feedback path. They find the conditions under which clock is required when and only when there is a change in data input to a register. Power compiler recommends particular RTL coding styles to enable clock-gating during synthesis.

Such local optimization opportunities are often found dynamically through simulation, and they do not exist in HLS. This is because (i) in HLS enable signals for clock-gating purpose can be generated at the compile time without any dynamic analysis or symbolic simulation, (ii) global enable for the registers is easily visible at the high-level because at lower level conditions under which a register is updated can not be determined, such as conditional execution of blocks using ‘if-then-else’, ‘while’ etc. [86] provides an overview of clock-gating technique and how tools like power compiler use it. They also show the advantages of clock-gating for register banks and how it helps in the place and route stage of the design. Many of these concepts are used in our work as well albeit at a higher level. However, this work requires RTL simulation, generation of VCD, and subsequent analysis, making the entire process extremely time consuming. At the higher level an effort is required to make the entire process few orders of magnitude faster, while power savings should be better or equal atleast.

High-level synthesis from C-like HDLs commonly use Control Data Flow Graphs (CDFGs) as intermediate representations during the synthesis process, and consequently, most research in the area of low-power high-level synthesis is targeted towards the CDFG-based synthesis. In the past, various power optimization techniques targeting the power reduction during synthesis have been proposed. High-level synthesis for C-like HDLs include stages such as scheduling, allocation and binding. Various techniques are proposed for different stages to affect the power consumption of the design once the RTL is created from HLS. Scheduling of various operations of a design can be exploited for generating power-efficient designs. The problem of resource-constrained scheduling for low-power has been addressed in [87, 88]. These approaches use CDFGs to first determine the mobility of various operations based on the ASAP and ALAP schedules. Using the computed mobilities and other relevant factors, priorities are assigned to various operations. Based on the assigned priorities, various operations of the design are then scheduled in each clock cycle such that the power consumption of the design is reduced.

During the allocation phase of a high-level synthesis process, functional units and registers are allocated to the design, whereas in the binding phase, operations and variables of a design are mapped to the allocated functional units and registers respectively. [89, 90, 11, 91] present techniques targeting low-power reduction during allocation and binding phases. [89] presents an allocation method for low-power high-level synthesis, which selects a sequence of operations for various functional units such that the overall transition activity of the design is reduced. [90] presents an algorithm targeting the minimization of the average power of a circuit during the binding phase of high-level synthesis process using game-theoretic techniques. In that work, binding of the operations to the functional units is done such that the total power consumption is minimized. [11] presents an efficient register-sharing and dynamic register-binding strategy targeting power reduction in data-intensive hardware designs. The experiments demonstrate that for a small overhead in the number of registers, it is possible to significantly reduce or eliminate spurious computations in a design. The proposed strategy handles this problem by performing register duplication or an inter-iteration variable assignment swap during the high-level synthesis process.

[86] provides an overview of clock-gating technique and how tools like power compiler use it. Many of these concepts are used in our work as well albeit at a higher level. [92] discusses how

observability don't care (ODC) conditions can be exploited for reducing the switch activity for redundant computations. Such an approach provides more advantage over the one provided by the commercial tools. Most of the RTL based approach shows that inability of designer to look into the control flow of the design at high-level leads to such opportunities. Also, many times, various different blocks of the design are created by different RTL designers because of which sometimes global clock-gating or enabling conditions are missed by the RTL clock-gating tools. [93] utilize input stability conditions to find more aggressive power saving opportunities than ODC based approach. Cong et al. [94] show the application of ODC and STC for HLS tool. In all the analyses, simulation based selection is not performed. Our approach can be considered as a co-operative solution because we provide a guidance on top of the selections made by such techniques.

Munch et al. discuss the opportunities to reduce power at the RTL using operand isolation based technique to reduce the dynamic power at the RTL [95]. Operand isolation is a technique, which helps in reducing the redundant activities around datapath unit and is considered as a complementary technique to clock-gating. Because clock-gating does not help in controlling the datapath activity, it just controls the clock toggles of registers.

[96] describes a system called CoDel, where they apply clock-gating at high-level to DSP benchmarks for power reduction. The reported power numbers show some clock-gating opportunities, which Power Compiler misses but their tool can find. They utilize state transition information and set of reads and writes for each state to obtain the clock-gating logic. Although, their approach helps in reducing power consumption, reported numbers show almost 20% area penalty and 15% timing penalty as compared to Power Compiler.

[97] also provides an approach to enable clock-gating in the HLS flow but their approach also lacks a simulation driven realistic power reduction feature. Their approach require RTL synthesis to insert the gating logic while we insert the necessary logic into the source code before generating the RTL. Singh et al. in [98] present algorithm for clock-gating and operand isolation to reduce the dynamic power consumption of the design at high level in Bluespec [8]. They propose a technique to automatically synthesize power optimized RTL using Bluespec compiler. Singh et al. present quite a few approaches for power reduction for concurrent action oriented synthesis tools such as Bluespec in [99, 100, 98, 101, 102, 103, 104, 105, 106, 107, 108]. They have presented approaches for dynamic power reduction using operand isolation and clock-gating. They discuss the complexity of power reduction using clock-gating and rescheduling of rules. They also present an approach in which rescheduling for rules is performed to reduce the peak power consumption of the design. Further, this may impact the functionality of the design for which they present a formal verification approach using a model checker SPIN. They check the equivalence between the design before and after performing the optimization.

Clock gating and operand isolation are two techniques to reduce the power consumption in state-of-the-art hardware designs. Both approaches basically follow a two-step procedure: first, they statically analyze a hardware circuit to determine irrelevant computations. Second, all parts which are responsible for these computations are replaced by others that consume less power in the average case, either by gating clocks or by isolating operands. Jens et al. [109] defines the theoretical

basis for adoption of these approaches in their entirety. They show how irrelevant computation can be eliminated using their approach. They present passiveness conditions for each signal x , which indicate that the value currently carried by x does not contribute to the final result of the system. After showing how their theory can be generally used in the context of clock gating and operand isolation a classification of many state-of-the-art approaches is performed and shown that most of the approaches in the literature are conservative approximations of their general setting.

[91] targets low-power design for FPGA circuits. It presents a simulated annealing engine that simultaneously performs scheduling, resource selection, functional unit binding, register binding and datapath generation in order to reduce power. [91] also proposes a MUX optimization algorithm based on weighted bipartite-matching to further reduce the power consumption of the design. Power management refers to techniques that involve shutting-down parts of a hardware design that are not being used for power savings. This can be done by disabling the loading of a subset of registers based on some logic [110, 111, 95, 112]. Operand Isolation (also known as *signal gating*) avoids unnecessary computations in a design by gating its signals in order to block the propagation of switching activity through the circuit. [111] discusses automation of operand isolation during ADL(Architecture Description Languages) based RTL generation of embedded processors. In [95], a model is described to estimate power savings that can be obtained by isolation of selected modules at RTL.

[112] defines power-island as a cluster of logic whose power can be controlled independent from the rest of the circuit, and hence can be completely powered down when all of the logic contained within it is idling. [112] proposes technique that eliminates spurious switching activity and the leakage power dissipation in a circuit through the use of power islands. The technique first schedules the design in order to identify the minimal number of resources needed under the given latency constraints. After scheduling, the functional unit binding is done based on the life cycle analysis of all the operations. The scheduling and binding steps are followed by a partitioning phase which uses a placement algorithm that performs partitioning such that the components with maximally overlapping lifetimes are clustered together, and assigned to the same power-island. After the partitioning phase, register-binding is performed such that both the total and the average active cycles of registers are minimized.

Recently power reduction techniques have also been used for security purpose. With increased outsourcing, confirming the genuineness of third party manufactured ICs has emerged as a major challenge. Researchers have effectively used various side-channel analysis techniques to fingerprint ICs viz. power, timing, EM-radiation. In [113], [114] authors have used circuit partitioning techniques to selectively exaggerate power consumption in targeted portions while reducing the overall power of the chip. In [115] a sustained vector scheme is employed to ensure that no transitions are initiated from the PIs so that the overall transitions occurring (that translates to overall dynamic power) in the circuit can be minimized. In [116] authors have used voltage inversion technique to change the functionality of the circuit so that malicious insertions are exposed off. This is augmented by a power profiling technique which is proven to be very effective in light of the changed functionality.

Some other low-power high-level synthesis works include [117, 118, 119]. [117] presents a high-level synthesis system for targeting reduction of power consumption in control-flow intensive as well as data-dominated circuits. The system uses an iterative improvement algorithm to take into account the interaction among the different synthesis tasks for power savings. [118] presents a high-level synthesis system for minimizing power consumption in application specific datapath intensive CMOS circuits using a variety of architectural and computational transformations. The proposed approach finds computational structures that result in the lowest power consumption for a specified throughput given a high-level algorithmic specification. [119] proposes a thread partitioning algorithm for low-power high-level synthesis systems. The algorithm divides parallel behaving circuit blocks (threads) of a design into subparts (sub-threads) such that gated-clocks can be applied to each part for power savings. As proposed in [120, 121, 122, 123, 124], high-level synthesis of hardware designs can also be used to target the peak power reduction of the generated designs.

Lakshminarayana et al. [125] present a methodology to quickly explore various hardware configurations and also to obtain relatively accurate design matrices for each configuration. They use C2R high level synthesis tool to quickly generate RTL description of the hardware. They show how HLS languages give a greater degree of freedom to do selective optimization by means of inserting the so-called ‘directives’ into the behavioral code (also known as restructuring). They present case studies to develop or modify behavioral IP descriptions and use standard FPGA boards to profile the IP in very short time. The difference in the measured and actual IP design matrices are not significant as one is more concerned with relative difference among various configurations. A variety of compute-intense benchmarks like AES is used to demonstrate how platform specific optimizations as well as higher level micro architectural optimizations can be done using a commercial HLS tool, Xilinx Spartan/Virtex boards and Xilinx EDK design suite. The results presented in this paper show how various architectures in hardware software codesign flow can be chosen keeping energy efficiency in mind. They also show how they reduce design cycle time to reach the optimal results.

2.3.1 Summary - Low-Power High-Level Synthesis Work

Dynamic power is one of the most important components of power consumption of a design, and thus its reduction is targeted during most power-aware high-level synthesis processes. Current approaches at RTL or higher level are designer’s knowledge dependent. Most of the approaches do not provide any support for power reduction from the behavioral specifications itself. In this thesis, we make an attempt in that direction. We propose approaches to enable clock-gating from the C description itself for various granularities of clock-gating such as fine grain at variable level and coarse grain at function or scope level. We also show how to extend this approach for sequential clock-gating. Finally, we present how to utilize power models to guide power reduction process at high-level. The advantage of such an approach is facilitation of power reduction features at the high-level. Also, the approaches for power estimation, which will be discussed later in detail, combined with power reduction approach can make the design flow completely at high-level. This

will help in providing power aware design methodology at high-level with faster turn around time.

Chapter 3

Background

We discuss the necessary background topics, which will help in understanding this dissertation. We briefly explain the components of average power consumption, because this dissertation addresses the problem of power estimation and reduction at the high-level. We provide an overview of FSM modeling using GEZEL, which helps in understanding the approach discussed in Chapter 4. We briefly provide an overview of the ESTEREL [18], PowerTheater [23], C2R [22]. These tools are used in this thesis and will help the reader in understanding the details in a few chapters. We briefly discuss various modeling styles used at the high-level under Transaction Level Modeling (TLM) section. We finally touch upon the basics of clock-gating and sequential clock-gating.

3.1 Average Power components

Average power dissipation in an hardware design consists of the following three basic components [1]-

1. *Dynamic Power* - It is related to switching activity occurring in the design. Dynamic power can be expressed as -

$$\text{Dynamic Power} = kCV^2fa$$

where, k is a constant.

C is the overall capacitance.

V is the supply voltage of the design.

f is the switching frequency of the component.

a is the switching activity.

Dynamic Power is the dominant source of power dissipation in hardware design, and is highly dependent on the application and architecture. A lot of opportunities to reduce dynamic power exist at high-level. Power reduction techniques discussed in this thesis concentrate on reducing the dynamic power consumption of the hardware design. Switching activity and capacitance are two most important factors to keep in mind at high level because small changes in hardware architecture may impact the capacitance and activity profile of the design.

2. *Leakage Power* - It is the static component of the power dissipation and occurs due to spurious currents in the non-conducting state of the transistors. Static power has two important factors, namely *Sub-threshold leakage* and *gate to source leakage*. As process technology node size is decreasing sub-threshold leakage is increasing. This factor exponentially increases with newer technology node. Some of the prominent techniques to reduce leakage power include power-gating and stacking of transistors. These techniques are applied after the netlist is finalized for the design. Gate to source leakage can be reduced by improving the gate insulation of a node of transistor.
3. *Short-circuit Power* - It occurs due to the stacked P and N devices in a CMOS logic gate that are in the ON state simultaneously, and can be minimized by reducing the signal transition times. It is hard to reduce this component of power through synthesis-time optimizations. After the netlist is finalized, a lot of care is given to place the design so that there is not much voltage drop for some transistors. Voltage drop on nodes may not only cause timing issues but also short circuit power issues. If short circuit power persists for longer time it may cause extremely high power consumption.

Various tools exist in industry that can perform power estimation accurately at the RTL or lower-level such as PowerTheater [23]. These tools require design information for models in Verilog/VHDL, simulation dump from the RTL or lower-level simulation (such as Value Change Dump (VCD) or FSDB) for activity analysis, and technology library etc. Such information helps in doing power estimation accurately and efficiently at the RTL and lower level.

3.2 PowerTheater

PowerTheater [23] is an RTL/gate-level power estimation tool, which provides good accuracy for RTL power estimation with respect to the corresponding gate-level and silicon implementation. We have used PowerTheater extensively for power estimation purposes. PowerTheater (PT) accepts design description in Verilog, VHDL or mixed verilog and vhdl. Other input required for average power analysis is value change dump in vcd or fsdb format dumped from RTL simulation. PT also requires power characterized libraries in .lib or .alf format for power analysis. Apart from doing average power estimation, PowerTheater also helps designers to perform the activity analysis for the testbenches, probabilistic power estimation, time-based power estimation, etc.

In simulation-based approach, PT first reads the design description provided in verilog/vhdl then it extracts activities from VCD and convert it into GAF (Global Activity File) and finally it maps the activities of the different signals obtained from RTL simulation of the design and does the power analysis. The most time consuming part in the simulation-based approach is activity extraction. If testbench is written for fairly large simulation time then size of value change dump will be very big and hence activity extraction will take longer time, which will further elongate the power analysis time.

Vectorless power estimation method is used when stimulus for the design is unavailable. In vectorless approach, PowerTheater reads the design description in Verilog/vhdl, then it assumes that activities of ports and internal signal will be provided by the user in VAF (Vectorless Activity Format). Once the activity of input-output ports and internal signal (if available) is provided, PT performs the probabilistic power estimation by propagating the activity of the missing signals from the activity provided for inputs and outputs. The accuracy of this method depends upon the accuracy of the activities of input, output ports and internal signals in vectorless activity file.

3.3 FSMD Modeling using GEZEL

The model of computation for GEZEL is hardware-oriented and is based on FSMD. GEZEL environment supports co-simulation with several instruction-set simulators as well as with SystemC and Java. In this dissertation, we utilize the modeling style of GEZEL for FSMDs and obtain state-wise toggle count of the design. Various features and advantages of the GEZEL are available at [25], here we list a few main advantages of the GEZEL:

- GEZEL uses cycle-true semantics with dedicated modeling for control structures and allows compact representation of the micro-architecture of the domain-specific processors.
- The simulation environment is scripted for fast edit-load-simulate cycles.
- GEZEL simulation environment provides a support to collect datapath activity per cycle during the simulation. This information can be further utilized for capturing total activity per state to create statistical model for average power computation.

3.3.1 An Example

In GEZEL, there exists a separate datapath and control corresponding to each FSMD. In the control path, we describe the control behavior of the design using finite state machine. The mapping of state machine to the description is self evident in the GCD example shown in Listing 3.1. The state machine performs certain computation functions invoking datapaths based on the input conditions in each state and moves to the next state. In the GCD design, state s0 of the FSM checks if ldr

signal is true and then invokes the signal flow graph “run” and moves on to state s1. The datapath “run” performs the core GCD computation.

Listing 3.1 shows the snippet of gcd example written in GEZEL. In fsm, we describe the controller behavior of the design using finite state machine. In every state, state machine performs some computation and goes to next state based on certain condition. For example, in the gcd design shown in Listing 3.1 the state s0 of the fsm checks if ldr signal is true, then perform operations in “run” signal flow graph (sfg), which does the core computation for gcd. GEZEL provides an environment where designer can express an FSMD and can do cycle-accurate simulation for the design as well. Detailed description of GEZEL is available at [25]. GEZEL supports advance datatypes such as integers etc. and helps in modeling FSMDs.

Listing 3.1: GCD code snippet in GEZEL

```

1 //Datapaths
2 sfg init { m = m.in;
3 n = n.in; }
4 sfg run { m = (m > n) ? m - n : m;
5 n = (n > m) ? n - m : n; }
6 }
7 //State machine capture for FSMD
8 fsm ctl_gcd(gcd) {
9 initial s0;
10 state s1;
11 //State machine information
12 @s0 if (ldr) then (run) -> s1;
13 else (init) -> s0;
14 @s1 if (doner) then (init) -> s0;
15 else (run) -> s1;
16 }

```

3.4 ESTEREL

ESTEREL is an imperative language for modeling synchronous reactive systems [20], especially suited for control-dominated systems. It has various constructs to express concurrency, communication and preemption, whereas data-handling follows the style of procedural languages such as C. Its semantics can be understood as a Finite State Mealy Machine (FSM), but it must ensure determinism, so that a program generates the same output sequence for the same input sequence. Internally, the ESTEREL compiler translates the control part of a program into a set of Boolean equations with Boolean registers.

3.4.1 Esterel Studio

Esterel Studio (ES) [18] is a development platform for designing reactive systems, which integrates a GUI for design capture, a verification engine for design verification and code generators to automatically generate target-specific executables. The GUI enables modeling through the graphical safe state machine specification called SSM or the ESTEREL textual specification [20]. ES

performs static as well as run-time consistency checks on this model for non-determinism, dead code, etc. The next most important ingredient of ES is the formal verifier (*esVerify*), which verifies the designer's intent. It allows both assertion-based verification as well as sequential equivalence checking. User-defined properties (expressed as assertions) and automatically extracted properties (out-bound, overflow, etc) are formally verified by generating the appropriate counter-traces (.esi files) to illustrate violations. An assertion is flagged as an **Assume** to notify *esVerify* that it is an assumption about the environment. It can be applied as a constraint during formal verification (FV) of the the other non-**Assume** assertions. Finally, ES performs multi-targeted code generation, which range from targets such as RTL (in VHDL/Verilog) to ESL (in C).

3.5 Multi-variate least squares linear regression model

Linear regression is a regression analysis in which a relationship between multiple independent variables and dependent variable is modeled by least square function. This equation is called as linear regression equation, it is a linear combination of one or more model parameters, which is called as regression coefficients.

3.5.1 Theoretical Model

Consider a sample of m observations done on n variables (X_1, X_2, \dots, X_n). If these n variables are assumed to satisfy a linear relation with the response variable Y then it can be represented as:

$$Y = \beta_0 + \beta_1 X_1 + \beta_2 X_2 + \dots + \beta_n X_n \quad (3.1)$$

For the regression model shown in Equation 3.1, there can be various ways to calculate the value of regression coefficients β such that the error between the predicted and measured value of response variable is the minimum. Lets' denote the i^{th} observation on these variables as $(X_{i1}, X_{i2}, \dots, X_{in})$, $i = 1, 2, \dots, m$. Lets say Y_i is value of the response variable Y for the i^{th} observation. *Least squares error* method can be used to reduce the difference between predicted and measured values, objective function can be represented as:

$$\min \left(\sum_{i=1}^m (Y - \beta_0 - \beta_1 X_{i1} - \dots - \beta_n X_{in})^2 \right) \quad (3.2)$$

3.6 Transaction-level Modeling

In this section, we briefly discuss the important terms used in the system-level domain specially related to transaction level modeling (TLM). There are a lot of references that are available on TLM, we briefly touch upon the terms and concepts explained in the paper by Cai et.al [2].

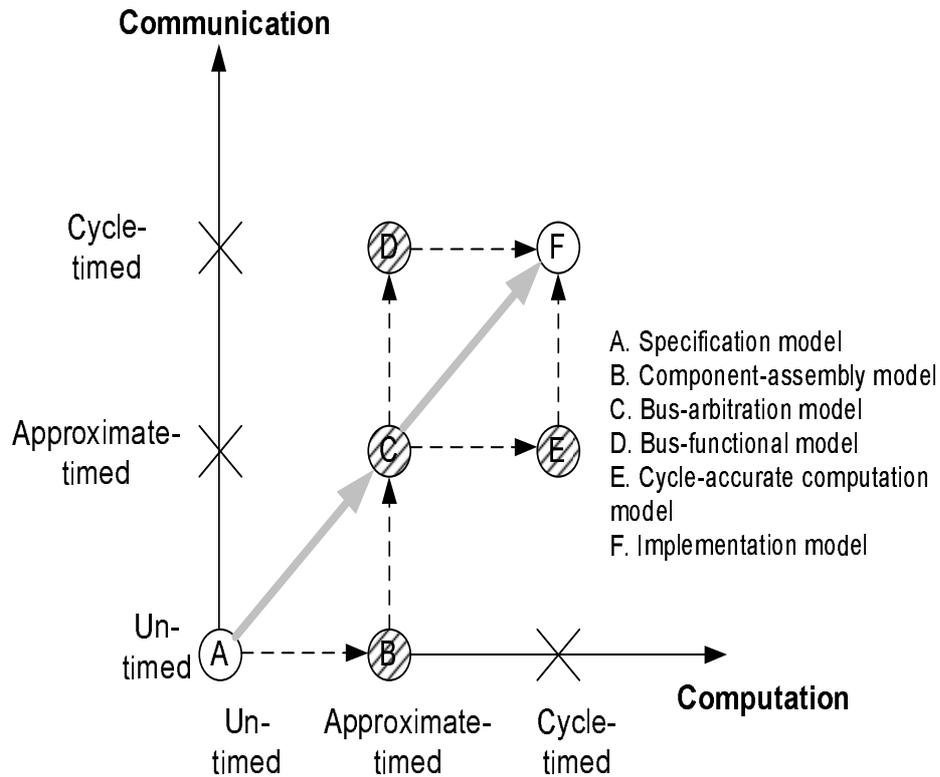


Figure 3.1: System modeling discussed by Cai. et. al. [2]

Transaction-level modeling (TLM) is used to increase the productivity of design flow. In such kind of modeling style models are presented at the higher level and can be used for quick design space exploration purpose. In TLM, details of communication models are separated from details of computation models. Figure 3.1 presents the system modeling approaches discussed in [2]. This paper briefly discusses the models commonly used in the design flow. Figure 3.1 provides a relationship between various TLM models through the relationship between communication (y-axis) and computation (x-axis). On each axis we can see three different time approximations namely un-timed, approximate-timed and cycle-timed. Un-timed model generally represents the functional model without any timing information. Cycle-timed accuracy can be considered as fine-grain accuracy for TLMs, while approximate-timed model can be any model having timing accuracy between untimed and cycle-timed models.

Various type of models used in design flow are captured in Figure 3.1. Specification model can be considered as the most abstract model, where computation and communication have no timing information and it is generally used for functional verification purposes. In component assembly models, computation is modeled in approximate time, while communication between processing elements (PEs) is not modeled with any timing information. Bus arbitration model represents approximate timing for the communication as well as the computation of the design. Generally the communication between PEs is modeled using channels without any cycle-accurate informa-

tion. A bus arbiter is inserted into the system to arbitrate any bus conflict. Bus functional model represents the cycle-accurate communication and approximate time computation. Constraints of communication are captured using timing information or clock cycles of bus clock. Cycle accurate computation model contains cycle-accurate computation and approximate time communication. This model can be generated from the bus arbitration model, it contains more detailed computation modeling at the cycle-accurate level. Finally implementation model generally captures PEs at RTL, communication channels are generally represented using wires. PE interface is pin accurate in the implementation model.

For SystemC various TLM models are discussed in [3]. TLM modeling shown in Fig. 3.2 is specific to SystemC, if we closely monitor modeling paradigm for SystemC. Its modeling accuracy also depends on timing accuracy and data granularity while in the earlier discussed paradigm, it was discussed more in terms of communication and computation. These paradigm discuss the need of various level of abstraction for modeling different systems and its importance at various level.

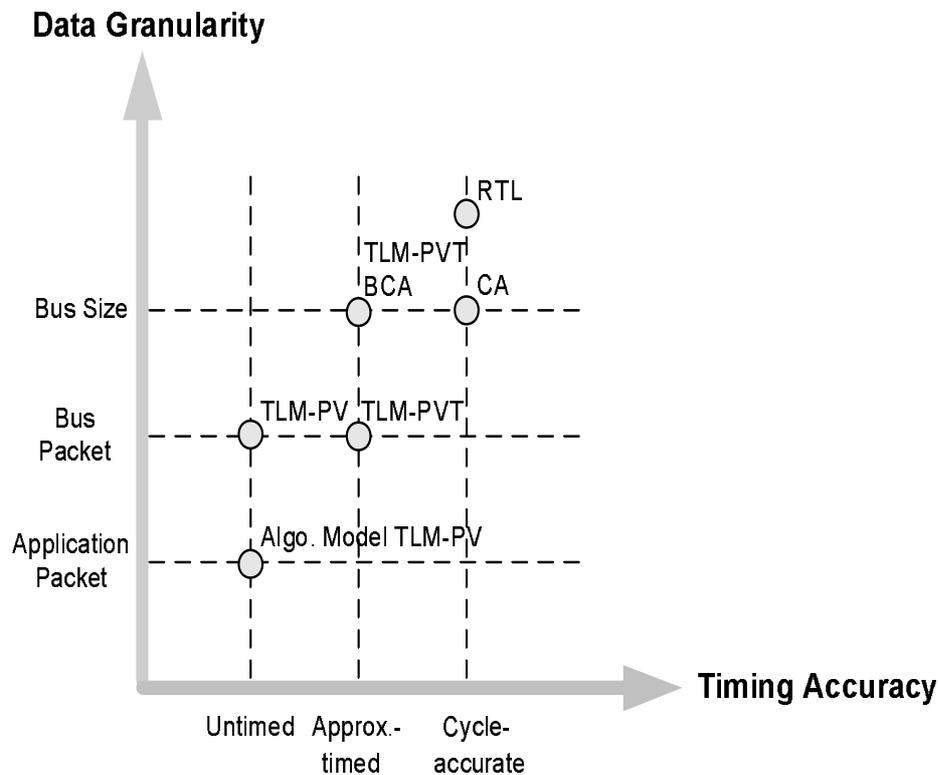


Figure 3.2: System modeling discussed by Frank Ghenassia [3]

To illustrate data granularity, one can take an example of video IP. If IP is modeled using frame based algorithm, then data transfer is done frame by frame at the coarse grain level, which can be considered as application packet level. If IP is modeled at the finer granularity then transfer can be represented as line or column based, or a transfer consisting of both line and column. While at the finest granularity it can be represented as pixel by pixel to transfer the video data. While

timing accuracy varies from untimed to cycle-accurate (cycle by cycle behavior), anything falling in between is approximate timed.

The purpose of untimed TLM is primarily creating functional software model and functional verification. Since this model is created primarily for software programmers, it is also named as transaction level model with *programmer's view* (PV). *Programmer's view plus timing* (PVT) models generally try to capture micro-architectural details containing essential timing annotations for behavioral and communication specifications. These models are generally used for checking the simulation accuracy for real time embedded software development and architecture analysis. Such timed TLMs are also known as PVT. Figure 3.2 captures the modeling accuracy of the untimed and timed TLM with respect to other modeling style such as bus cycle accurate (BCA), cycle accurate (CA) and register transfer level (RTL) models.

3.7 High Level Synthesis using C2R

C2R [22] is a high-level synthesis tool, which takes ANSI C specifications as input and generates synthesizable RTL as output. We have used this tool extensively in our studies. Here we discuss in detail, the design methodology based on C2R and salient features of this tool.

- **Specification:** As we know that for any design flow, specification is the starting point of the design flow. In the C2R flow, ANSI C specification is taken as input. Since C is one of the most widely adopted languages, it is relatively easy to obtain specifications for various important applications such as compression, image processing, security related IPs. Such specifications in C helps the designer to understand the application domain better and helps in fixing design requirements at the specification stage itself.
- **Restructured Specifications:** This is one of the most important stages of the design flow, which requires the highest effort from the designer. In this stage, the designer works on the macro-architecture of the design based on the specification and design-requirements and provides a restructured design. In this context, the restructuring phase includes inserting the necessary parallelism in the appropriate parts of the code. C2R provides user a way of compiling the restructured specification using gcc. This helps the designer to ensure the functional correctness of the design and understanding if various ways of introducing parallelism can affect the functionality of the design.

To enable sequencing and parallelism in the input C language, C2R requires restructuring to be performed by the designer using some directives. This restructured language is taken as an input by the C2R compiler. Restructured design from the macro-level can be understood as concurrent computation units, C2R provides the directives to include the parallelism. One can go through the code-flow (which is sequential C code) and introduce parallelism based on his needs. By default, the compiler selects a particular implementation based on control flow. To change the default flow, the designer has to provide some extra information. For

example, at the highest level one can extract independent processes using the `c2r_process`. Inside a thread, a designer can further exploit the parallelism using `c2r_fork` or `c2r_spawn` directives, in which instead of sequentially executing the code one can exploit the inherent parallelism in the code/design and use it at process-level. C2R also provides an opportunity to go a level deeper to enable parallelism, in which the user can ask the C2R to produce the micro-architecture utilizing lesser clock or more clock based on the requirement.

Such parallelism can be exploited using forwarding of clock, this technique utilizes the local data dependency, if specified then compiler will put the clocks in the generated RTL Verilog. For all the stages, compiler has a default behavior in which it tries to find the maximum parallelism, while other variants of macro-architecture can be obtained by including directives such as `c2r_fork`, `c2r_process`, etc.

- C2R: There are three main stages/passes in the compiler. We will briefly touch these passes and provide an overview on how compiler understands the input provided and how it processes them.
 - Frontend Pass: Compiler first does the lexical processing for parsing the input specifications. Front end pass can be understood as the pass in which the compiler mainly gathers the information from the restructured C and creates the abstract syntax tree (AST) for further processing. In this stage, the compiler checks the correctness of the input language and passes through the expression-typing stage to check the lexical and syntactical correctness. Also, the compiler starts building the symbol-table to record all variables, functions, declarations, etc. This symbol table is used by the later passes to extract the information related to types, size, scope, attributes, etc. The compiler also performs syntax checking, type checking etc.
 - Middle Pass: This pass can be considered as a pass where AST is massaged based on the macro-architecture definition provided by the designer. In this pass, the compiler does processing for loop-unrolling, linearizing (i.e. flattening constructs, such as if-then-else, while, switch, for, etc.). Every stage introduces some additional details on the AST or elimination such as dead code elimination, etc. This middle pass can also be considered as a stage where compiler goes through some optimization passes as well such as clock-insertion pass. Finally before the backend code generation, compiler extracts the control/data flow graph (CDFG) from the AST. Finally an appropriate finite state machine is extracted from the CDFG of the compiler for backend processing to generate the synthesizable Verilog.
 - Backend Pass: Backend pass mainly works on the CDFG provided by the middle pass and creates the Verilog/ cycle-accurate SystemC. In this stage, the compiler finds out appropriate constructs of Verilog to be put on various nodes, which are already scheduled and allocated.
- Backend Flow: In the context of high-level synthesis, the flow which utilizes synthesizable RTL can be considered as backend flow. The generated verilog can be used in any FPGA/ASIC flow using tools provided by xilinx, synopsys, etc. There is a difference in the backend

pass and backend flow. Backend pass is the pass to perform post processing on the CDFG, while backend flow is applied on the HLS generated RTL. Backend pass is internal to the high-level synthesis tool.

3.8 Power Reduction using Clock-gating

Dynamic power consumption remains to be the biggest contributor to the total power consumption of a hardware design. Register power consumption is one of the highest contributor to the dynamic power, making register power as one of the biggest contributor. One of the most prominent technique to reduce power consumption of a hardware register is clock-gating. In this section, we discuss the basics of *clock-gating* and its variant known as *sequential clock-gating*.

3.8.1 Clock-gating

A typical register in hardware is shown in Fig. 3.3. For a register shown in Fig. 3.3, the input to the register is the output of a multiplexer behind it. This multiplexer output is dependent on its EN input, which is the control signal for the multiplexer. If the EN is true then register takes IN as input else it retains the previous value. If the EN signal does not change frequently, (i.e new data is not coming frequently to the register), the register still consumes dynamic power every clock cycle. The clock-gating gates the clock based on the EN signal, which stops the clock driving the register when it is not needed. There are many ways to gate a clock, and one common way is the latch based clock-gating, pictured in Fig. 3.4. More details on clock-gating can be obtained at [86].

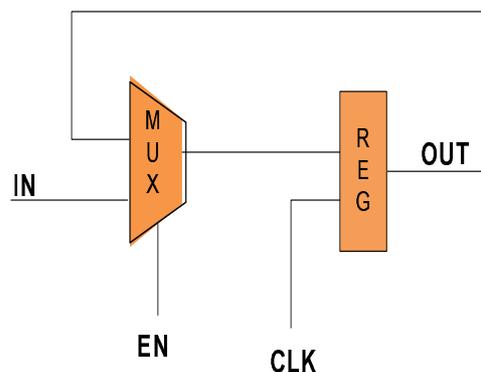


Figure 3.3: Typical register representation in hardware

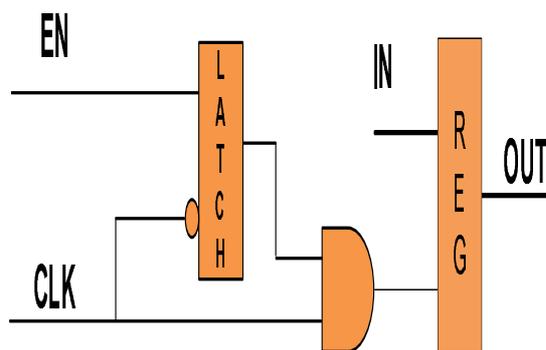


Figure 3.4: Clock-gated register representation in hardware

3.8.2 Sequential Clock-gating

Sequential clock-gating is a technique which is applied across the clock boundaries. Sequential clock-gating is applied to gain further power savings [4], [126], [127]. At the RTL, sequential clock-gating opportunities may exist in many different ways, some of which are discussed below with examples:

De-assert enables in the earlier cycles if forward stage is clock-gated

In the example shown in Fig. 3.5, three registers `d_out`, R1 and R2 are connected in a pipeline. The register (`d_out`) has feedback path from a multiplexer. This multiplexer's controlling input is `en`. If `en` is true then `d_out` accepts value from the preceding combinational logic, else `d_out` retains its old value. Other registers do not have any multiplexer before the input, hence no enabling conditions to control the register value.

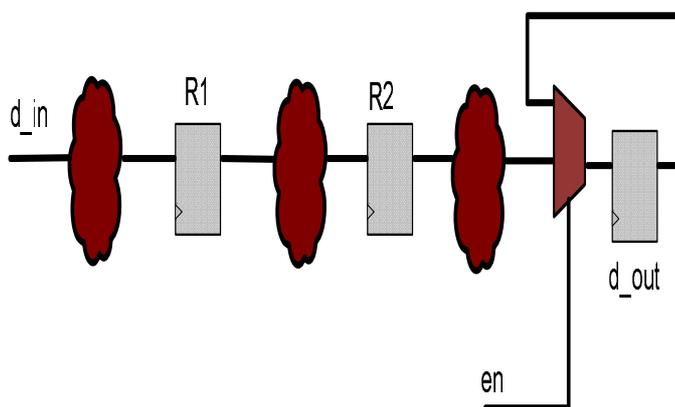


Figure 3.5: An example of de-asserting the datapath registers to find out the sequential clock-gating opportunities [4]

Suppose we find that the register R2 and its preceding combinational logic do not feed into any other part of the design except the logic preceding d_out. We also find that register R1's fanout cone does not include any other logic than shown in the figure. Now, if en is false, at cycle t, any activity on R1 in cycle t-2 (2 cycles earlier) will not affect the d_out. Similarly any activity on R2 at cycle t-1 will not be relevant either. Discovery of these kinds of relationships between the registers can be exploited for clock-gating registers R1, R2 to save clock-power wasted in the redundant computations. Fig. 3.6 shows the power aware circuitry after utilizing the relationship for these enabling conditions. Since an enable signal for d_out already exists, RTL synthesis tool can recognize the clock-gating opportunity for d_out. On the other hand, clock-gating of R1, R2 requires the above analysis because sequential gating opportunities for these registers are not visible within a single clock boundary.

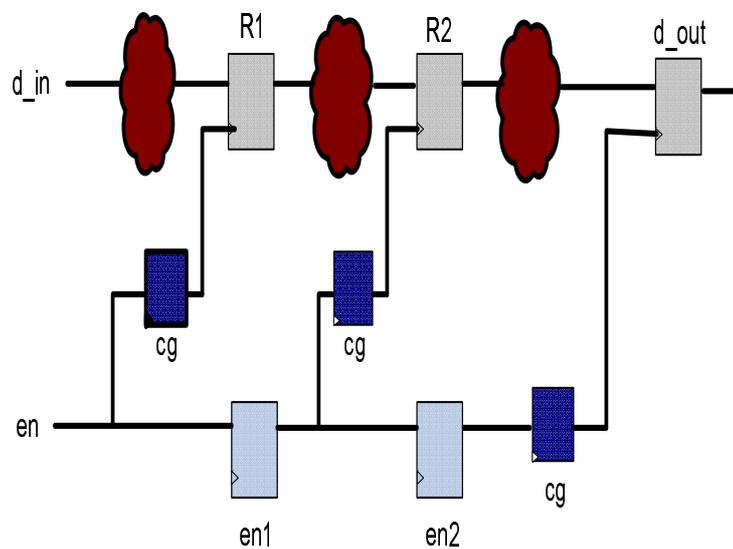


Figure 3.6: After applying the power reduction technique on the example circuit [4]

De-assert the forward stage of enable in later cycles if the current stage is clock-gated

Fig. 3.7 presents another example where the power reduction technique discussed above can be applied. In the example circuit shown here, datapath starts from the inputs d1_in, d2_in and ends at the output d_out. The two inputs are enabled by en1 and en2 respectively, while there is no enabling condition for d_out. If there is no change in value of R1/R3 at an arbitrary clock cycle t, then there will be no change in the value of R2/R4 at clock cycle t+1, and no change in the value of d_out at cycle t+2. Of course, this has to mean that that the fanout cones for each of these registers are limited to what is shown in this figure. For this scenario en1, en2 can be propagated to control the clock power of registers R2, R4 and d out. Fig. 3.8 shows the final design after applying sequential clock-gating optimization.

Now consider another example in the Snippet 1. If register r changes in cycle t , then register

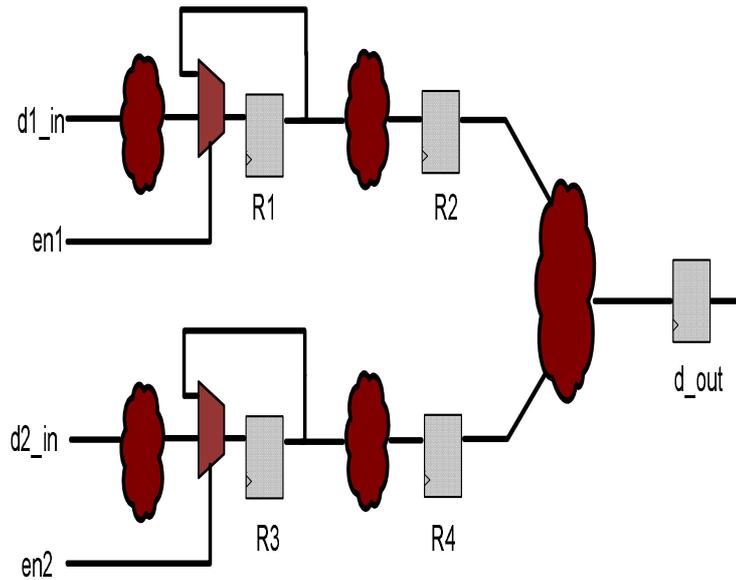


Figure 3.7: An example of de-asserting the datapath registers to find the sequential clock gating opportunities [4]

q changes in cycle $t + 1$. Moreover, if r does not change in cycle t , then q does not change in $t + 1$. This information cannot be inferred from structural information as were the cases in the previous two examples. However, a model checker can easily infer this information, and hence we can eliminate the computation of u , z (provided the only usage of u , z were creating the enabling condition for q 's updation), and clock gate q a cycle later if r is gated in the current cycle. Since this is a simple enough example, one could see it as obvious. But for more complex scenarios, only a model checker or state space exploration based method can infer these kinds of information. This also gives us a motivation to further research sequential clock-gating technique and adopt it at higher abstraction level.

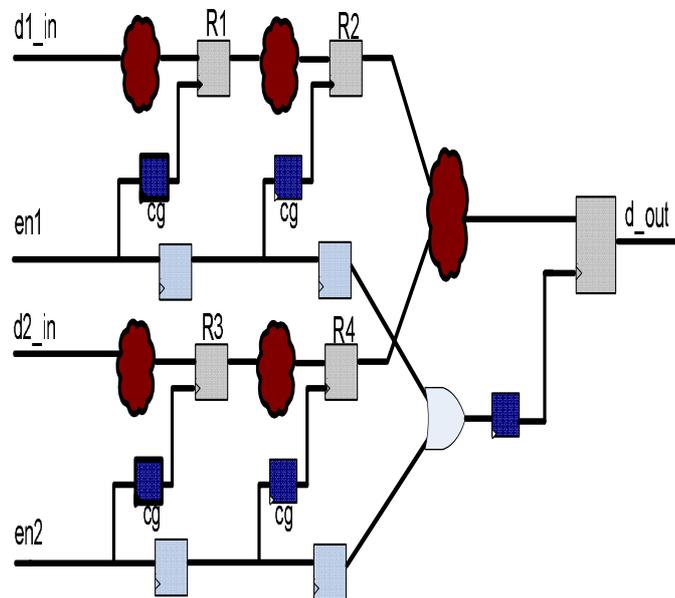


Figure 3.8: An example of de-asserting the datapath registers to find the sequential clock gating opportunities [4]

Snippet 1 Dynamic Opportunity of sequential clock-gating for registers r, q

```

Input int in1,in2;
wire int x, y;
register int u,z,r,q;
....
x = in1 + 1;
y = in2 - 2;
next(u) = in1;
next(z) = in2 -3
....
if (x-y > 10 ) then next(r) = expression;
....
if (u-z > 10) then next(q) = expression2;

```

Chapter 4

Statistical Regression Based Power Models

Accurate power estimation is one of the most important ingredients for any successful design methodology in embedded system or system on chip (SoC) design. Ability to estimate power accurately and early in the design flow helps in controlling form factor, battery life and size, etc. Current power estimation techniques are well established at the RTL and lower levels of abstraction, while the techniques above RTL suffer from accuracy problem. One of the main reasons for such inaccuracy of power estimation at higher levels is the lack of technology specific information and implementation details of a design. Power estimation at the RTL or above, requires design information in the form of Verilog/VHDL/SystemC/C model, technology libraries, simulation dump, etc. However, the more detailed the model is, the slower the simulation and hence the slower the estimation. In order to make this process faster it is desirable to use power estimation techniques that do not require all this information directly exposed in the model. In the past various techniques have been proposed to achieve this goal. These techniques involve capturing the activity by simulating a high-level model and feeding it to the lower level power estimation techniques [28, 30] using the library of blocks containing power related information [44], and using efficient power models/macro-models to estimate the power consumption [45, 46], etc.

The focus of this chapter is on creating an efficient power model to estimate the power consumption of a particular design/design-block without burdening the simulation model with all the design details and technology related information. The goal is to maintain a balance between the estimation speed and the accuracy achieved. In this work, we target a cycle-accurate FSMMD based modeling style and figure out what information must be collected from their simulation to obtain sufficiently accurate power estimation without having to simulate the RTL of the design in a full-chip simulation. Therefore we experimented and verified the model's accuracy on 90 as well as 180nm technology libraries to show the applicability of our model. The cycle-accurate FSMMD level of abstraction is less detailed than RTL. This FSMMD model can be co-simulated with other FSMMD models, and even C-models.

This provides the benefit of power aware exploration at the highest possible level of abstraction avoiding an RT level full-chip simulation. Even though the power consumption of an implementa-

tion depends upon the technology libraries and other implementation details, signal toggle plays a statistically significant role in power consumption. So, if the signal toggles at the various states of the design can be statistically correlated to the power consumption over a sizable number of simulation examples, one can learn the dependence of various statistics available at the FSMMD level on the actual power consumption of the implementation. Since, learning is aided by a detailed model of the design with specific technology library information, such information is folded into the regression model in the methodology presented in this Chapter. We will refer this methodology as Statistical regression based power models for Co-processors Power Estimation (SCoPE) hereafter in this Chapter.

SCoPE methodology utilizes the activity/toggle information of each state and corresponding datapath explicitly available in the FSMMD modeling style. Note that in the RTL model, the distinction between state and datapath activities for specific states is not explicit. As a result, state specific toggle counting is harder. The FSMMD abstraction comes to the rescue here. Learning based on the activity of the design at its various states is done by inferring a regression equation relating such statistics of the abstract FSMMD model and the power consumption of the implementation model. At the end of such a learning phase, the power model is trained/learned, and it can be used for estimating power consumption based on state specific activity information extracted from the co-simulation of the FSMMD and the rest of the chip (other parts of the chip could also have FSMMD substitutions for various components). In the *SCoPE* methodology power numbers can be obtained just after co-simulating the co-processor models with the simulation model of ARM processor.

There are quite a few advantages of *SCoPE* methodology: a-) it utilizes the explicit partitioning of design in states and datapath provided by the FSMMD modeling style (which helps in extracting relevant run-time information), b-) once the model is ready it requires only activity/toggle information per state (which can be obtained from the simulation of the FSMMD) to get an estimate on power consumption, c-) these models can directly be attached with the simulation model of the processors. Main features of *SCoPE* methodology can be summarized as follows:

- A two phase power estimation technique for components in an SoC design using multivariate regression utilizing a cycle-accurate FSMMD model of the design.
- A compositional methodology for power estimation of SoCs utilizing the statistical power models of its components while simulating at a higher abstraction level.
- Experimental demonstration of the accuracy of this technique for some co-processor examples.
- Demonstration of a flow based on this methodology which incorporates GEZEL [25] for FSMMD modeling and activity capture, PowerTheater [23] for learning, and JMP [128] for regression modeling.

4.1 Regression based Power Model for FSMs

The core idea is based on the fact that activities associated with a particular state of the FSM model corresponds to the specific part of the implementation model as shown in Fig. 4.1. Note that methodology discussed in this chapter is applicable to the language/design-flow capable of representing cycle-accurate FSM models such as Esterel [18], SystemC [16], etc.

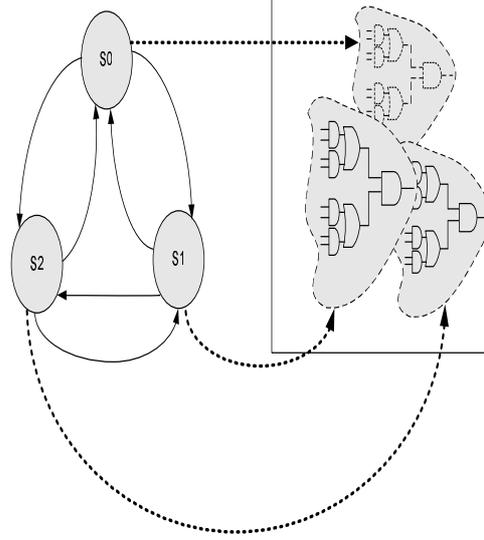


Figure 4.1: Conceptual mapping of activity to the real design.

Let's consider an FSM \mathbf{M} with state set \mathbf{S} for a test run for T clock cycles, α_s is the total number of times design visits state \mathbf{s} during the test. Let τ_i^s be number of signal toggles during the i^{th} time the design visits state \mathbf{s} , then total number of signal toggles in state \mathbf{s} can be represented as:

$$\Gamma_s = \sum_{i=1}^{\alpha_s} \tau_i^s \quad (4.1)$$

While the design is in state \mathbf{s} during the test, energy spent due to toggles in state \mathbf{s} , $E_s = k\Gamma_s C_s V^2$, where k is a proportionality constant, C_s is the capacitance, and V is the operating voltage. The total dynamic energy spent in design is assumed to be the sum of energy spent in each state and can be represented as:

$$E_{dyn} = \sum_{s \in \mathbf{S}} k\Gamma_s C_s V^2$$

Now, if clock frequency of design is f then the average dynamic power can be represented as:

$$\begin{aligned}
P_{dyn} &= E_{dyn} / (T/f) \\
&= \sum_{s \in S} k \Gamma_s C_s V^2 f / T \\
&= \sum_{s \in S} k (\Gamma_s / T) C_s V^2 f \\
&= \sum_{s \in S} K_s (\Gamma_s / T), \text{ where } K_s = k C_s V^2 f
\end{aligned}$$

Total power of a design is sum of dynamic and leakage power, then the equation for average power can be represented as:

$$P_{total} = P_{leakage} + \sum_{s \in S} K_s (\Gamma_s / T) \quad (4.2)$$

In the Equation 4.2, K_s is a constant for state s under the assumption that the state-wise toggles at FSM level are proportional to the toggles occurring on the implementation model. K_s can be obtained by taking the ratio of the dynamic power numbers obtained from RTL estimation to toggle count information obtained from FSM simulation and number of clock cycles for which the test was run. Our Experiments with 90 and 180nm technology libraries show that leakage dependence is captured in this model and power numbers obtained using this model are close to RTL power estimation numbers.

Relationship between toggles and total power, maps well to the regression model discussed in the Section 3.5. Variables of the regression model are obtained by taking the ratio of state-wise toggles and simulation duration in number of clock cycles. Regression coefficients can be obtained by applying least squares error objective on these variables and power number obtained using RTL power estimation on the implementation model. Using this model, we can reduce a lot of complexity in creating power models/macro-models because most of the time these models require a lot more information than toggle/activity of inputs and outputs etc.

4.2 Steps for Power Modeling

There are two models involved in this power modeling methodology: the FSM model and the power model. The FSM model is a cycle accurate abstraction of the implementation model. This model is the input for our power modeling methodology. The power model captures a relationship between the toggle statistics of an FSM model and power estimation done on implementation model. This regression model is the output of our power modeling methodology.

4.2.1 Learning Phase

Following steps are needed to perform learning phase of the *SCoPE* methodology (also shown in Figure 4.2):

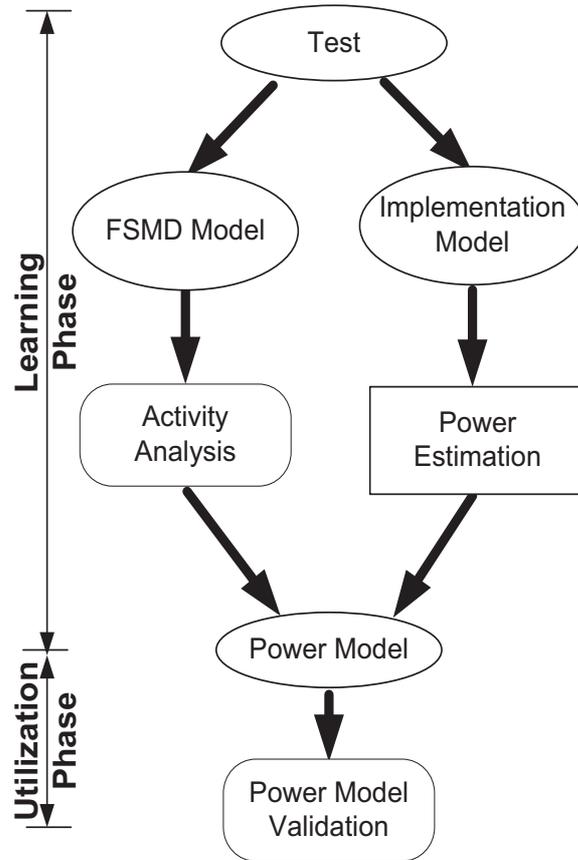


Figure 4.2: *SCoPE* Methodology for creating and utilizing power models of FSMs

- Test setup:* *SCoPE* methodology utilizes statistical modeling tool JMP [128]. Based on the experience of expert users, it requires at least 30 experiments for creating accurate statistical model using JMP. In this study we found that regression coefficients for power model converged to a stable value for 80 tests. Random tests are utilized in our experiments. As we know that power consumption is very much dependent on the input vectors/simulation, the better stimulus provided for the analysis the better power model (more accurate) we obtain. There exist quite a few approaches in the research where techniques to find the appropriate input vectors based on the hamming distance, data correlation, etc. are presented [129]. However the main focus of *SCoPE* methodology is on creation and utilization of the power model. The stimulus generation approach presented in [129] can also be utilized in our framework for training power model during test setup stage of the methodology.
- FSM Model:* We use FSM models written in GEZEL. An FSM is a cycle-accurate model of a controller with a datapath [130]. FSM modeling in GEZEL clearly distinguishes between control and datapath operations. Modeling style for datapath and control is very different but that is captured well in GEZEL using state transitions for controller and

expressions on signals or registers for datapath operations. FSM model should capture the information on every state and corresponding datapath.

- *Activity Analysis:* In this step, we mainly concentrate on obtaining the toggle count associated with every state. GEZEL simulator provides a facility to get a toggle count for every datapath unit for each cycle. For getting activity count associated with every state, we instrumented the GEZEL code and processed its simulation output using a c-shell script; this collects the sum of toggles for every state for the whole simulation duration. This activity information is used for training and validation of the power-model. GEZEL environment is capable of doing activity estimation associated with each datapath unit. In the script, we made sure that for every cycle each datapath unit can be associated to its corresponding state. This information is further processed to get the count of total toggle for every state during the simulation.
- *Implementation Model:* We use an RTL model written in Verilog as an implementation model for our analysis. One can also use the gate-level or even more detailed model in SCoPE methodology. We performed the analysis using RTL power estimation numbers. We assume that the finite state machine for the RTL model is same as the one used in FSM model. RTL contains more detailed information on the datapath and implementation of the hardware design. Hence RTL model is assumed as a reference point for our analysis. Power numbers are measured from the RTL power estimator PowerTheater [23], keeping in view that such an estimator provides fairly quick and accurate results for RTL circuit as compare to its silicon implementation.
- *Power Estimation for the Implementation Model:* Power estimation for the implementation model requires design information (Verilog description), activity information (we use Value Change Dump (VCD) generated by the RTL simulator) and technology library (we use 90 and 180 nm power characterized libraries). Power estimator provides the average power numbers corresponding to VCD of every training vector. As we go in lower level the power estimation will be very much time consuming but accuracy of power numbers obtained will increase and this will further help in improving the accuracy of regression based power model.
- *Power Model:* As discussed in section 4.1, we use regressions for creating the power model, JMP [128] is used in our framework for creating regression model. We utilized activity information from activity analysis step along with the measured power information from the power estimation step for training the model. We used model fitting feature of JMP using multiple regression model and objective as least squares error to get accurate value of the regression coefficients of the power model (as discussed in Section 4.1).

4.2.2 Utilization Phase

Validation of regression model is very important and we perform the validation of power-model during utilization phase.

- *Power Model Validation:* once the regression model is ready, we validate the power model using random tests. In this stage, we check how closely the output produced by the predicted model matches to the power numbers obtained from the power estimation stage discussed earlier. Most of the reported work use root mean square (rms) error (such as [44, 46]) which can only be helpful for average power over various tests. In our analysis, we are calculating error for each test which gives us insight on worst case scenario (rms error is less than the worst case error).
- *Usage:* Total toggles for every state is given as an input to the power model. Power model provides a quick estimate on power consumption using regression equation and does not require power estimation at the implementation level.

4.3 Results and Conclusions

4.3.1 Tool Flow

We have used PowerTheater [23] for power estimation, VCS [131] for RTL simulation, GEZEL [25] for FSM modeling and JMP [128] for creating regression model. Here we discuss in detail about the tools and wrapper we built for the *SCoPE* methodology. Fig. 4.3 shows the tool flow used in the *SCoPE* methodology. Grey divider line shows different phases. GEZEL simulation environment, c-shell script is required in both phases. From the GEZEL simulation environment toggles for every datapath units can be obtained. The C script wrapper collects toggles per state for the whole simulation duration and associate the appropriate toggles to each datapath computation unit. In the learning phase we provide the Verilog RTL model, value change dump and technology library information to PowerTheater for power estimation. These power numbers along with the state-wise toggle information are then passed to JMP for regression modeling. Once the model is ready the power estimation is performed in the utilization phase using the regression model (as discussed in Section 4.1, 4.2).

4.3.2 Experimental Results

We applied *SCoPE* methodology to various designs shown in Table 4.1, 4.2. In our study, designs have several states varying from as low as 4 different control states in the EUCLID design to 10 different control states (one for each of the 10 rounds of encryption) in the AES design. Complexity

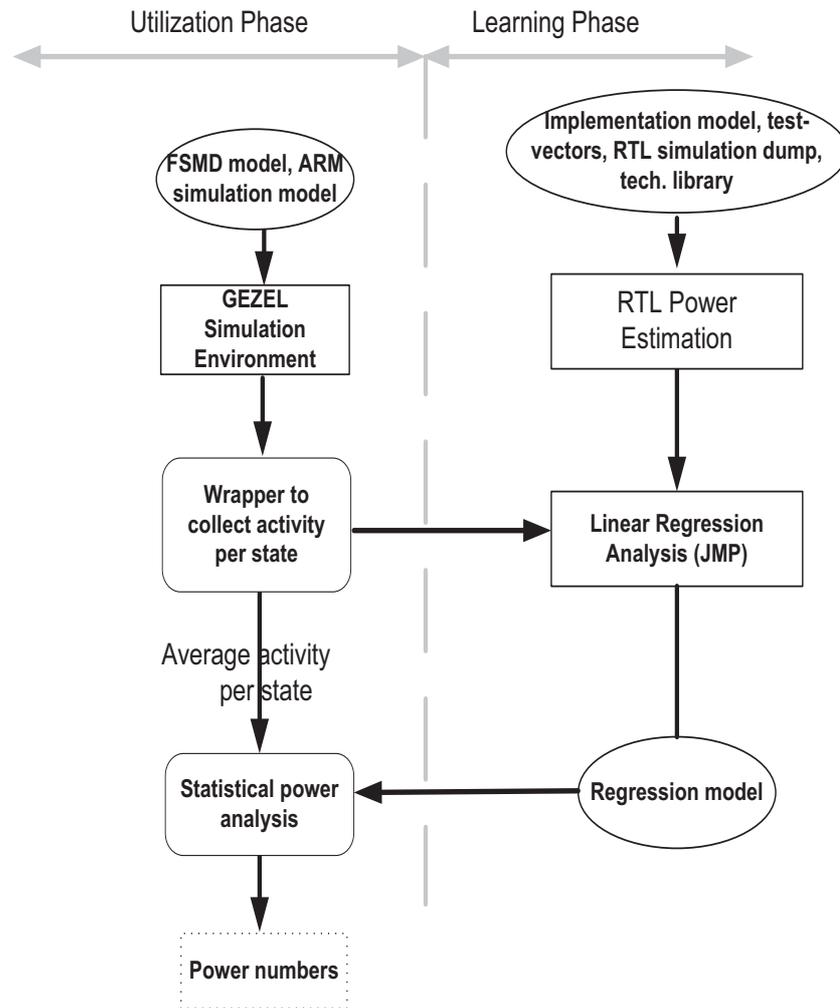


Figure 4.3: Various tools used in different phases of the *SCoPE* methodology

of the designs in terms of their size, ranges from 1099 NAND gates (for EUCLID) to 47166 equivalent NAND gates (for AES). Table 4.1, 4.2 capture power numbers, worst case and RMS error for various designs. Fig. 4.4, 4.5 provide a comparison of power numbers for different simulations of our approach with the PowerTheater for 90 and 180 nm respectively. Experimental results show that our model shows a good accuracy for different technology nodes.

Design	Pred. Power (mW)	RTL Power (mW)	Worst Error (%)	RMS Error (%)
AES	145	138.99	4.14	2.26
EUCLID	0.497	0.500	0.66	0.387
UART transmit	0.485	0.525	8.28	4.15
XTEA cipher	0.530	0.538	1.58	0.607
XTEA decipher	0.541	0.551	1.99	0.78

Table 4.1: Comparison of measured power (90 nm) and predicted power for various designs

Design	Pred. Power (mW)	RTL Power (mW)	Worst Error (%)	RMS Error (%)
AES	455.3	468	2.79	1.49
EUCLID	1.86	1.98	6.09	3.81
UART transmit	1.59	1.68	5.67	2.95
XTEA cipher	1.765	1.75	.86	0.25
XTEA decipher	1.79	1.827	2.08	0.79

Table 4.2: Comparison of measured power (180 nm) and predicted power for various designs

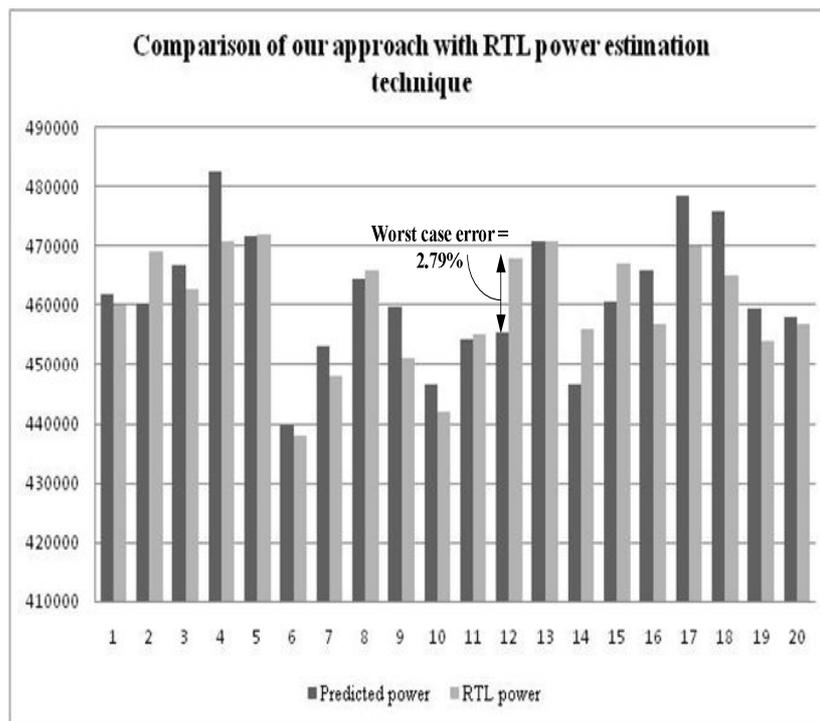


Figure 4.5: Comparison of measured power and predicted power for AES at 180nm

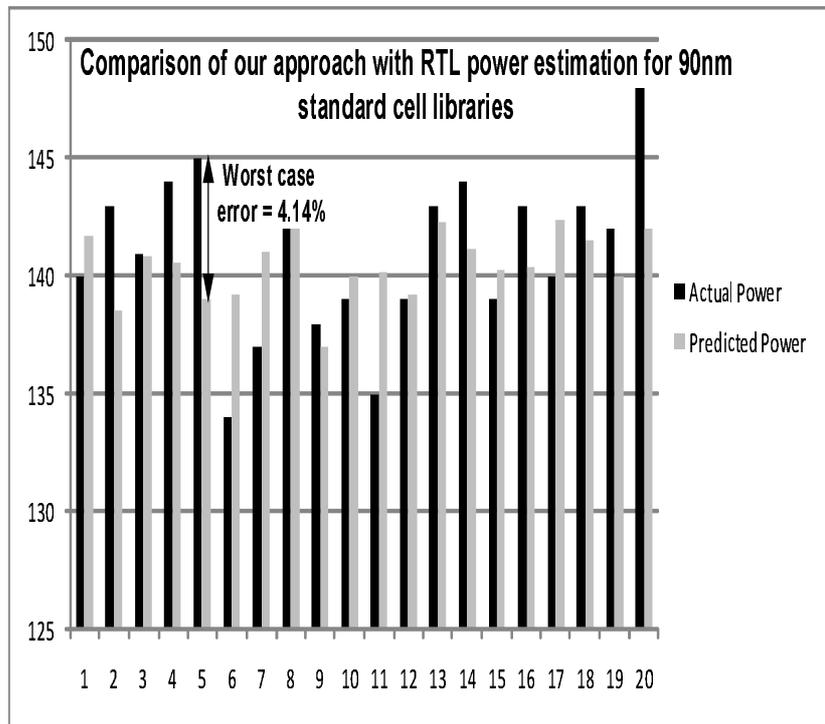


Figure 4.4: Comparison of measured power and predicted power for AES at 90nm

4.3.3 Discussion

Comparison with state of the art RTL power estimation techniques

Power estimation utilizing power models presented in this Chapter provides the following advantages over the simulation based RTL power estimation methodology:

- Once the power model is ready, it can be utilized for accurate power estimation requiring FSM simulation only.
- Power model presented in this chapter is a simple regression model which does not require going through the overheads involved in RTL power estimation such as processing RTL simulation-dump, technology libraries and design information, etc.

Factors affecting accuracy

Accuracy of such power models is dependent on various factors. These factors include test vectors, training of regression models, etc. Test vectors are very important because if the model is created using the test vectors that are being used for most common usage scenario then the accuracy of the

power model can be further improved. In our analysis we have used random vectors to train and validate the power model. Earlier related works [46] report RMS error which do not provide per sample average power accuracy and is not a good measure for average power analysis per sample/test vector. As we know that RMS error is less than the worst case error, so for the cases where one needs RMS error our framework will provide even better results. Training of regression model can also be further improved by doing the training on more test vectors. Finally the regression model we have used is multi-variate least square error based regression model, which is more suitable in reducing the average deviation of power consumption. Such kind of models may not be suitable for the designs with very high power variation for different states.

Secondly in our power model leakage dependence is also captured with the same regression model. We utilized the power numbers from the implementation model to train the power model (in our case RTL model of the hardware design). The power numbers obtained from the implementation model include the numbers coming from state dependent leakage as well. In the proposed power model, state dependent leakage numbers are not taken into account, which has further affected the accuracy.

Tools, Environment and Speed up

We utilize GEZEL, JMP and PowerTheater tools for experimentation. We anticipate such a methodology can be used with other framework/tools capable of representing FSM. Simulation environment may require some additional work to obtain the toggle count associated with each state and corresponding datapath of the FSM. As it is visible in our approach, power estimates can be obtained at the end of GEZEL simulation. In this approach, we have attempted to increase the abstraction level to cycle accurate FSM and still maintained a sufficient accuracy. We can remove a part of design cycle (that is the need of lower level power estimation), which leads to reduction of possible delays in the design cycle. To give an idea, we collected the time spent for RTL simulation and power estimation of AES design; which is found to be around 11 minutes of CPU time (7 minutes for power estimation and 4 minutes for RTL simulation for a test vector of almost 40,000 clock cycles). While in our case it is only around 5 minutes of CPU time. We believe that if such an equation based model can be inserted in the C source code of system-level model, it can show very good speedup. We have inserted a power model in the SystemC model in the different context [33] and found that the power models can provide a very good speedup (upto 2-3 order of magnitude). However such an analysis is not performed here.

Scalability of our framework/methodology

In this chapter, we presented our approach on co-processors used for acceleration of certain data intensive task. In our experiments we found that design with 10 or more control states at FSM level takes almost same or lesser execution time as compared to RTL. Learning phase for AES design with 10 states and 80 samples (test vectors) took less than 5 minutes of CPU execution

time. Most of the time in learning phase was spent in GEZEL simulation, once the state specific toggles were collected then regression model was created within a few seconds. Hence, we think our approach can be scaled to the co-processor of reasonable size. In our case, we utilized existing framework to extract state specific datapath activity information from the GEZEL simulation, while such feature might not be available with other frameworks. Other frameworks based on C/SystemC will require some instrumentation of design source code or environment.

Impact of RTL or lower level power reduction techniques

Power model presented in *SCoPE* is trained using RTL power estimation method. RTL power estimation tools may not provide the impact of power reduction techniques such as clock-gating, power-gating, etc. Now a days power estimation tools such as PowerTheater [23] can perform optimizations such as clock-gating, operand isolation at RTL abstraction of the design. In such a scenario it is recommended to train the power model after applying power reduction. To verify the claim we trained the power model for xtea cipher using power estimations performed after applying clock-gating on the design. We found that for the same test vectors worst case error was 1.55% and rms error was found to be .56%.

Similarly, multiple voltage domains inside a co-processor are not considered in the *SCoPE* methodology. However this methodology would still be useful in case a user wants to use multiple voltage domains at the granularity of co-processors. This type of power model would require more learning if a lot of power reduction is obtained by techniques which are available after gate-level design stage. Such optimizations include leakage power reduction techniques using stacking to reduce the power consumption.

In this Chapter, we presented a means to solution using GEZEL, JMP and PowerTheater. We assume such a methodology can also be used with other framework/tools capable of representing FSM/D. Such a methodology is capable of saving various iterations, which design/EDA engineers might have to go through otherwise to obtain the power numbers in their design flow.

Chapter 5

High Level Simulation Directed RTL Power Estimation

5.1 Introduction

In this chapter, we present a *high-level power estimation* methodology, which is based on a high-level synthesis framework and supports sufficiently accurate power estimation of hardware designs at the high-level. For early and accurate power estimation, the proposed methodology utilizes register transfer level (RTL) probabilistic power estimation technique controlled by the high-level simulation. Furthermore, our methodology does not require a designer to move to the traditional RTL power estimation methodology, thus facilitating easy and early power analysis and aiding the cause of adoption of high-level design practices in ASIC design flow. This chapter provides detailed description of our methodology including tools used, algorithm for extracting activity from high-level value change dump and finally mapping this information for RTL power estimation.

Current high-level power estimation methodologies are based on -

1. Relative power estimation [34]: *Relative power estimation* at high-level helps the designer to perform early relative trade-offs for power and performance but does not provide accurate power numbers, rendering such a methodology to be unsuitable for handling design problems where accuracy is a concern.
2. Power-model reuse [49], [50]: In a *power-model reuse* based methodology, power numbers are borrowed from power-models either available from the previous generation of the chip or created from scratch. Since creating new power-models involves building a large infrastructure, such a power estimation methodology is suitable only for design problems where power-models of high accuracies are already available.

Thus, an ideal high-level power estimation methodology involves building a complete infrastruc-

ture with facilities to perform accurate power analysis and optimization. Such a methodology will target design problems that need accurate power numbers but do not have the scope of utilizing the already existing power-models. However, such a methodology is difficult to build because: (i) High-level design flows and tools have not sufficiently matured and (ii) infrastructure required to enable the high-level power estimation is extensive - that is where our technique helps.

In the industry, accurate power estimation is currently performed mostly at the RTL and gate-level. Typically, RTL power estimation techniques require following inputs: (i) design description in Verilog/VHDL or other hardware description language, (ii) RTL simulation trace in the format of Value Change Dump (VCD), Fast Signal Database (FSDB), etc. and (iii) power characterized libraries (e.g. standard cell power libraries) to correctly estimate the power consumption of a design. However, it is not feasible to utilize these techniques for high-level architectural exploration and optimization because such techniques and flows are targeted to be specifically used by the RTL design teams. On the other hand, a good high-level power estimation methodology should allow a high-level designer, architect or modeler to quickly and accurately gauge the effect of various high-level modifications on the power consumption of the design without being required to completely move to RTL power estimation flow. Such a methodology will significantly reduce the complexity of the power estimation process at the high-level.

Our Approach: In this chapter, we propose a high-level power estimation methodology which reuses an RTL power estimation technique guided by the high-level inputs for providing reasonable power estimates. In this sense, the proposed technique assimilates the best of both the worlds and provides a feasible and efficient solution to the high-level power estimation problem.

In this approach, inputs are similar to an RTL power analysis methodology with the exception that the RTL simulation trace of the design is replaced with its high-level simulation trace for accurate power estimation at ESL. One requirement to enable such a flow is that a mapping must be established between the variables in the high-level model of the design and corresponding signals (such as those at the port boundaries) in its RTL implementation. For such a mapping to exist, we recommend utilizing a high-level synthesis (HLS) engine.

5.2 Rationale for Our Approach

In general, total power analysis time T_{total} for a design is given as,

$$T_{total} = t_d + t_a + t_p. \quad (5.1)$$

In Equation 5.1 -

1. t_d represents the time taken by the power-analysis tool in extracting the design information. This includes the time spent by the tool in elaborating the design and generating its intermediate representation.

2. t_a represents the activity extraction time taken by the tool for collecting power analysis related information for various signals of the design from a simulation dump (like a VCD file). This time is proportional to design size and simulation duration. Algorithm *VCDExtraction* corresponds to this part of the power analysis process.
3. t_p represents the time taken by the tool in performing various power calculations based on the power characterized libraries, activity values and other design information (such as fanouts).

A design can be considered as a top-level module, which instantiates m child modules. We assume that each child module consists of m child modules and h denotes the hierarchical depth. Assume that the design is simulated for t_{sim} duration. Now, consider the following notations -

n_s - Number of variables at high-level per module.

n_i - Number of intermediate signals.

$t_{s.a}$ - Activity extraction time at high-level.

$t_{s.p}$ - Power calculation time at high-level.

$T_{s.total}$ - Total power analysis time at high-level (based on high-level simulation).

Corresponding notations at RTL can be denoted as n_r , $t_{r.a}$, $t_{r.p}$ and $T_{r.total}$. Moreover, transforming a model from high-level to RTL involves the introduction of additional intermediate signals which are required to represent the model at RTL. Such intermediate signals are represented as n_i .

Thus, at the high-level $t_{s.a}$ will be proportional to m , h , n_s and t_{sim} , and can be given as,

$$t_{s.a} = k_s * n_s * m^h * t_{sim}, \text{ where } k_s \text{ is a constant.} \quad (5.2)$$

Similarly, for the RTL, the corresponding equation can be given as,

$$t_{r.a} = k_r * n_r * m^h * t_{sim}, \text{ where } k_r \text{ is a constant.} \quad (5.3)$$

Let's assume that on an average every variable at high-level can be represented in β bits. Thus, we can establish the following relationship between the number of high-level variables and RTL signals,

$$n_r = n_s * \beta + n_i. \quad (5.4)$$

Time taken in extracting the design information t_d can be assumed to be similar at high-level and RTL. Thus, using Equation 5.1 at the high-level, we get,

$$T_{s.total} = t_d + (k_s * n_s * m^h * t_{sim}) + t_{s.p}. \quad (5.5)$$

$$T_{r_total} = t_d + (k_r * n_r * m^h * t_{sim}) + t_{r.p}. \quad (5.6)$$

In Equation 5.1, power calculations time t_p is proportional to the number of signals of the design. However, for long simulation durations, activity extraction time is usually the bottleneck and is much larger than t_p . This is because size of the simulation dump increases drastically with the simulation time, and extracting information from such larger simulation dumps is a time-consuming process. Thus, we assume that $t_s \gg t_p$.

Comparing Equation 5.5 and Equation 5.6, overall power analysis speed up obtained by using the proposed methodology can be given as,

$$N_{PE} = n_r/n_s. \quad (5.7)$$

Applying Equation 5.4 in Equation 5.7, we get,

$$N_{PE} = \beta + (n_i/n_s). \quad (5.8)$$

Value of β depends on the TLM specification style, for most of the designs $n_i > n_s$ which also contributes to the overall speedup offered by the proposed power analysis methodology with respect to RTL power analysis approach. Our experiments as shown in Table 5.1 validate the above arguments. Certain assumptions used in this analysis may not hold in the practical scenario, however they are used here to prove the applicability of our approach.

This proposed approach for power analysis can be called as mixed probabilistic approach. During this process, activities of the input-output ports and some intermediate signals are correctly imported from the high-level simulation dumps. Using this information, faster probabilistic power estimation is done at RTL which contributes to the overall speedup. Based on the accuracy of probabilistic power estimation, overall results can be improved. Furthermore, an accurate mapping of the high-level variables to the corresponding RTL signals also contributes to the accuracy of such a power analysis methodology. Some high-level synthesis tools generate such a mapping file which can be used to further improve the power estimation accuracy of our approach.

5.3 Our Methodology

Our high-level power-estimation methodology starts with the high-level model of a hardware design and generates the corresponding power numbers. As shown in Fig. 5.1, the overall methodology can be divided into the following steps-

1. Convert the high-level model of the design to an equivalent cycle-accurate RTL model. This can be done manually or using an appropriate high-level synthesis engine. Most high-level synthesis tools [7, 8] facilitate capturing the specification of a design at a level of abstraction

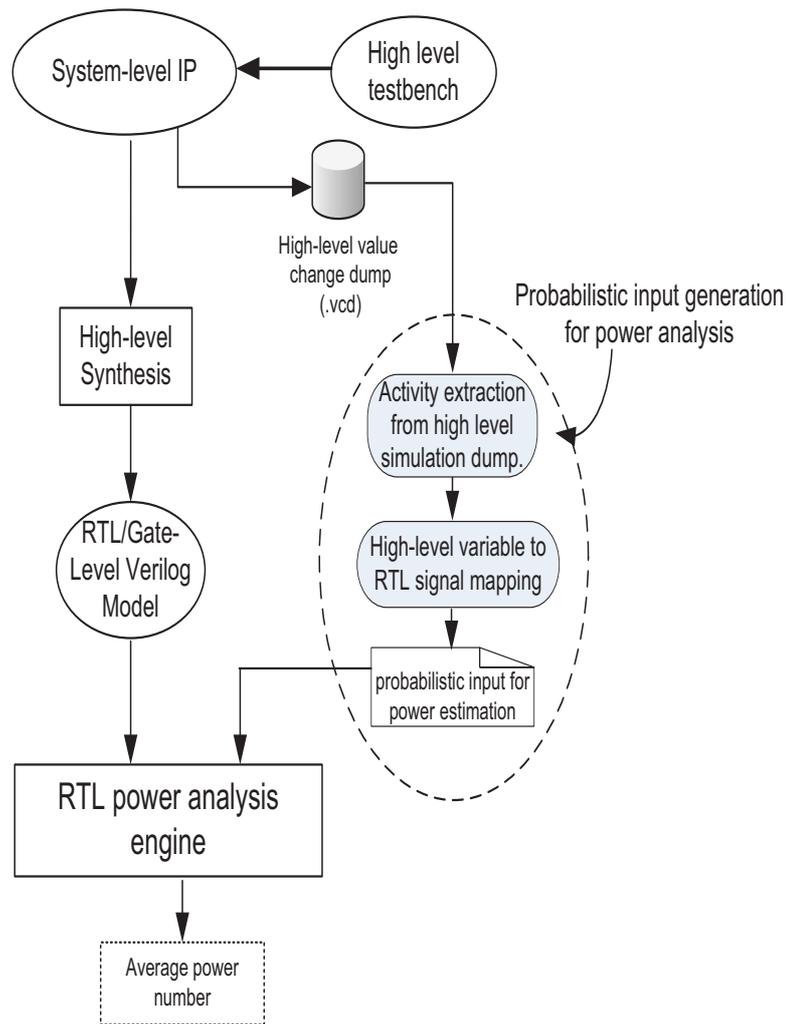


Figure 5.1: Detailed view of our power estimation methodology

above RTL followed by automatic generation of the corresponding RTL code. In our experiments, we have used Esterel Studio to synthesize the high-level ESTEREL models to their RTL implementations.

2. Simulate the high-level model and generate its VCD. For better accuracy of power numbers, all high-level variables of the model should be captured in the VCD.
3. Apply Algorithm *VCDExtraction* 5.3.1 on the VCD file generated during high-level simulation in Step 2. Algorithm *VCDExtraction* is used for activity extraction from the high-level simulation dump.
4. Generate the mapping of high-level variables to RTL signals as described in Section 5.3.2 and use the outputs of Algorithm *VCDExtraction* to create appropriate inputs for performing

probabilistic power analysis using RTL power estimator (such as *PowerTheater* [84]). Such inputs include power-related information for the mapped RTL inputs, outputs, and other intermediate signals, which can be then probabilistically propagated by the RTL power estimator for finding the activity information of the remaining signals and generate the power numbers.

In our framework, high-level power-related information extracted from the algorithm *VCDExtraction* is provided as input to *PowerTheater* in a particular format known as global activity format (GAF). Note that in pure RTL probabilistic power-estimation technique, such inputs (information related to the activity of ports and some intermediate signals) are usually provided by the user, and in some cases default values of activity factors are used. On the other hand, for the mapped variables/signals of a design that are visible at both high-level as well as RTL, their high-level simulation information extracted in our framework is exactly same as the information which would have been generated during the RTL simulation, and hence our approach can provide better results than pure probabilistic approach.

5. Collect the power numbers reported by the RTL power estimation tool.

5.3.1 Activity Extraction from the High-level VCD

Activity and duty-cycle extraction of various variables of a design is an integral part of our proposed methodology. Consider a high-level design S with the following notations-

$$M = \{ m_i : m_i \text{ is a module of design } S \}; V_i = \{ v_{ij} : v_{ij} \text{ is a variable used in module } m_i \in M \}.$$

Algorithm *VCDExtraction* extracts the hierarchy, activity-factor and duty-cycle related information from a VCD file generated during high-level simulation. Activity-factor and duty-cycle are computed using Equation 5.9 and 5.10 respectively.

$$\text{Activityfactor}, a_k^{ij} = (N_k/N_h) \quad (5.9)$$

$$\text{Dutycycle}, \delta_k^{ij} = (t_k/t_{total}) \quad (5.10)$$

In Equation 5.9, N_k denotes the total number of times a variable's value changed during the simulation and N_h represents highest toggles that can occur on a variable's value (in general, clock of the design has the highest toggles). In Equation 5.10, t_k represents the time duration for which the k^{th} bit of a variable remains in a boolean value **True** and t_{total} represents the total simulation duration.

Input to the Algorithm *VCDExtraction* is a high-level VCD file and its output is the duty-cycle, activity-factor and hierarchy information associated with each signal of the RTL design. The algorithm can be divided into two parts-

1. **Part a** - Extraction of following information: (i) hierarchy for each variable, (ii) for each value of a variable, number of simulation ticks for which the value remains unchanged and (iii) order in which various values of a variable changed during the simulation.
2. **Part b** - Calculate the activity-factor and duty-cycle of various signals using Equation 5.9 and Equation 5.10 respectively.

Algorithm VCDExtraction

Input : High-level VCD file.

Output: Hierarchy, Activity and Duty-cycle of various RTL signals.

Procedure *Part (a)* {

$\forall m_i \in M$ **do**

$\forall v_{ij} \in V_i$ **do**

- a.1- Store v_{ij} 's hierarchical name w.r.t the topmost module.
- a.2- Collect the changes occurred in the value of v_{ij} and the order in which its values changed.
- a.3- Estimate the simulation time spent in each value.

}

Procedure *Part (b)* {

\forall bit b_{ijk} of v_{ij} from LSB to MSB **do**

- b.2- Using Equation 5.10, and information estimated in *Part a.3*, estimate the duty-cycle of b_{ijk} .
- b.1- Using Equation 5.9 and information collected in *Part a.2*, calculate the activity factor of b_{ijk} .

}

Return hierarchical information, duty-cycle and activity-factor for various signals of the design.

5.3.2 High-level Variable to RTL Signal Mapping

High-level model of a design is an abstract model capturing the functional behavior of the design, and hence it contains much lesser information compared to the corresponding RTL or gate-level implementations. An ESL designer's job is to make sure that high-level model is compliant to the specification of the given design and to ensure, via architectural trade-offs, that the implementations generated (using a high-level synthesis tool) from such a model meets all the constraints on the area, power and latency of the design. Our proposed power-estimation methodology can be used in such scenarios. When synthesis tools create RTL models of designs from their high-level

models, they provide either a mapping file containing the information about high-level variables and corresponding lower-level signals or some guidelines to find such relationships between the high-level variables and lower-level signals. Esterel Studio provides such a facility for mapping inputs, outputs and intermediate variables of the high-level model to the RTL signals. In this methodology, we used this mapping to map activity information of all the high-level variables to RTL signals for power estimation purpose.

Variables at high-level may or may not have same name in the generated RTL code, and thus proper mapping of high-level variables to RTL signals should be created for accurate power estimation. We briefly provide an overview of the information we extracted from our experimental setup. In this paper, we extracted this information based on certain rules or synthesis guidelines which are specific to Esterel Studio synthesis. More detailed information exists in the user manual of Esterel Studio [18]. It should be noted that generating such a mapping is one time process for the analysis of a design because once this information is available then it can be used for many simulation runs. Activity of high-level variables extracted from algorithm *VCDExtraction* is then used for the generated signals/nets (in case their name gets transformed because of HLS).

Specifics for our experiments

We briefly provide overview of the information we extracted for our experimental setup. Note that different synthesis tools may provide such information in different formats, and for our experiments, we extracted this information based on certain rules or synthesis guidelines which are specific to Esterel Studio synthesis. More detailed information exists in *monolithic HDL code generation chapter in user manual of Esterel Studio* [18]. Below, we capture the salient points considered for our analysis.

- A signal that is exchanged between a module and its external environment is called an interface signal (similar to input output in HDLs). Once the RTL code is generated the interface signals are flattened in generated code. Listing 5.1 shows an example:

Listing 5.1: Interface in ESTEREL

```
1 interface A:
2   input i;
3   output o;
4 end interface
5 interface B:
6   input j;
7   port p: A;
8 end interface
9 .....
10 port q: B
```

In the example shown in Listing 5.1, B has a port p of type A and port q is defined as of type B. While generating the RTL code for port q, Esterel Studio will flatten the input and output variables and in generated RTL code these variables looks as shown in Listing 5.2:

Listing 5.2: Interface signals in generated Verilog

```

1 input q-j;
2 input q-p-i;
3 output q-p-o;

```

- Signal names such as $V7_sb_X$ may appear for intermediate signals used for connecting logic-gates, etc. in the generated RTL corresponding to the high-level variable X .
- A module M having an interface or intermediate variable X at high-level will be translated into two signals. Suppose the variable is represented as :

input X : bool;

then the generated RTL code will see the following signals:

input X;

input X_data;

ESTEREL has the notion of full signals; generated RTL code will see two signals corresponding to the status (X) and actual value (X_data), where status is an extra control information, if status is *true* then actual value will be updated else it will remain unaffected.

One should note that generating such a mapping is one time process for the analysis of a design because once this information is available then it can be used for many simulation runs. Activity of high-level variables extracted from algorithm *VCDExtraction* is then used for the generated signals/nets (in case their name gets transformed because of HLS).

5.4 Results

Tools Used: In our experimental setup, we have performed RTL and high-level simulation using VCS [131] and Esterel Studio [18] respectively. We have used 180nm power characterized *typical case* process libraries. We have used VCS 2006.06 release for RTL simulation, Esterel Studio 5.4.4 (build i2), and PowerTheater 2008.1 release for power estimation purposes. We have performed all our experiments on Dell optiplex (GX620) machine (operating system redhat enterprise linux (RHEL)) with 2 GB RAM and 3.39 GHz CPU speed.

Experiments: Table 5.1 captures the results obtained on different design IPs modeled at high-level. Entries in Column 1 (Design) represents the design type, column 2 (P_{RTL}) represents the power numbers obtained using RTL power estimation technique, Column 3 (P_{SPAF}) represents the power numbers obtained using our methodology and column 4 (E) shows the error with respect to the RTL power estimation approach. Column 5 (G_{VCD}) represents the ratio of the size of VCDs obtained at RTL and high-level to provide an overview of the difference of information needed at lower-level design flow.

Among various design IPs, we have used VeSPA processor model (which is a 32-bit processor with five stage pipeline) [132], and implemented the design in behavioral style, for this case study model is written as instruction set simulation model in ESTEREL. Other examples in this case study include power state machine (PSM), which is a prototype of a controller/state-machine used in modern processors for power management purposes. Other designs are finite impulse response filter (FIR), fast fourier transform (FFT), universal asynchronous receiver and transmitter (UART), etc.

$$E = |(P_{RTL} - P_{SPAF})/P_{RTL}| * 100 \quad (5.11)$$

$$G_{VCD} = G_{RTL.VCD}/G_{SPAF.VCD} \quad (5.12)$$

In Equation 5.11, P_{RTL} is the power number obtained using traditional RTL power estimation approach and P_{SPAF} is the power number obtained using our approach. In Equation 5.12, $G_{RTL.VCD}$ represents the disk space used by VCD in RTL power estimation and $G_{SPAF.VCD}$ using 5.3.

Design	P_{RTL} (mW)	P_{SPAF} (mW)	E (%)	G_{VCD}
PSM	1.23	1.16	5.7	168
FFT	12.6	13.7	8.7	50
FIR	21.5	22.2	3.3	1189
RAM	.451	.419	7.12	8
VeSPA	132.4	125.95	4.87	60
UART	11.7	10.8	7.69	33

Table 5.1: Results of our approach on different designs

Instruction Type	P_{RTL} (mW)	P_{SPAF} (mW)	E (%)
ADD	159	153	3.77
AND	119	107	10.08
NOT	130	131	0.77
XOR	151	150	.66
OR	125	123	1.6
CMP	159	153	3.77

Table 5.2: VeSPA processor instruction wise power numbers

We can see difference in accuracy (Table 5.1) for various benchmarks is ranging from 3-9% We attribute these differences to various facts a-) these models are modeled by four different people so modeling style might be different among the modelers for the same synthesis guidelines; b-) probabilistic activity propagation may suit particular type of logic/design, e.g. most of the instructions in VeSPA processor except AND instructions for random test patterns are showing good

accuracy (within 4%); c-) there can be scenarios where intermediate signals at RTL simulation (that cannot be related with the high-level model) have very high activity factor but in our methodology power estimation tool is propagating the activity value based on fixed probability. But one should note that results are consistently within 10% and can be further improved if one can establish better ways of establishing activity relationship between generated intermediate RTL signals and high-level variables.

Earlier presented work for the high-level models (such as [44], [46]) also provide the similar results. Approach discussed in these references utilize the power models for power estimation (mostly in specific scenario or for particular application). However our approach is an automation solution, which aids high-level synthesis framework to get RTL like accuracy with some extra effort.

Power Estimation Speed-up: We have seen 2-12 times speedup for our benchmarks. Power estimation time speedup for FIR design using our approach is 12 times and speedup in activity extraction time $t_{s.a}$ is 26 times. We compared the speedup numbers by running the designs for few minutes of CPU time to obtain these speedup numbers. In case of VeSPA design where we got the simulation bug, we ran few seconds simulation to get the accuracy numbers (random tests of 10000 clock cycle duration). We found that for smaller simulations tool takes almost constant time in seconds, which can be attributed to the fact that tool require some initial setup time to run these simulations, hence bigger runs are required to compare the speedup numbers. Such bigger simulations caused a simulation bug (out of memory) at high-level and we could not report the speedup numbers for all the designs. Efficiency of our approach is very much dependent on the amount of information being processed by the power estimator. As shown in the Table 5.1, upto 1100 time disk space reduction is possible for simulation dump size.

Chapter 6

Applying Verification Collaterals for Accurate Power Estimation

6.1 Introduction

In this chapter, we present a methodology for accurate power estimation at high-level, which utilizes the existing verification or validation resources in the design flow. This methodology can help in developing an infrastructure to estimate the power consumption at higher-level. Current power estimation methods are limited to RTL or lower level, in our approach we utilize best from both the worlds (RTL and high-level) and try to get the RTL like accuracy using high-level simulation only. We provide steps, examples of different properties using assertions, and methods to create directed testbenches as verification resources in the design flow for power estimation purpose. The approach presented in this chapter is specific to a high-level synthesis framework. We use Esterel Studio and Power Theater (to show the efficacy of our results). Our case study is performed on a finite state machine (FSM) based design. This case study presents detailed description of properties to generate the directed testbench for reaching a particular state of the FSM. Once we have the stimulus for the particular state of the FSM, we estimate the power of that particular state of the design at RTL, which is further utilized for the high-level power estimation.

Our approach utilizes verification/validation resources existing in the design flow to activate different key states and capture power related information. Simulators, checkers, coverage models, assertions, test content and tools developed during a typical validation flow are collectively called *verification collaterals*. Verification collaterals are employed to enhance correctness and to improve confidence in the high-level model that it satisfies the design requirements. Our focus in this chapter is specific to the creation of directed testbenches or the development of assertions and how they can be utilized downstream to characterize power numbers for different aspects of a high-level design such as states and transitions. The advantage of our approach is the ability to: (i) reuse existing verification collaterals and (ii) estimate power at high-level without rebuilding the complete

infrastructure. Such an approach helps in quickly narrowing down power-specifics earlier in the design cycle, which facilitates a realistic and rapid power exploration for a design.

6.2 Our Methodology

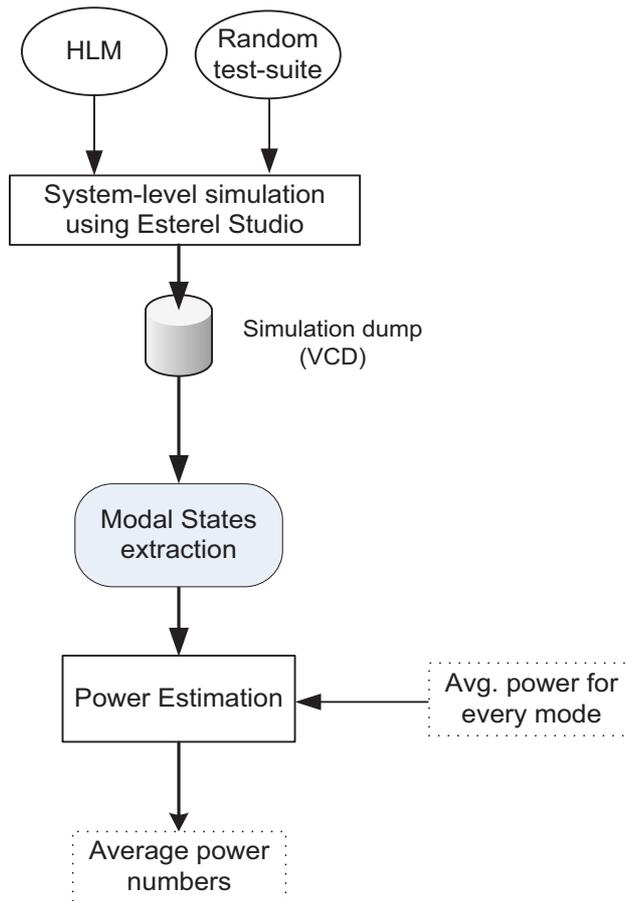


Figure 6.1: Detailed view of system-level power estimation technique

Fig. 6.1 shows the detailed view of our technique at system-level utilizing the average power associated with each mode and transition. Figure 6.2 provides the detail on how to obtain the average power number associated with each mode and transition. For applying our power estimation technique, designer require a high-level designing and verification framework. We are using Esterel Studio for designing and verifying the High-level models (HLM). Our HLMs are written in cycle-accurate transaction level (CATL) style in ESTEREL. From high-level simulation (performed in Esterel Studio) we extract the total time spent in each state, transitions and the total number of transitions. We also utilize the average power consumption in each mode at RTL (using PowerTheater)

and calculate the power at system-level. The main advantage of our approach is that we do not require to do the power estimation at RTL for every testbench from the testsuite for the whole design. We utilize average power associated with each mode for estimation purpose at system-level.

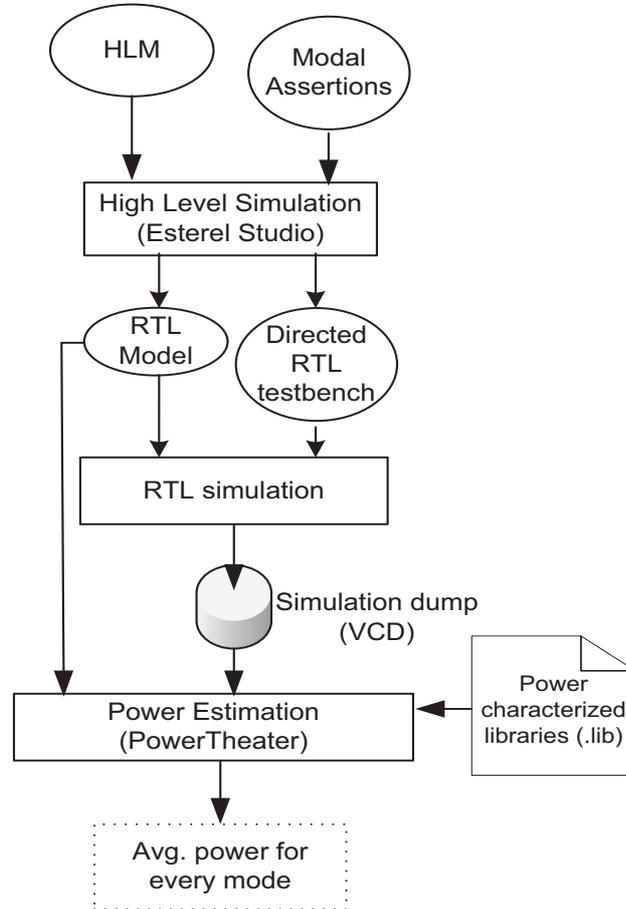


Figure 6.2: Power estimation for each mode at RTL

Following are the steps that needs to be followed for power estimation in this framework (Fig. 6.1) :

1. Simulate the High Level Model (HLM) written in ESTEREL using Esterel Studio and store the Value Change Dump (VCD) for further analysis.
2. Extract different modes and transitions, total number of transitions, total time spent in each mode and transitions from the simulation dump as discussed in section 6.2.2.
3. Once we have the estimate on time spent in each state and transition, power estimate of each state and transition then calculate the average power of the design using system-level

simulation as discussed in section 6.2.3. This utilizes the average power number for different modes using assertions as shown in Fig. 6.2. Here are the steps to obtain the average power number for each mode/transition (Fig. 6.2):

- (a) We create the cycle-accurate RTL model corresponding to the system-level cycle-accurate description in ESTEREL using the high-level synthesis engine provided by Esterel Studio (ES). ES is also used to convert the modal assertions (as discussed in section 6.2.1) to directed RTL testbenches.
- (b) Once the RTL model and testbenches are ready, we simulate the RTL model using RTL simulator (VCS [131]) and generate the simulation dump.
- (c) The dump generated at RTL, power characterized libraries and all the required inputs are then passed to the RTL power estimation tool (PowerTheater). The average power numbers calculated at this stage are then utilized for power estimation during system-level simulation. This is one time process for as these numbers can be utilized for multiple system-level simulations for power estimation.

6.2.1 Assertions for Finding out Particular Mode of Design

In a typical design flow, properties are written for verifying the behavior of the design. In a design flow where design can work in different modes, one can write assertions to excite the different modes (states) of a design and to trigger a mode change (transition). We assume that the properties are expressed as assertions and enforced during the verification stage of the design. Such that they are available during the power estimation stage, if not then we write the assertions for testing the reachability of states and transitions. We illustrate some reachability assertions written for a simplistic modal design shown in Listing 6.1. Note that the design is specified as a pseudo-code and not in any specific language.

Listing 6.1: A simple modal design

```

1
2 module XYZ
3 begin
4   bool c1, c2;
5   string state; // Possible values A, B, C, D
6
7   if(c1 == 0 && c2 == 0) then
8     state = "A";
9   else
10    if(c1 == 1 && c2 == 0) then
11      state = "B";
12    else
13      if(c1 == 0 && c2 == 1) then
14        state = "C"
15      else
16        state = "D"
17      end if
18    end if
19  end if
20 end

```

The design shown in Listing 6.1 has a variable `state` that is used to capture the mode of the system, which basically has four different values. The Boolean control variables `c1` and `c2` are used to trigger the required mode change. Assertions written for state reachability properties are shown in Listing 6.2.

Listing 6.2: Assertions to verify the reachability of states

```

1
2 // Property specifies the condition that reaches state A
3
4 assert always ( (c1=0 && c2=0) -> (state='A') );
5
6 // Property specifies the condition that reaches state D
7
8 assert always ( (c1=1 && c2=1) -> (state='D') );

```

An assertion that triggers the system to change from one state to another is shown in Listing 6.3.

Listing 6.3: Assertions to verify the reachability of a transition

```

1
2 // Property specifies the conditions that causes a transition from state D to C
3
4 assert always ( (c1=1 && c2=1 && next (c1=0)) -> ((state='D') && next(state='C')) );

```

The assertions are utilized for creating directed test cases such that each testbench either drives the system to a mode or causes the system to change its mode of operation. We utilize Esterel to express the assertions and Esterel Studio to generate the corresponding testbenches. Note that these directed testbenches are given as input to PowerTheater to estimate the expected power in the various operating modes of the system. These power numbers are reused at system-level to compute the overall power consumption of the system.

6.2.2 Extraction of Modes from the Simulation Dump

As discussed earlier at high-level there are control signals associated with each of the mode, any change in the value of these control signals is dumped in the simulation-dump. Signals for these modes can have value depending upon its type e.g. if there is a boolean associated with mode then we check when the signal is true, else if it is short then we check the exact value of the variable associated with that mode. We extract all the possible time-stamps for which the mode signal remains in the expected value and how many number of times design goes to particular mode for calculating the total time spent in each mode. This can easily be done as VCD contains all the information related to the value of the mode. Similarly we calculate the total number of transitions (from one mode to another) occurring in the dump. This knowledge is then used in doing system-level power estimation of the design. We wrote a C-shell script to extract the total time spent in each mode/transition, and the number of times design comes to that mode/transition.

6.2.3 High-level Power Estimation

To calculate the high level power we first try to establish a relationship using energy spent in each mode during the full simulation duration and then we establish a relationship for the power at each state. Lets say there are n different states and m different transitions in the design. Energy spent for the state i can be represented as E_i and total time spent in the state i can be represented as t_i . Similarly E_j is the energy associated with transition j , t_j is the total time spent on the transition j . If E_{total} is the total energy consumed in the design then we can establish the following:

$$E_{total} = \sum_{i=1}^n E_i + \sum_{j=1}^m E_j \quad (6.1)$$

If total average power is P_{total} and total simulation duration is T then from Equation 6.1 we can establish the following:

$$P_{total} * T = \sum_{i=1}^n P_i * t_i + \sum_{j=1}^m P_j * t_j \quad (6.2)$$

If a_i, a_j is represented as a fraction of total simulation duration spent in state i and transition j , we establish a relationship between a_i, a_j, t_i, t_j and T as shown in the Equation 6.3, 6.4.

$$a_i = t_i/T \quad (6.3)$$

$$a_j = t_j/T \quad (6.4)$$

In Equation 6.2 if we divide both the sides by T then we will get the following:

$$P_{total} = \sum_{i=1}^n P_i * a_i + \sum_{j=1}^m P_j * a_j \quad (6.5)$$

Which can further be simplified as:

$$P_{total} = \sum_{i=1}^n P_{Hi} + \sum_{j=1}^m P_{Hj} \quad (6.6)$$

In Equation 6.6, P_{Hi} and P_{Hj} represents the component of expected average power spent in each state and transition respectively. Equation 6.5 and 6.6 establishes the relationship between average power of each state, transition and total power calculated at high-level. In our approach P_i and P_j is calculated at RTL by utilizing the assertions written during verification stage (as discussed

in section 6.2.1) and a_i and a_j is calculated as discussed in section 6.2.2. Finally P_{Hi} , P_{Hj} is calculated from the values of P_i , P_j , a_i and a_j .

6.3 Case Study

PSM Specification

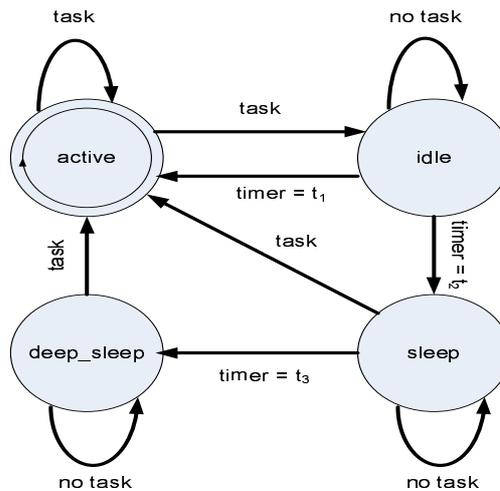


Figure 6.3: Macro-state diagram of PSM

State diagram of the PSM is shown in Fig. 6.3, more details about the design is available at [133].

PSM is an essential ingredient in reducing the power consumption of the system by regulating the power states. It is distributed along the design; hence the validation of the PSM requires simulating almost the whole system. In Figure 6.3, we show the four different states of the PSM namely `active`, `idle`, `sleep` and `deep_sleep`. Its four different constituents namely *Queue*, *Timer*, *Sampler* and *Service Provider* are shown in Figure 6.4.

- **PSM:** The system remains in `active` state if it is receiving a task. If no task is received for t_1 clock cycles, then it goes to the `idle` state. In the `idle` state, if it does not receive any task for t_2 cycles, then it goes to the `sleep` state. In `sleep`, if the system does not receive any task for t_3 cycles, then it goes to the deep sleep state `deep_sleep`. From any of these states, the system returns to the `active` state if it receives a task.
- **Sampler:** All power states have a sampler that samples a task every t_4 cycles.
- **Queue:** A task received by the system is stored on a finite queue, which is accessible to every state.

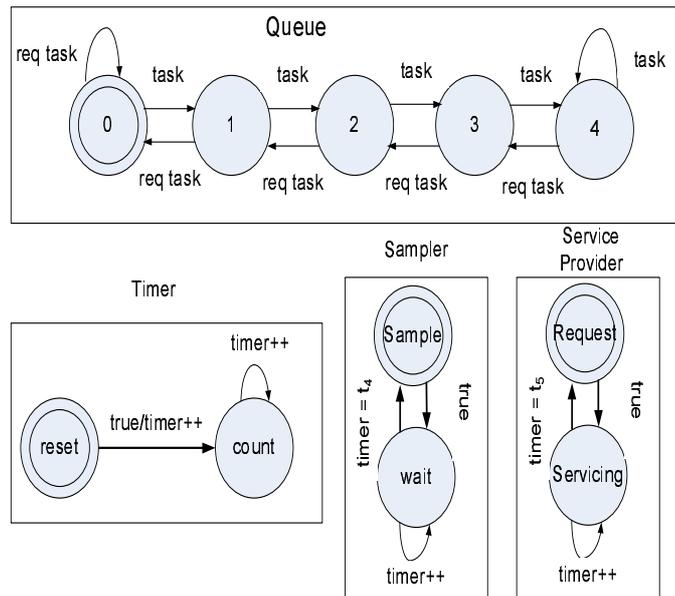


Figure 6.4: Specification of the different PSM ingredients

- **Service Provider (SP):** In the `active` state, the SP processes a sampled task for t_5 cycles. The SP requests for a task and receives a task that is de-queued.
- **Timer:** Every state except the `deep_sleep` has a timer embedded, which tracks for how long the system is idle without receiving a task.

Verification of the PSM is very important because such a component is often used in system level power management for reduction of power in embedded systems and this type of controller fits the design type we are targeting in this chapter.

One can reach to a particular state by putting the property in form of assertion on the state. Same stimulus can be created manually by writing the directed testbench from the specification of the design. We explain both the approaches in detail in this section.

6.3.1 Directed Testbench Creation from the Specification

We illustrate this approach by the example specification discussed in the Section 6.3. Let us assume that we want to obtain the idle state power numbers of the design. As discussed in the Section 6.3, design goes to idle state (from active state) if it does not receive any input task for t_1 cycles. Design will remain in the idle state for t_2 cycles if it still does not receive any task. Hence, one can write a testcase in which design will not receive any input task for t_1 cycles, after that design will remain in the idle state from t_1+1 until $t_1 + t_2$ cycles. To capture the power information of the idle state one should run power estimator from (t_1+1) to $(t_1 + t_2)$ clock-cycles to capture

the average power consumption of the design in idle state, although stimulus should atleast have continuous ($t_1 + t_2$) clock-cycles on which design does not receive any input task.

Similarly, one can write the testcases for all the valid transitions. For example, in the case discussed above at cycle t_1 , the transition from active state to idle state will occur. If, there is a good power management system working in the system then lot of computation unit in the design won't be in active state and hence in such a situation there will be sudden drop in the power consumption of the design. Now, to obtain transition power from one state to another instead of measuring the average power over certain clock cycles, we did clock by clock power measurement for the clock cycles when these transitions are occurring and collected the power consumption for the transitions. Such measurements help us in improving the accuracy of the average power numbers.

We can observe there are four different macro states i.e. active, idle, sleep and deep_sleep state respectively. From the state transition diagram, we see that one can write testcases for all the valid transition scenarios that include active to idle, idle to sleep, sleep to deep_sleep, idle to active, sleep to active and deep_sleep to active, etc.

The detailed explanation on how one can obtain the power estimation at high-level is provided in [29]. In this section, we briefly provide the description on how we have used that method in our case. As discussed above, we can first create directed testbenches from the specification and obtain the average power numbers for these scenarios using RTL power estimator PowerTheater [84].

We generate the high-level simulation dump (VCD), which has the information of these state variables, for any arbitrary testvector. We calculate the time spent on each state during the simulation, From the VCD. We then calculate the total energy spent on each state, which is the product of power and time spent i.e. clock-cycles in a particular state, with the information extracted from the earlier steps. Then we sum up the energy for each state/transition and average it over whole simulation duration to calculate the average power.

6.3.2 Utilizing Assertions for creating Directed Testbench

Various properties to reach/stay in a particular macro-state can be written for the design. The purpose of this step is to provide an approach to automate the testbench generation process using simple assertions. We propose two type of approaches A-) By writing a verification monitor for the whole design from which one can monitor input and output values and then insert a property to see weather design follows it. B-) As a part of design, writing some properties on states or timer values to see if design can ever reach to those states or values.

Writing Properties using Verification Monitor

For verification of certain properties, we create a new toplevel in which we put the design as *design under verification (DUV)*. From the verification monitor, we can control the inputs, observe

the outputs, and verify the properties of the design.

- *Assertion to reach a particular state of the design:* As explained earlier to reach to idle state of the design, it should not receive any input task for t_1 cycles and if we want the design to stay in this particular state then input task should be kept false for $(t_1 + t_2)$ clock-cycles. Listing 6.4 shows how one can write the assertion whose negation will force the Esterel Studio to provide a counter case that design will go to idle state. In the listing shown we force input task not to receive any input for $(t_1 + t_2)$ clock-cycles using repeat construct. One can use the for loop, etc. Just after that we put an assertion saying that input task can never be true. Note that we set this property as an always true property and then ask Esterel Studio to prove this property, for which the verification engine generates a testcase showing that input can be false for more than $(t_1 + t_2)$ clock-cycles. State specific power estimation in this case will be similar to the one explained in Section 6.3.1.

Listing 6.4: Assertion for state reachability

```

1 repeat ?t1+?t2 times
2   emit ?input_task = false;
3   await tick;
4 end repeat;
5 assert idleStateReachability = not(?input_task = false);

```

- *Assertion for a transition from a state:* Example explained above will contain the transition from active to idle state because no task was received till $(t_1 + t_2)$ clock-cycles. similarly idle to sleep and sleep to deep_sleep transitions can be captured using the property shown above. For transitions from idle or sleep or deep_sleep state to active state transition cannot be generated using the property above. Listing 6.5 shows the transitions scenario that can be captured using Esterel Studio which is written as never true property. Now the never true property is written in such a way that we force the verification monitor not to receive any input task for certain clock-cycles and then we ask that design cannot receive any task further. The counter case of this case is putting no input task till certain cycles and then allow the input task, which forces the design to move to active state from idle/sleep/deep_sleep state etc. Listing 6.5 shows the transition to active state from the deep_sleep state.

Listing 6.5: Assertion for state transition

```

1 repeat ?t1+?t2+?t3 times
2   emit ?input_task = false;
3   await tick;
4 end repeat;
5 assert idleActiveTransitionStateR = not(?input_task = true);

```

In this case study, we show controlling one input because that can help us in identifying the state but this approach can specially be helpful for the cases where user wants to control the multiple inputs in parallel. In that case, such an approach can be used to automatically generate testcase using assertions.

Writing Properties on Design States/Variables

- *Generating a testcase to reach a state:*

As discussed earlier, one can write assertions in the design and can set the property as always true or assumed using Esterel Studio. In case the property is not true then Esterel Studio generates a counter case for the scenario. For PSM, the code snippet for the design from high-level utilizing the macro-state behavior looks as shown in Listing 6.6.

Listing 6.6: Overview of the code for states in Esterel

```

1 if (?State = "active") then
2 .....
3 .....
4 else
5     if (?state = "idle") then
6         .....
7         .....
8     end if;
9 else
10    if (?state = "sleep") then
11        .....
12        .....
13    end if;
14 else
15    if (?state = "deep_sleep") then
16        .....
17        .....
18    end if;
19 end if;

```

Now, to create a testbench to reach to a particular state, one can write an assertion inside the design code of that state and putting a property that design can never remain in that state itself. For example if we want to see if design can reach to the idle state of the design then inside the code for the idle state, we need to write an always true property, which says that in idle state, design will never remain in the idle state (which is a contradiction). So the code for such a scenario will look as shown in Listing 6.7.

Listing 6.7: Assertion for state reachability

```

1 if (?state = "idle") then
2 .....
3 .....
4 assert idle_state_check = not(?state = "idle");
5 end if;

```

For the scenario stated above Esterel Studio generate a counter example that contains input which forces the design to remain in the idle state.

- *Testcase to reach a particular state and then remaining inside the state for some cycles using invariant:*

We have seen in the specification section that design remains in particular state for some cycles. We can write properties, in which we can verify, if design remains in that state for some time or not. For PSM example one can do the same by using the invariants for the

timer. In our high-level design we have a separate timer module. In Esterel studio we can write an assertion shown in Listing 6.8.

Listing 6.8: Assertion for staying in a state for certain cycles

```
1 assert idle_timer_check = not(?timer_count = ?t2);
```

Now for the case shown in Listing 6.8, we can set the Esterel Studio to test an always true property for the instantiation of the timer for idle state. From the specification document, we know that timer will always go to t_2 , if it does not receive any input task. Esterel Studio will be able to produce a counter example for the negated property.

- *Testcase for generating Transition Scenario:*

One can write the assertions to generate testcases in which design goes through the state transitions. Listing 6.9 shows such an example of the transition condition that can be used for creating the directed testbench which can be used to capture the macro state transitions.

Listing 6.9: assertion for extracting transitions

```
1 assert state_test = not ( (?state = "active") and (next ?state = "idle"));
```

Assertion shown in the Listing 6.9 represents the valid transition extracted from the specification document. Now this assertion is again negated and Esterel Studio option for always true is set to test such a property. From specification, we know that such case is trivial transition case hence Esterel Studio generates a negated scenario that can be used to capture the state transition of the design. The Listing 6.9 shows that design can never reach to idle state from active state, which is not true. In this case, Esterel Studio generates a counter case showing that design can reach to idle state from active state.

6.4 Results

In this chapter, we presented a case study of our approach on power state machine. We presented various properties one can write to obtain state specific typical stimulus input. We also presented an approach in which one can write those tests manually based on the specification provided for the design. Our results utilizing the manual approach shows good results with good accuracy. We automated the manual testbench creation using the approach discussed in this chapter. We presented our case study on small design, while the scalability of such an approach still needs to be tested on bigger designs. For example, using invariant based approach might be very costly sometimes, if design has a lot of states. Similarly capturing transitions for different states may take many cycles in that case, we will require instrumentation of the source code according to that specification.

As discussed in the Section 6.2.3, we calculate the time spent in each state, transition (a_j, a_k) for the whole simulation and the results are captured in the Table 6.1 and Table 6.2. P_j and P_k in

PSM state	active	idle	sleep	deep_sleep
% of total-time spend in each state	56.68	10.44	19.63	13.23

Table 6.1: Percentage of simulation time spent in each state at system-level

transitions	% of total-time spent in each transition
active→idle	.013
idle→sleep	.013
sleep→deep_sleep	.007
deep_sleep→active	.007
sleep→active	.007
idle→active	.000

Table 6.2: Percentage of simulation time spent in each transition

Equation 6.5 represent average power of each state and transition captured in Table 6.3 and Table 6.4. P_{H_j} and P_{H_k} in Equation 6.6 represents the weighted per state power and transition power, these numbers are reported in Table 6.5 and 6.6. We calculated the total average power using our methodology applied at system-level simulation and got the power number 1.253 mW for a randomly selected testbenches. To verify our result we created the equivalent RTL testbenches of the design for which we calculated the power using system-level testbenches and found out the power number P_{RTL} as 1.27 mW. Accuracy (E) of our power estimation methodology in this case is coming out as 98.66%, E is calculated using the Equation 6.7. All these results are captured in Table 6.7.

$$E = (P_{total})/P_{RTL} * 100 \quad (6.7)$$

From the result tables we can see that the minimum and maximum bound on power for this design ranges from .833 mW to 1.57 mW. We can also notice that if design has got a lot of transitions from idle, sleep, deep_sleep to active state or vice versa then we will also see a lot of variation in average power of the whole design because these transitions consume a lot of power. Transitions in this case are taking a lot of power, which is also understandable because if the system is coming from sleep state to active state then most of the design will see a lot of activity and hence a huge amount of power. Overall effect on total power is not too high because of transitions. Total time spent in transitions is less than 1% as shown in the Table 6.2.

state	active	idle	sleep	deep_sleep
power consumption (mW)	1.57	.834	.833	.833

Table 6.3: power estimation of the states at RTL

transitions	power consumption (mW)
active→idle	1.61
idle→sleep	1.61
sleep→deep_sleep	1.62
deep_sleep→active	12.72
sleep→active	12.72
idle→active	12.72

Table 6.4: Power estimation at RTL for each transition

state	active	idle	sleep	deep_sleep
power consumption (mW)	.8898	.0871	.1635	.1103

Table 6.5: power numbers at each state using our approach

transitions	power consumption (mW)
active→idle	.00021
idle→sleep	.00021
sleep→deep_sleep	.00011
deep_sleep→active	.00089
sleep→active	.00089
idle→active	.00000

Table 6.6: Percentage of simulation time spent in each transition

RTL average power (mW)	System-level average power (mW)	Accuracy w.r.t RTL (%)
1.270	1.253	98.66

Table 6.7: Estimation accuracy as compared to RTL

Chapter 7

Power Reduction using High-Level Clock-gating

Hardware co-processors are used for accelerating specific compute-intensive tasks dedicated to video/audio codec, encryption/decryption, etc. Since many of these data-processing tasks already have efficient software algorithms, one could reuse those to synthesize the co-processor IPs. However, such software algorithms are usually sequential and written in C/C++. High-level Synthesis (HLS) helps in converting software implementation to register transfer level (RTL) hardware design. Such co-processor based systems show enhanced performance but often have greater power/energy consumption. Therefore, the automated synthesis of such accelerator IPs must be power-aware. Downstream power savings features such as clock-gating are unknown during HLS. Designer is forced to take such power-aware decisions only after logic synthesis stage, causing an increase in design time and effort. In this chapter, we present a design automation solution to facilitate various granularities of clock-gating at high-level C description of the design.

7.1 Introduction

Adoption of Electronic System Level (ESL) based methodologies in hardware design flow is on the rise. C/C++/SystemC are the description languages used to represent the design behavior at the ESL. One of the main ingredients of ESL based design flow is high-level synthesis (HLS). HLS helps in creating RTL/gate-level description of the design from the behavioral or system level description. In the recent past, use of hardware co-processors was advocated for specific domains such as digital signal processing (matrix multiplication, filtering operations), security (encryption/decryption), etc. These signal processing or security related tasks are focused on data-intensive computing for which efficient software algorithms exist in C/C++. Earlier these tasks were exclusively done in software. With the availability of more gates on the silicon die, use of co-processor/accelerator to perform these data-intensive computations using co-design method-

ologies is on the rise. While such co-processor based acceleration can enhance performance, it might impact power/energy consumption of the system. Therefore, the automated synthesis of such accelerators must be power-aware.

HLS of such efficient software algorithms will greatly help in reducing design cycle time as well. It is beneficial to reuse these time-tested and validated software algorithms. It is even more beneficial if one can use automated synthesis to convert these co-processor IPs from their software implementations, rather than manually recasting the algorithm into low level hardware implementation. One hurdle in such a strategy is that these software algorithms are targeted usually for sequential processing, written in C/C++. In contrast, the hardware implementation has to be highly parallel, clocked, leading to both efficiency and low power requirements. The hardware acceleration require intensive parallelism/pipelining/buffering techniques which mandates appropriate choice of their micro-architectures (parallel, pipelined, etc.).

The process of making this choice may require multiple passes through refinement/synthesis followed by performance/power estimation cycles. In a power-aware methodology, the power estimation must be accurate, and at the highest possible level of abstraction for faster convergence of the selection process. However, down-stream gate-level power savings features unseen at register transfer level (RTL), may render the exaggeration of the power estimates and force the designer to take decisions only after logic synthesis and estimation. This increases design time and effort. In this chapter, we show how to endow the HLS itself with the ability to generate RTL with the power-saving features that are normally inserted during gate-level synthesis. This allows realistic power exploration at the RTL, resulting in faster convergence of the micro-architecture selection, without compromising the quality of the generated hardware co-processor.

In an HLS based design flow, tool generated RTL model is passed through logic synthesis tools that would add various power saving features such as clock-gating. These power reduction features are often unseen at high-level. Therefore, the RTL power estimation may be grossly exaggerated, leading the designer to put more efforts to reduce power consumption at high-level. To alleviate this problem, we present a design automation solution to facilitate clock-gating at high-level. This will help a designer in controlling various granularities of clock-gating from the C description of the design.

We try to answer following questions in this Chapter:

- Is there a way to enable RTL power aware trade offs from C specification?
- If answer to the question above is yes, can a designer/architect make such trade offs with existing automated HLS flow?
- If answer to the question above is yes, how to measure the efficacy of the approach?

We try to answer the problems/questions posed above by providing our perspective on the need for register clock-gating at high-level C description in Section 7.2. In Section 7.3, we discuss how to

enable various granularities (such as function, scope, etc.) of clock-gating in C description, our algorithm, etc.

Main Contributions

The contributions of the presented methodology are following:

- Introducing a methodology implemented in a co-processor synthesis flow using C2R [22], for exploring power saving opportunities at the C-code level by
 - Exemplifying the complexities/issues related to assignment of clock-gating at a behavioral description level;
 - Strategizing the selection of granularities of clockgating opportunities in a behavioral description written in ANSI-C; and
 - Demonstrating a specific priority algorithm to handle various conflicts for enabling clock-gating in functions, scopes, variables, etc.
- Comparison of the quality of the synthesized RTL (after such exploration) against the quality of results obtained by gate-level power optimization tools such as power compiler using a few industry strength benchmarks such as AES, GZIP etc. design examples.

7.2 Why is clock-gating needed at High-level?

Timing, power or area can be controlled at architecture level by inserting directives in the C code or by applying intelligent scheduling and resource allocation schemes for synthesis. In case of power reduction, availability of tools to provide an estimation at the RTL with reasonable accuracy has opened up the opportunity to make macro/micro architectural power aware trade offs before RTL is finalized. Also clock-gating based power reduction sometimes leads to very high fidelity, it is very important to generate clock-gated RTL for these co-processors. One of the case studies presented in [134] shows that upto 70% of power reduction in design can be achieved with clock-gating alone.

Clock-gating is extensively supported at the RTL because designers have insight on how many registers are allocated to the datapath and controller of the design. If we compare it with respect to the C2R based co-processor synthesis, where only macro-architecture is fixed, enabling clock-gating is not as easy as it looks like. We introduce pragmas that can be inserted in behavioral code to enable clock-gating of registers during HLS. One should note that clock-gating can be enabled at different granularity from very coarse grain to fine grain such as register bank of 32bit to a single bit register, etc. By macro architecture we mean the architecture where design contains minimal

information about the implementation. Most of the times this design information contains various directives, processes to include parallelism and clock stages in the sequential C code. While the micro-architecture selection is done by the Cebatech C2R by applying design specific FSM to map the behavior of high-level specification.

In the high-level C description of the design, an architect can control the register power consumption by performing the gating for global/static variables. Similarly a case might arise to do the clock-gating of a register of a particular function (a hardware module at the RTL after synthesis) or even a scope of function. Such situations once handled at high-level will facilitate architect to control the register power of the design for various granularities. Design automation solutions, where HLS can be guided for register power reduction, will greatly help in providing early power-aware realistic design trade-offs.

Let us consider an example of pseudo C code as shown in Listing 7.1. In this example code a, b, c, and d are pointers to integers and foo is a function which is called by reference at two different places in the main function (line number 4 and 6). Now, Designer's/Macro-modeler's intent is to parallelly execute the two foo modules in the generated hardware. After doing quick power estimation, designer comes to know that he just wants to clock-gate first foo block but not the second foo. There is no distinction between the two foo calls at the C description level. However at the RTL, two different foo modules can be instantiated as foo1 and foo2.

Listing 7.1: C source code to enable clock-gating

```

1 int main ( )
2 {
3     int *a, *b, *c, *d;
4     foo ( a , b ) ;
5     . . . . .
6     foo ( c , d ) ;
7     . . . . .
8 }
9
10 void foo ( int *x , int *y )
11 {
12     int yyy ;
13     int xxx ;
14     . . . . .
15 }

```

Any C-based synthesis tool would require such an indication before generating RTL. On the other hand, to clock-gate various registers of foo1 and foo2 during logic synthesis designer just needs to provide information of these registers based on design hierarchy. In case of C description, we propose to use clock-gating(ON/OFF) directives, which will help the macro-modeler to control the granularities of various clock-gating decisions. Listing 7.2 shows the change in the code after applying the clock-gating directives. Details of different granularities and issues involved are available in Section 7.3.1.

Listing 7.2: C source code after inserting clock-gating directives

```

1 int main ( )
2 {
3     int *a, *b, *c, *d;

```

```
4 clock_gating(ON) foo ( a , b ) ;  
5 . . . . .  
6 clock_gating(OFF) foo ( c , d ) ;  
7 . . . . .  
8 }  
9  
10 void foo ( int *x , int *y )  
11 {  
12     int yyy ;  
13     clock_gating(ON) int xxx ;  
14     . . . . .  
15 }
```

7.3 How to enable Clock-gating at High-Level?

It is very important that the high-level designer can exploit clock-gating benefits from the source code of the design. Techniques for power estimation presented in earlier chapters can help in finding out the hot spots of the design at high-level. These techniques can provide 6-20 times speedup as compared to the RTL power estimation techniques with acceptable accuracy. Once at the high-level, hot spots of the design are known, designer would like to gauge the impact of clock-gating on the power consumption of the design. One should note that such an analysis is performed across various components, and granularity may vary from a register-bank to a whole module. A High level designer may just want to enable it at C description level because various scope/functions in the design are well known in the C description. Keeping this view in mind, below we present a few scenarios/situations that can be exploited from the C description to avail clock-gating at different granularities.

7.3.1 Application of clock-gating for various granularities at the source-code level

For the C specification, designer can explore the granularities of clock-gating in many ways. Such an exploration is important because it helps the designer to take various decisions to reduce power consumption of a specific part of the design. At the C specification level, the designer has a visibility on various functions, scopes, variables, global variables, etc. Once designer knows which part of the design needs to be clock-gated, he can utilize directive such as `clock_gating(ON)` or `clock_gating(OFF)`. Such directive used as macro in C environment won't affect functionality of high-level design in C simulation environment such as gcc [22]. It will direct the HLS to generate clock-gated RTL. Instead of using the directives, a script can also be used to direct HLS to clock-gate particular part of the design. Most of the times designer may want to clock-gate at coarse grain level such as scope, function (as discussed in Section 7.2). There exist some cases for fine grain decisions for which a careful analysis is required. We capture a few of these cases in this section and suggest a priority algorithm to take care of all the possible granularities for clock-gating. We

discuss in detail how different directives can be applied for various granularities in the restructured C specification.

- *Function level:* As discussed earlier in the Section 7.2 (Listing 7.2) different calls of functions can be clock-gated ON or OFF using the clock-gating directives to get the function level control to clock-gate the design.
- *Scope level:* In a high-level description in C, another possibility to clock-gate the design comes at scope level. Sometimes, at scope level, designer wants to create hardware for the part of design (corresponding to the scope) and at the same time he wants to clock-gate every register variable in that scope. Listing 7.2 shows designer can write a directive such as `clock_gating(OFF)` for not clock-gating a scope of function `foo`.
- *Variable level:* This is the simplest case in which user wants to enable clock-gating for particular variable. At compile time, if compiler sees a particular variable needs to be clock-gated because a `clock_gating(OFF)` directive is used so it will not clock-gate that variable, as shown in Listing 7.3.
- *Register bank:* Clock-gating every bit is sometimes very expensive because of its impact on area and timing. By clock-gating bank of registers such as 4 bits or 8 bits at a time, a lot of area can be saved because only one gating logic (a latch with and gate) can replace 4 and 8 multiplexers respectively.

7.3.2 Priority for clock-gating decisions

As can be seen in the descriptions above, we enable clock gating before RTL is generated. In our approach, we provide the user with a flexibility so that he can turn clock gating ON or OFF at different granularities (e.g., at function level, scope level, variable level, or based on register bank size). Such provision without any priority will lead to conflicts. For example consider the example shown in Listing 7.3. In this example first call to `foo` contains the `clock_gating(ON)` directive while the other call contains `clock_gating(OFF)` directive. Even inside the function `foo` description `clock_gating(ON)` is used for `yyy`, while for `xxx` `clock_gating(OFF)` is used. Synthesis tool will see this as a conflict because for first call of `foo`, `clock_gating(ON)` directive is used, which suggests that all registers of `foo` (at line 12) will be clock-gated. On the other hand for `xxx` `clock_gating(OFF)` directive is used. In such a scenario high-level synthesis should know which directive should take precedence. We prioritize the clock-gating decisions based on the flow-chart shown in Fig. 7.1.

Listing 7.3: Example showing Conflicting Clock-gating directives

```
1 int main ( )
2 {
3   int *a, *b, *c, *d;
4   clock_gating(ON) foo (a , b) ;
5   . . . . .
```

```

6 clock_gating(OFF) foo ( c , d ) ;
7 . . . . .
8 }
9 void foo ( int *x , int *y )
10 {
11 clock_gating(ON) int yyy ;
12 clock_gating(OFF) int xxx ;
13 . . . . .
14 {
15 clock_gating(OFF);
16 . . . . .
17 }
18 }

```

As shown in the flow chart in Figure 7.1, if clock-gating is enabled/disabled for a particular variable then it will get the highest priority; then we check if the clock-gating is enabled/disabled at scope/function level and finally at global variable level. Finally we check if the user sets a limit for register bank size, and if so, only those register-banks having size below that limit, will be enabled/disabled for clock-gating. In the example shown in Listing 7.3, for first call of foo, xxx will not be clock-gated while yyy will be clock-gated. Similarly for second call of foo, yyy will be clock-gated while xxx will not be clock-gated. One should note that high-level synthesis tool can be targeted with other priority but for our analysis we have kept the priority according to the flow-chart.

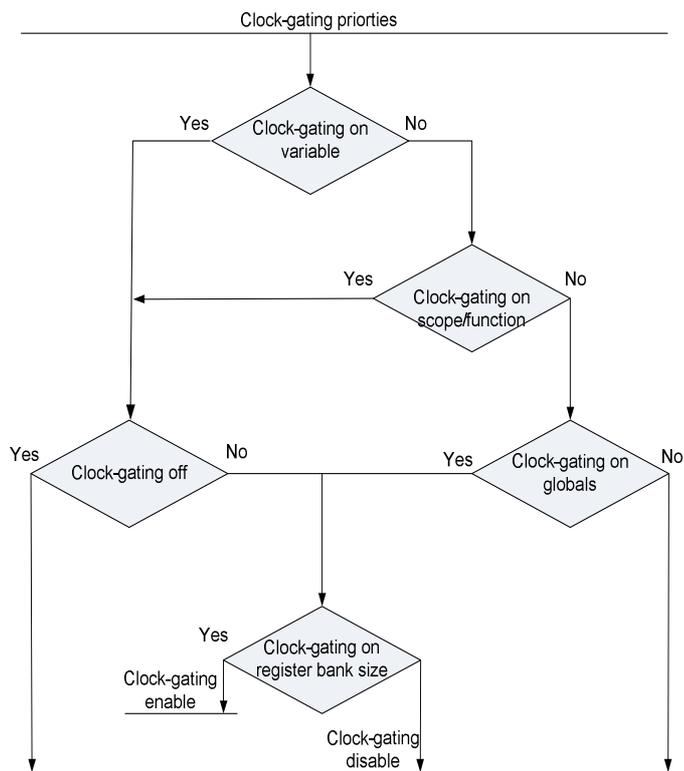


Figure 7.1: Flow chart for gating clocks of registers at high-level

7.3.3 Algorithm for enabling clock-gating in C based co-processor synthesis framework

Our methodology takes advantage of the internal representation of C2R HLS tool. C2R compiler internally creates a Control Data Flow Graph (CDFG), which is passed through scheduling and resource binding after several optimizations. At the last pass when the registers in the design are fixed, we apply the algorithm for clock-gating. This enhanced scheduled CDFG is then used to generate the RTL Verilog containing gated clocks. We use the latch based clock-gating discussed in [86]. Our algorithm first finds out total number of registers in the design. Then it tries to find out clock-gating candidates based on the flow-chart or priority algorithm discussed in Fig. 7.1. Once registers/register-banks are finalized for clock-gating, our algorithm tries to find out the common enable conditions that are used to clock-gate the registers. This helps in further reducing the area (as discussed in [86]). Once the enabling conditions and registers are finalized, registers are changed to clock-gated registers.

To describe the algorithm we need to first fix a few notations as following: Let R be the set of registers in a design D ; $NCG(r)$ be a predicate which is decided based on the flowchart in Figure 7.1, and it evaluates to true if r needs to be clock gated, else false. Also, let $R_{cg} = \{r \mid r \in R \wedge NCG(r)\}$. Thus $R_{cg} \subseteq R$. Let C be the set of all clocks in D and $Clk(r)$ return the name of the clock of a register r . Let $Clkgated(r)$ return the name of the clock of r after clock-gating. For a register bank E having the same enable signal, $Clkgated(E)$ returns the clock of the entire bank after clock gating. Let En_r represent the enable signal for register r . Also, let $EN = \{En_r \mid r \in R_{cg}\}$.

Algorithm *clock-gating of registers*

Input : Scheduled CDFG in the last phase of HLS

Output : synthesizable RTL Verilog

Initially : $R_{cg} = \emptyset$

$\forall r \in R$ **do**

 if $NCG(r)$ then $R_{cg} = R_{cg} \cup \{r\}$.

 Find the enable signal En_r for r ;

Partition $R_{cg} = \{R_i^{Clk}\}_{i \in C}$

such that $R_i^{Clk} = \{r \mid r \in R_{cg} \wedge i \in C \wedge Clk(r) = i\}$.

R_i^{Clk} is the set of all registers having the same clock i .

$\forall R_i^{Clk} \subseteq R_{cg}$ **do**

 partition $R_i^{Clk} = \{E_j\}_{j \in EN}$

 such that $E_j \subseteq R_i^{Clk}$ and E_j is a register bank with all registers having the same enable signal En ;

$\forall E_j$ **do**

Compute $Clkgated(E_j)$

Replace clock (Clk) of all the registers in E_j with the $Clkgated(E_j)$.

Generate the hardware in Verilog (containing the gated clock for the intended registers).

7.4 Results

We present the results of our approach on a number of designs as shown in Table 7.1. We performed our experiments using our enhancement of a C synthesis compiler [22], PowerTheater [84] for RTL power estimation, and TSMC 180nm libraries used in [135] for logic synthesis. Column 1 in the Table 8.1 shows the names of the designs experimented on, column 2 shows the power reduction achieved by applying clock-gating on every register using our approach, column 3, 4, 5 represents the clock-gating applied on register banks of sizes greater than or equal to 4, 8, 16 bits respectively. Table 8.1 shows that clock-gating on every register bit may not give the highest power savings. This is because our clock-gating introduces a *latch* and an *and* gate per enable signal, which also consume power. A trade-off point may be achieved by clock-gating register banks instead of individual registers. For example, in the caesar cipher case, almost 37% power reduction with respect to a non-clock-gated synthesis by individual register clock gating was achieved, but the savings further improved by clock-gating register bank of size 4 or more instead. NA in the Table 8.1 represents *not applicable*, because there was no register bank with size equal to or more than 16 bits.

Benchmark	% power reduction using clock-gating			
	Every bit	4 bit bank	8 bit bank	16 bit bank
Fibonacci	16.32	16.41	16.32	NA
Caesar	37.78	38.22	38.22	NA
FIR	36.1	36.2	36.2	NA
AES	30.74	30.74	29.83	26.42
GZIP	16.73	16.39	14.96	12.29

Table 7.1: Percentage power reduction as compare to the original design without clock-gating

We also compared the area and timing results of our approach against an automated RTL clock-gating technique using power compiler [85] shown in Table 7.2. In Table 7.2 we can see 4 cases where our approach found better power reduction opportunities, while impact on area and timing as compared to the lower level techniques is minor.

For example, for AES design, about 3% area saving was obtained while there was no impact on timing as compare to the results generated by power compiler. Our approach shows almost 25% extra power savings as compared to power compiler. This shows that our algorithm exploits the

Benchmark	% comparison of our results against RTL clock gating		
	Area	Timing	Power
Fibonacci	89.7	100	94.89
Caesar	112.10	100	91.03
FIR	100.2	100	100.2
AES	95.77	100	74.58
GZIP	100.36	100	100.36

Table 7.2: Comparison of our approach with automated RTL clock-gating

visibility of the registers and their control inputs especially for bank or registers. The purpose for this table is to provide a view to the reader that using our approach a very realistic design exploration is possible at the high-level in shorter time, without compromising quality, and some times with more power savings. In this table, we normalized the values from power compiler as 100.

7.4.1 Clock-gating Exploration at High-Level

In the Section 7.4, we discussed how a designer can filter the clock-gating register candidates and compared the results of our approach with respect to RTL clock-gating techniques. Here, we present a few example scenarios that can be easily explored at high-level. To explore the same scenario at the RTL, designer would require a lot more information from HLS to correctly clock-gate registers. It is expected that the person working at C source level would have more detailed knowledge about the architecture, and micro-architecture is produced by the synthesis tool. Hence relying on the RTL clock-gating techniques would increase the design time.

Let us consider Fibonacci example. As we know, in Fibonacci series, the number n is provided and corresponding series is emitted by the design/software. Now, depending upon the architecture/constraints on the design for a number n , design may take several clocks to produce the series. We noticed that n remains constant till design creates complete sequence. Generally value of n changes based on the input signal which suggest that data input is valid on the port (using an interface function). This suggests clock-gating this variable and other variables/scopes that are using this variable may lead to good power savings. Similarly in the case of AES design, it is a symmetric cipher and the *key* to encrypt/decrypt data remains same for several data iterations. Hence it make sense to just clock-gate variables saving value of key and the variables using value of key, as there might not be a lot of toggling happening on these registers. Now, applying such changes for clock-gating is very easy at source code level using our technique.

As discussed in Section 7.3, designer would require to put a directive to guide HLS to clock-gate the registers we are interested in. We observed almost 10% power savings for fibonacci example while in case of AES we could see little less than 5% power savings. In AES design there are a

few functions which can be separately clock-gated to see which function has the maximal impact on power consumption of the design. Such optimizations if applied at RTL would require a lot of effort especially in an HLS based design flow. In these cases effect on timing was negligible and we observed little saving in the area as well because of replacement of mux for a register bank with clock-gating logic.

7.5 Summary

In this chapter, we answered a few questions discussed earlier on enabling clock-gating in ANSI C description using clock-gating directives. This will guide HLS to generate power aware hardware. We first show how different directives can enable clock-gating at high-level starting from a variable, fixed bit-width, scope to a function. Secondly we provided an algorithm on how such directives would be taken care by the HLS tool. Finally, we provide example scenarios that an architect can easily see in the high-level description. Enabling such reductions at the RTL in HLS based design flow would require back and forth interaction between HLS and lower level tool, making the power aware exploration task very tedious and at the lower-level of abstraction.

We implemented the strategy into a commercial co-processor synthesis tool and applied it to synthesize a number of co-processors. We show that a significant power reduction is possible, sometimes beyond what a lower level automatic clock-gating tool would provide, without compromising area and timing. We also present various ways of controlling the clock-gating granularity. Precise control on various variables in ANSI C description helps in getting rid of redundant clocking of the registers of the design for different granularities.

Chapter 8

Model-Checking to exploit Sequential Clock-Gating

Dynamic power reduction techniques such as sequential clock-gating aim at eliminating inconsequential computation and clock-toggles of the registers. Usually sequential clock-gating opportunities are discovered manually based on certain characteristics of a design (e.g. pipelining). Since manual addition of sequential gating circuitry might change the functionality of the design, sequential equivalence checking is needed after such changes. Tools for sequential equivalence checking are expensive, and based on recent technologies. Therefore, it is desirable to automate the discovery of sequential clock-gating opportunities using already existing and proven technologies such as model checking and thereby a priori proving that the changes will not affect the required functionality. Model Checking Based Sequential Clock Gating (MCBCG) method formally proves particular sequential dependencies of registers on other registers and logic, thus sequentially gating such registers will not require further validation. An automation scheme for MCBCG methodology is also proposed in this chapter. Our experiments show up to 30% more savings than the traditional (combinational) clock-gating based power reduction techniques. We further experimented the approach to find out the savings opportunities exist in HLS tools. This further suggests that at high-level finding such opportunities might be very beneficial for HLS based design flow.

8.1 Introduction

Register-power is generally one of the biggest contributors [86] to the total dynamic power consumption. To obtain maximal power reduction of the design, it is advisable to reduce the register power of the design. Clock-gating [86] is one of the techniques which has been extensively used to reduce the dynamic power consumption. The main idea behind clock-gating is to reduce the unnecessary toggles of registers when the register update is not required or its value is irrelevant to the computation. In the absence of clock-gating, at every clock cycle, each register gets updated

even when it remains unchanged in value. One can utilize the enable/control signal for a register to stop the clock from making the register toggle. The computation of the enable/control signal to indicate that the clock needs to be stopped can be done in two ways. When the information is temporally localized and only based on the current cycle, it is known as *combinational clock gating* described in the Chapter 7. However, more benefits can be obtained if it is known that the register value is not going to be used in the future cycles, or if certain events from some previous cycles have indicated that register update is unnecessary. Utilizing such information to gate a register's clock is commonly known as *sequential clock-gating* [127].

Keeping such information across clock cycles requires relevant circuitry to propagate the information temporally into future cycles. Also, knowing whether register's value update is irrelevant to future computations depends on the dynamic relationship between the register and other events during the execution. If this information can be inferred from the structure of the design, then designers can manually insert the relevant gating circuitry. In pipelined designs these kind of inferences are often not that difficult [126]. It is hard, when the structure of the design does not provide an obvious opportunity, or when the designer may suspect that such relationships across clock cycles exist, but cannot be sure due to the complexity of the design. Tracing the fanin or fanout cones of registers can be often helpful, but requires automated analysis [92].

In any case, such optimizations that can affect the functionality of the design over the span of time/clocks require considerable effort on verification, usually best done by checking sequential equivalence. In traditional methodologies, first these optimization techniques are applied manually and then equivalence checking is performed. However, sequential equivalence checking technologies are new, and hence using a more proven technology would be more desirable. Keeping this in mind, we propose a novel methodology called as *Model Checking Based Sequential Clock-Gating (MCBCG)*. MCBCG utilizes model-checking to apriori prove that sequential clock gating opportunity exists for a register r , based on other registers, whose activities in past cycles would always indicate irrelevance of updation of this register in a future cycle. Since we are using model checking, we can try to prove this for any arbitrary set of registers if we wish, and hence we do not depend on the special structures such as pipelines. Of course, designer should apply this in a more informed manner than arbitrarily. Note that most of the current approaches we found in the literature, implement sequential clock-gating based on the architecture of the design such as pipelined designs. In contrast, our approach and the one presented in [92] can be applied to any RTL model of the design. Secondly, since we already prove that sequentially clock gating the registers found using this approach would not affect the computation, we do not need expensive sequential equivalence checking afterwards.

Main contributions of this work for the presented methodology are following:

- A novel methodology to investigate the sequential clock-gating opportunities in hardware designs.
- Demonstration of some temporal properties that can be used for model checking to infer these opportunities.

- Experimental results showing opportunities to save substantial power over combinational clock gating based power reduction.
- Extending the HLS framework to generate sequentially clock-gated RTL.

8.2 Our Approach and sample properties

8.2.1 An Illustrative Example

We briefly discuss here a few possible scenarios that will help in understanding the rationale for our approach. Let us assume two register a , b and b is dependent on a . In the Snippet 2 $in1$, $in2$ and en can be considered as input. Snippet 3 shows how the values of the registers change. We can see from this example that if value of a is not utilized elsewhere then a is not required to do the unnecessary computations.

Snippet 2 Representation of two registers p , r

```
.....
next(a) = in1 + in2;
.....
if(en)
  next(b) = a + 5;
```

Snippet 3 Clock by clock view of the changes in value of various registers

```
Initially a = 0; b = 0;
cycle 1 in1 = 5, in2 = 5; en = 0; a = 0, b = 0
// results are updated a cycle later
cycle 2 in1 = 7, in2 = 6; en = 1; a = 10, b = 0
// b gets updated because en was true
cycle 3 in1 = 8, in2 = 7; en = 0; a = 13, b = 15
// en was false, b won't change
cycle 4 in1 = 7, in2 = 2; en = 0; a = 15, b = 15
```

Similarly let us consider the example where a can be clock-gated while b cannot be but they have a relationship as shown earlier. Snippet 4, 5 show register b gets unnecessary clock-toggles since there is no value change but still clock is supplied. In this case after applying sequential clock-gating, functionally design would not exhibit any change but dynamic power still can be reduced. While in the earlier example, if value of register a is used in other parts of the design then it is necessary to check that value change in the register a should cause a value change on some other register after a clock-cycle. If it is the case we can't do a sequential clock-gating. While for the second case we can obviously apply sequential clock-gating because it will not affect the design functionality.

Snippet 4 Example of operations on registers a, b

```

.....
if(en)
    next(a) = in1 + in2;
.....
next(b) = a + 5;

```

Snippet 5 Clock by clock view of the changes in value of various registers

```

Initially a = 0; b = 0;
cycle 1 in1 = 5, in2 = 5; en = 0; a = 0, b = 0
// results are updated a cycle later
cycle 2 in1 = 7, in2 = 6; en = 1; a = 0, b = 5
// en was false, a won't change
cycle 3 in1 = 8, in2 = 7; en = 0; a = 13, b = 5
// a gets updated because en was true
cycle 4 in1 = 7, in2 = 2; en = 0; a = 13, b = 18
// b changed a cycle later, a did not change because
// en was false

```

8.2.2 Our Approach

As discussed in Section 3.8.2, sequential clock-gating can be applied if two registers are dependent on each other in certain ways. One way to find these dependencies is to go through the structure of the design and find out such relationships (explained later in this chapter). This requires traversal of the netlist. However, in some cases, for example, in the case of Snippet 1, it cannot be inferred through such traversal. Fortunately, by utilizing a model-checker we can find such intra-register relationships across clock boundaries, both in the case of structurally inferable dependences, and data dependent relations. Based on the discussion in the Section 3.8.2, 8.2.1, we can map this relationship to a property of the design. Such a property is basically run as a query to the database of the design, asking if a relationship between registers exist or not. Once this property is verified by the model checker and the answer to the query is 'yes'. Then we know how to exploit this opportunity to utilize sequential clock-gating in the design.

Let us assume we have the implementation model and a formal model (for model checking) for a design under consideration. Some design groups in industry/academia utilize tools that can convert the RTL description to the formal models for model-checking purpose. The formal model must capture all the registers and register transfer logics from the implementation model. Now we will see some sample properties that one can write on the formal model. If these properties are verified then we can use the implied relationship for clock-gating the registers, where redundant activities are happening. Let us consider two arbitrary registers r , p in the design. Suppose the value of p changes at an arbitrary clock cycle t , and the value of r changes at clock cycle $t + 1$. We want to verify if r changes at cycle t , then p must have changed at cycle $t - 1$. In the formal model, we add

three extra state variables for register r and p . Let us name them as p_pre2 , p_pre and r_pre with the same types as p and r . Initialize p_pre2 , p_pre with the initial value of p and r_pre with the initial value of r . Then make sure to add the transitions for p_pre2 , p_pre and r_pre . We will show the transitions using the syntax of Cadence SMV [136] for encoding value changes of registers. In Cadence SMV, the transition of register s is written as shown in Snippet 6.

Snippet 6 Registered value representation in Cadence SMV

```
next(s) = expression;
```

The expression may have if-then-else to encode the multiplexer logic. When the formal model has registers r and p , then it will look as shown in Snippet 7.

Snippet 7 Representation of two registers p, r

```
next(p) = expression1;
next(r) = expression2;
```

Change the model shown in Snippet 7 to the one shown in the Listing 8.

Snippet 8 Changes in the model to verify the properties

```
next(p_pre2) = p_pre;
next(p_pre) = p;
next(p) = expression1;
next(r_pre) = r;
next(r) = expression2;
```

Therefore at cycle $t + 1$ p_pre , r_pre contains the value of p and r from cycle t . So if $(p_pre \sim= p)$ is true that means at cycle $t + 1$ value of p changed. Similarly, if $(p_pre2 \sim= p_pre)$ at cycle $t + 1$, that means at cycle t p changed from its previous value at cycle $t - 1$. Let us consider the following two properties:

1. At cycle t , if r changes, p must have changed in $t - 1$. Verify the property shown below. If this property is true we can see if p can be clock-gated then r is a potential candidate for clock-gating in the next cycle.

$$G((r_pre \sim= r) \rightarrow (p_pre \sim= p_pre2));$$

2. At cycle $t - 1$, if p changes, then r must change at cycle t . Verify the property shown below. If this property is true then we have to check if the value of p is not used by some other register and r can be clock-gated then p is a potential candidate of sequential clock-gating.

$$G((p_pre \sim = p_pre2) \rightarrow (r_pre \sim = r));$$

By checking these and other similar properties on the formal model, we can be sure that de-asserting enable of p at any cycle t , would imply that r will not change its value at $t + 1$. Hence we can apply sequential clock-gating for the related registers r and p . G is the linear-time temporal logic (LTL) operator used in the example listings to represent always. We can write properties till n clock cycles, where n is the number of clock stages we are interested in. We call this as window size. Practically this number won't be too big; because every time n increases we introduce extra registers and clock gating logic to reduce the power consumption, which may affect area largely.

8.3 Our Proposed Methodology

As shown in Figure 8.1, MCBCG methodology starts with the formal model of the design. Once such a model is available, we can start writing the properties of the design as discussed in the Section 8.2.2. Then pass the model along with the properties (queries) to be verified to a model checker. We have used Cadence SMV for our analysis but other model checker such as VIS [137], ABC [138] can also be utilized for the similar analysis. VIS, ABC can directly read in the RTL model and facilitate user to write properties on the implementation model. We use SMV because of the familiarity with the tool.

Once the properties are verified by the model checker then based on the relationship between the registers apply the power reduction specific changes. These changes requires some changes in the RTL verilog and very simple to implement at RTL. In this paper, we applied latch based clock-gating scheme. The advantage of such a scheme is that combinational glitches can be avoided in the clock port of the register. If property is false then apply the similar property to other set of registers. Once RTL model is ready with the applicable changes for sequential clock-gating, utilize this model for power estimation. One should note that we apply clock-gating using power compiler in both the models but the difference is the enabling condition for the registers which can be sequentially clock-gated in our approach. This whole process can be automated by first finding out the structural relationship between the register pair and then utilizing our technique to find out if register pair are good candidates for clock-gating based scheme. Finally once the power optimized RTL is ready, we validate whether our approach provide any further reduction or not.

These reductions may show very good results for one set of stimulus and in other cases it may not. It is very important to choose right stimulus for the analysis. To measure the differences with respect to the original model, power consumption for both the models is estimated using same technology libraries and the same test vectors. Finally the difference in area, timing and power can be measured after completing the design synthesis, etc.

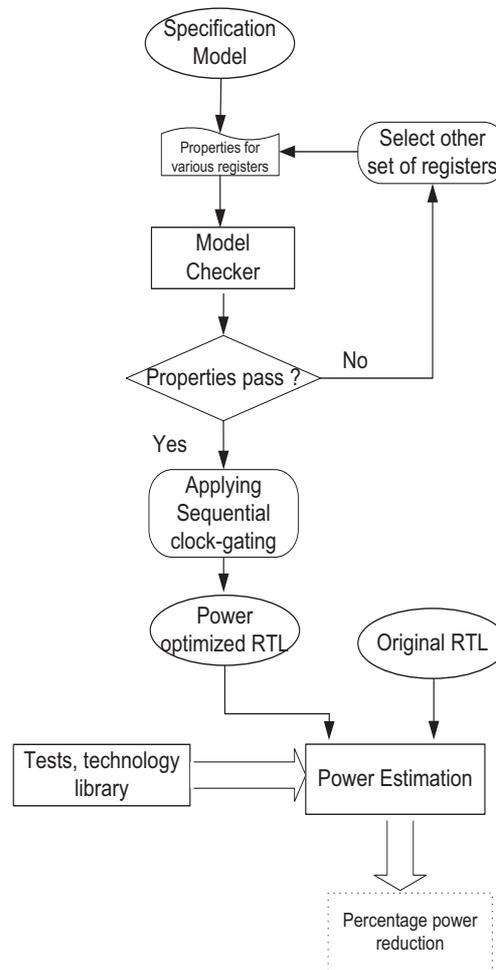


Figure 8.1: Our proposed methodology for applying sequential clock-gating at RTL

8.3.1 Changes applied on the RTL model

We propose minimal changes for the RTL modeling stage. The code changes are shown using Verilog hardware description language (hdl), similar changes would be required with VHDL or any other hdl. After including these changes tools such as power compiler can directly be utilized. Such a facility will help the designer to stay in the current design flow. Listing 8.1, 8.2 shows the changes that can be applied to the code. In Listing 8.1, we show how two different registers q , $q1$ may look like in Verilog. After doing the analysis using MCBCG methodology, we know that these two registers are the candidates for sequential clock-gating. We utilize the enabling logic of the two registers. One should note that a change in value of $en1$ a cycle before is captured first in $en1_tmp$ register. Finally we utilize the combinational logic of the two register's i.e. $en1_tmp$, $en2$ and store it in a wire (using assign statement in Verilog) and change the $en2$ with $en2_tmp$.

All code changes are illustrated in the Listing 8.2. We can also include the clock-gating specific changes directly in the RTL model but for our experiments we utilized power compiler to apply combinational clock-gating.

Listing 8.1: Verilog code segment before applying MCBCG technique

```

1 //register q with enable en1
2 always@(posedge clk or negedge rst)
3 begin
4 // Reset logic skipped
5   if (en1)
6     q <= r;
7   .....
8
9 end
10
11 //register q1 with enable en2
12 always@(posedge clk or negedge rst)
13 begin
14 // Reset logic skipped
15   if (en2)
16     q1 <= r1;
17   .....
18
19 end

```

Listing 8.2: Verilog code changes after applying MCBCG technique

```

1 //register q with enable en1
2 always@(posedge clk or negedge rst)
3 begin
4 // Reset logic skipped
5   if (en1)
6     q <= r;
7   .....
8
9 end
10
11 always@(posedge clk or negedge rst)
12 begin
13 // Reset logic skipped
14 // this will delay the en1 by a clock
15 // and latch it in a register
16   en1_tmp <= en;
17   .....
18
19 end
20
21 // New enable dependent on the previous
22 // enable condition
23 assign en2_tmp = en1_tmp and en2;
24
25 //register q1 with enable en2
26 always@(posedge clk or negedge rst)
27 begin
28 // Reset logic skipped
29   if (en2_tmp)
30     q1 <= r1;
31   .....
32
33 end

```

8.3.2 Automation of the flow

Algorithm 9 present a scheme to automate the methodology. To describe the algorithm we need to first fix a few notations as follows: Let R be the set of registers in a design D ; W be the window size (different stages of the registers to find the clock-gating opportunities), T be the set of register bank of the same type, and P be the set of properties for every register pair.

Algorithm 9 Automation of *sequential clock-gating of registers*

```

Procedure SeqCG{
   $\forall w \in W$  do
    execute Procedure WindowCG

}
Procedure WindowCG{
  Find out all register bank of same type (bitsize)  $t$ 
   $\forall t \in T$  do
    collect all register pairs
    Find out the relationship between the registers using properties discussed in Section 8.2.2
     $\forall p \in P$  do
      if property holds
        For both the cases discussed in Section 8.2.2 find out candidate register pairs for sequential
        clock-gating
        collect all the unique register pairs in a list

  Apply the sequential clock-gating optimization and change the verilog as discussed in Section
  8.3.1 for each register pair of the list

}

```

In the above procedure based on the window size, every property is verified for register bank of the same type. Doing such a task requires a knowledge of the register bank and would require addition of extra register stages in the formal model. This task can be incrementally performed for all the registers in the design. We have implemented the procedure using c script. Once property is verified on a register bank we save this information in a separate list, from there we finally post process the original RTL based on the steps discussed in Section 8.3.1.

8.4 Results and Conclusions

Our investigation in this work show promising results. We tried the approach on two sample designs FIR filter and 3-stage pipelined packet processing design. We could achieve more power

savings than combinational clock-gating using power compiler. Table 1 shows the comparison of our approach against power compiler. Our approach shows that almost 30% more power savings can be obtained. We used VCS tool for RTL simulation [131], power compiler [85] for clock-gating and power estimation, and 180nm libraries from [135].

Table 8.1: Percentage power reduction summary on the tested Designs

Design	Total Power Consumption	Power Consumption after c.g.	Power Consumption after seq. c.g.
FIR	1106uW	522uW	426uW
3-stage pipe	301uW	239uW	147uW

8.4.1 Extending the framework for HLS

This chapter presents application of model checking to find out a relationship between registers. However, model-checker may not always be available for HLS based framework. In this section, we briefly present the steps used for structural traversal for sequential clock-gating opportunities. As discussed earlier, sequential clock-gating technique relies on the fact that enable condition of one register a cycle or two cycles before can be utilized in earlier/later stages of the design. These steps are captured in Listing 10 to perform structural traversal of the intermediate Control/Data Dependency Graph (CDFG).

Snippet 10 Steps for finding sequential clock-gating opportunities at high-level

Step1: Obtain the dependency graph of the design and identify the registers of size atleast 4-bits, 8-bits (or bank size specified by the user). In sequential clock-gating an extra register is introduced for sequ, so this scheme needs to be implemented for a register-bank.

Step2: Save every register and its dependency on other registers and mark the dependency as levels while traversing the control flow.

Step3: foreach (level)

 Apply Step 4;

 Step 4:

 a-) From the dependency of every register from step2, find out the relationship of every register for the future stages

 b-) If the register in the future stages are dependent on more than one current state registers than use a combinational logic of all the enable conditions and delay it to clock-gate future register stage

 c-) else delay the gating logic by a clock using a register

 end foreach

Step 5: Generate the Verilog RTL of the new design

Results

The steps presented in the earlier section show the structural traversal of registers on the CDFG of the design. As discussed earlier in the Chapter 2 model-checking based approach can help in finding some of the behavioral relationships which structural approach may not find. In our experiments, we first verified the generated designs for all the benchmarks. We used simulation based approach where correct functionality of the design was verified with and without sequential clock-gating for the same test-vectors. We observed that every benchmark has not shown the opportunity to sequentially clock-gate because of the nature of the design. In our experiments, we found that designs having a pipelining nature or direct dependency between two registers have shown some power savings and could be optimized further. Table 8.2 captures the results obtained.

Table 8.2: Percentage power reduction summary at high-level

Design	Total Power Consumption	Power Consumption after c.g.	Power Consumption after seq. c.g.
AES	123 mW	90.6 mW	89.2 mW
3-stage pipe	1.11 mW	1.16mW	.752mW
FIR	2.9mW	1.85mW	1.79mW

Chapter 9

System Level Simulation Guided Approach for Clock-gating

Clock-gating is a well known technique to reduce dynamic power consumption of a hardware design. In any clock-gating based power reduction flow, automatic selection of appropriate registers and/or register banks is extremely time-consuming because power analysis is performed at the RTL or lower level. In a high-level synthesis (HLS) based design flow, to achieve faster design closure, one must be able to decide the appropriate set of registers to clock gate even before generating RTL. System-level simulations are known to provide faster simulation, yet there is no solution, which utilizes system-level simulation to provide guidance to HLS to create clock-gated RTL. Since predicting power reduction at higher levels of abstraction is difficult due to the dependence of power on physical details, an accurate and efficient *relative* power reduction model is required. In this chapter, we propose a novel system-level design methodology, which utilizes a ‘relative power reduction model’ that can help in predicting the impact of clock-gating on each register/bank quickly and accurately, by simulating the design at a cycle accurate transaction-level. As a result, our approach can automatically find the appropriate registers to clock-gate, guided by the system-level simulation.

9.1 Introduction

Currently, most of the HLS based design flows include power reduction pass through RTL or gate-level tools such as Power compiler [85], PowerTheater [84], etc. However, these tools rely on designers’ insertion of clock-gating directives for the registers at the RTL. In other words, it is the designer’s responsibility to select if all registers should be clock gated, or registers belonging to specific modules or specific parts of the design should be clock gated. If a register is updated very frequently with new values, inserting clock gating logic causes unnecessary overhead with adverse effect on power consumption. The other impact is on the timing closure. If unnecessary

clock-gating is done, more timing issues will arise at the lower levels of abstraction. Clock-gating should be applied only if a significant amount (i.e. above a certain threshold) of power savings for every register can be observed [139].

A few aggressive techniques to perform clock-gating utilize observability don't care (ODC) to find signals whose value change is don't care, and hence can be gated [92], [93]. However, the power consumption by additional circuitry can overwhelm the savings if such a technique is applied for every register. Most recent work such as [92], [93] suggest the registers for which more power saving opportunities may exist but cannot determine from the simulation if clock-gating based reduction will provide savings for every register/register-bank. The savings are presented on an average for the whole design and at the netlist level. Furthermore, to be effective at high-level, this information should be available even before RTL is generated.

In this chapter, we show that system-level simulation can be utilized to generate clock-gated RTL using HLS. To achieve this, we propose a novel power reduction model, which measures the activity of a register value change at cycle-accurate transaction level and then guides the HLS to generate the clock-gated RTL. This approach helps in finding whether each register would save or cost power if an aggressive power reduction approach is used in the design flow. We have benchmarked the accuracy of our estimation of power savings on the industry strength co-processor design blocks, and found that the accuracy of our estimation is encouraging (approx. 7%) and speedup to perform power analysis is in 2-3 orders of magnitude. We also show that for some designs, upto 45% of the registers in the design may not require clock-gating.

9.2 Power reduction model

We present power reduction model for clock-gating in this section. As we know the idea behind clock-gating is to reduce the unnecessary clock toggles. We introduce two power models in this section and the reason comes from the availability of technology libraries in the early design flow. If early in the design flow, good technology libraries are not available and yet we want to measure the impact of power reduction on the generated RTL. For such a scenario, extraction of effective clock updates from simulation of clock-gated design is required. Section 9.2.1 provides detailed analysis of this approach. A lot of times, during high-level synthesis, technology library information is available. In that case we can enhance the power model with such information and get a better visibility on the reduction that can be achieved. In that regard, we require information about integrated clock-gating cell and multiplexers. Section 9.2.2 provides details of technology dependent power model.

9.2.1 Activity based power reduction model

Our relative power model depends on 'savings duty cycle' estimation. In our HLS flow, cycle accurate transaction level (CATL) and register transfer level (RTL) models are generated from the

same control flow graph (CFG) hence the correct information on enables and registers are available while generating the CATL model. Let us assume that the enabling condition for a register update is denoted by en . To extract the useful update of a register, we insert a counter for each register of the design in the CATL model. Snippet 11 shows a very simple code after including a counter for every ‘enable’ of a register in the design. This counter helps in capturing all the effective updates happening to a register/register-bank (in this case, register q). This model can be simulated in a SystemC simulation environment.

Snippet 11 Counter inserted in the generated SystemC code

```

if(en){
    q = some computation;
    // counter for q will be inserted here
    counter_en = counter_en + 1;
}

```

Let us divide the total dynamic energy (E_{total}) for a register into two parts E_{useful} and E_{waste} (represented in the Equation 9.1). Let R be the set of registers in a design D ; let N_r represent the inserted counter for ‘enable’ of each multiplexer in front of a register r and N_{clk} a counter for the clock. E_{useful} can be represented as shown in the Equation 9.2. Similarly, E_{waste} can be represented as shown in the Equation 9.3. k is a constant of proportionality dependent on area, capacitance etc.

$$E_{total} = E_{useful} + E_{waste} \quad (9.1)$$

$$E_{useful} = k * (N_r/N_{clk}) * V^2 * f_{clk} * T \quad (9.2)$$

$$E_{waste} = k * (1 - (N_r/N_{clk})) * V^2 * f_{clk} * T \quad (9.3)$$

Thus, $(1 - (N_r/N_{clk}))$ is the ‘savings duty cycle’ for the register r . This power model can help us in evaluating the impact of clock-gating on power savings. If the information on technology libraries and clock-gating cell is not available, designer can provide a threshold to select the clock-gating candidates as shown in Equation 9.4. This inequality checks if savings are above certain limit ($a_{threshold}$). If it is, then clock-gate the register. Since the number of registers before and after clock-gating are same, such a relative power model will help in finding out the power reduction opportunities even before applying the clock-gating changes in the RTL. $a_{threshold}$ is a fraction between 0 and 1.

$$(1 - (N_r/N_{clk}) - a_{threshold}) \geq 0 \quad (9.4)$$

9.2.2 Power reduction model with technology specific information

Roughly, the power model presented in earlier section shows how much clock toggle savings exist in a design. In this section, we propose a more detailed power model of the design to aid the analysis. It is well known that when clock-gating is applied to a design then only modification in the design comes by removing a multiplexer from the feedback path of a register; also a clock with gating logic is supplied to the clock port. Our power model captures these changes. We utilize an integrated clock gating cell (icgc, generally the logic contains an *and* gate with a latch) to provide the gated clock. For the sake of clarity, we limit our discussion to single bit register bank but for a register bank of size N and icgc of fanout strength p , a few minor modifications would be required. As we know that power consumption is dependent on capacitance, which further depends upon the area; and the area information of these cells can be extracted from the technology library to improve the accuracy of the power model.

In our case, we used standard cell library from TSMC and obtained area information (details on our experimental setup are provided in the results section). Generally, all the standard cell technology libraries contain this information. One should note that we do not require information for all the cells rather a few cells, which are intended to be utilized for clock-gating purposes only. Equation 9.7 shows the coefficient, which is computed for various simulation vectors, if this coefficient is positive, tool will clock-gate the register, else it will not. The negative number shows that there would not be any power savings using clock-gating for a particular register. This way tool will automatically infer the information of registers to be clock-gated directly after the simulation.

Let us represent the energy of a register, mux, and icgc as E_r , E_{mux} , and E_{icgc} respectively. Frequency of updates for register and mux can be denoted as N_r and for icgc N_{clk} . Dynamic energy of a register before and after applying clock-gating can be denoted by E_{r1} , E_{r2} as shown in the Equations 9.5, 9.6.

$$E_{r1} \propto A_{reg} * N_{clk} + A_{mux} * N_r \quad (9.5)$$

$$E_{r2} \propto A_{reg} * N_r + A_{icgc} * N_{clk} \quad (9.6)$$

Difference in the energy values (which represents energy savings/wastage before and after clock-gating) with respect to energy of a register before clock-gating can be represented as shown in the Equation 9.7:

$$p_c = \frac{E_{r1} - E_{r2}}{E_{r1}} = 1 - \left(\frac{N_r \left(1 - \frac{A_{mux}}{A_{reg}} \right) + N_{clk} * \frac{A_{icgc}}{A_{reg}}}{N_{clk}} \right) \quad (9.7)$$

Equation 9.7 clearly shows p_c will not always be positive but it depends on a lot of factors. This

equation also shows that introduction of an icgc may cause power loss because it is always fed with the 100% activity of the clock port. This is also the point in the clock-tree from where the clock is branched and duty cycle is reduced. On the other hand, there will be some savings because of removal of muxes. In case of register banks the power savings would be higher because an icgc would be introduced for the bank but its drive strength would be computed based on the size of the register bank. System-level simulation can make the process of computing N_r , N_{clk} very fast. Values of A_{mux} , A_{reg} , A_{icgc} will change according to the standard cell library used in the flow. For N bit register bank and p bit fanout of an icgc the power model will look as shown in the Equation 9.8. After including the threshold provided by the designer, relative savings look as shown in the Equation 9.9.

$$1 - \left(\frac{\left(N_r + \left(\frac{A_{icgc}}{A_{reg}} \right) * N_{clk} (1/N + 1/p) - N_r \left(\frac{A_{mux}}{A_{reg}} \right) \right)}{N_{clk}} \right) \geq 0 \quad (9.8)$$

$$p_c - a_{threshold} \geq 0 \quad (9.9)$$

9.3 Rationale for Our Approach

Let us consider that for a design (at RTL) R and for vector i simulation takes T_i^R time. Let us represent the clock-gated RTL as R_c and the time to simulate vector i as T_i^{Rc} . Let C represent different configurations such as fine grain clock-gating for every register, register-bank, or to a particular module (coarse grain), etc. To check the efficacy of power reduction technique on the design for n different simulation vectors, the power estimation time required T_E^c can be calculated as per the simple procedure shown in the Procedure PowerEstTime.

```

Procedure PowerEstTime() {


---


   $\forall c \in C$  do
     $\forall i \in n$  do
      Simulate RTL and Estimate power for the design R,  $R^c$ 
    end do
  end do
}

```

Equation 9.10 represents the total time to simulate RTL with and without clock-gating for each simulation vector as per the procedure *PowerEstTime* discussed above. T_i^R , $T_{E_i}^R$ represent the time taken in performing simulation and power analysis for the design R before applying clock-gating, and T_i^{Rc} , $T_{E_i}^{Rc}$ represents the same after applying clock-gating.

$$T_E = \sum_{c \in C} \sum_{i=1}^n T_{E_i}^c = \sum_{c \in C} \sum_{i=1}^n (T_i^R + T_{E_i}^R + T_i^{Rc} + T_{E_i}^{Rc}) \quad (9.10)$$

Let us assume, there exists a system-level approach that can help in measuring the power savings using high-level power model (such as the one discussed in the last section) and analysis time remains close to system-level simulation and can be represented as T_i^s for vector i and total system-level simulation time can be represented as T^s . As we know, system-level simulations are faster than RTL, we can assume:

$$T_i^s \ll T_i^R \quad (9.11)$$

Experimentally, we found that cycle accurate transaction level (CATL) simulation is at least 10-15 times faster than corresponding RTL simulation. This can be attributed to quite a few reasons, for example, CATL simulation is applied to data-types such as char, integer as compared to bit-level data-types in RTL simulations. Also one can use functions, pointers, which makes the simulation faster than the RTL simulation. Relationship between system level and RTL simulation time can be represented as shown in the Equation 9.12.

$$\frac{\sum_i T_i^R}{\sum_i T_i^s} = \mu \quad (9.12)$$

As per the discussion above:

$$\mu \gg 1; \quad (9.13)$$

The ratio of RTL and system-level power analysis time can be represented as:

$$\frac{T_E}{T^s} = \frac{(\sum_i (T_i^R + T_{E_i}^R) + \sum_c \sum_i T_i^{Rc} + T_{E_i}^{Rc})}{\sum_i T_i^R / \mu} \quad (9.14)$$

$$\mu + \frac{\mu((\sum_i T_{E_i}^R + \sum_{c \in C} \sum_i (T_i^{Rc} + T_{E_i}^{Rc})))}{\sum_i T_i^R} \quad (9.15)$$

At the RTL, power estimation tool would need to extract activity from the simulation dump and map it to the design for getting power numbers. This is one of the most time consuming processes of a power estimation tool. Based on our experience, we find that all these factors in an RTL power analysis tool are almost 10-100 times slower than the system-level simulation [28]. This would make the speedup in the range of 2-3 orders of magnitude. We experimentally verify this claim in the results section. The rationale to use system-level approach can dramatically improve the design cycle time from the low power design perspective.

9.4 Our Methodology

A typical HLS based design flow involves generation of RTL, insertion of clock gating based on user inputs at the RT level, followed by logic synthesis, and then a pass through power estimation tools. These estimation tools require design information in verilog/vhdl, simulation dump in vcd/fsdb formats, and technology library information (.lib), etc. Power estimation is very time consuming mainly because RTL/gate-level simulation is slow and the generated vcd/fsdb is large and needs re-processing for power estimation purposes. After this time consuming single pass, if power closure is not achieved, another time consuming pass is made, and such passes are repeated until power closure is achieved. Figure 9.1 illustrates the iterative process in a standard HLS enabled design flow.

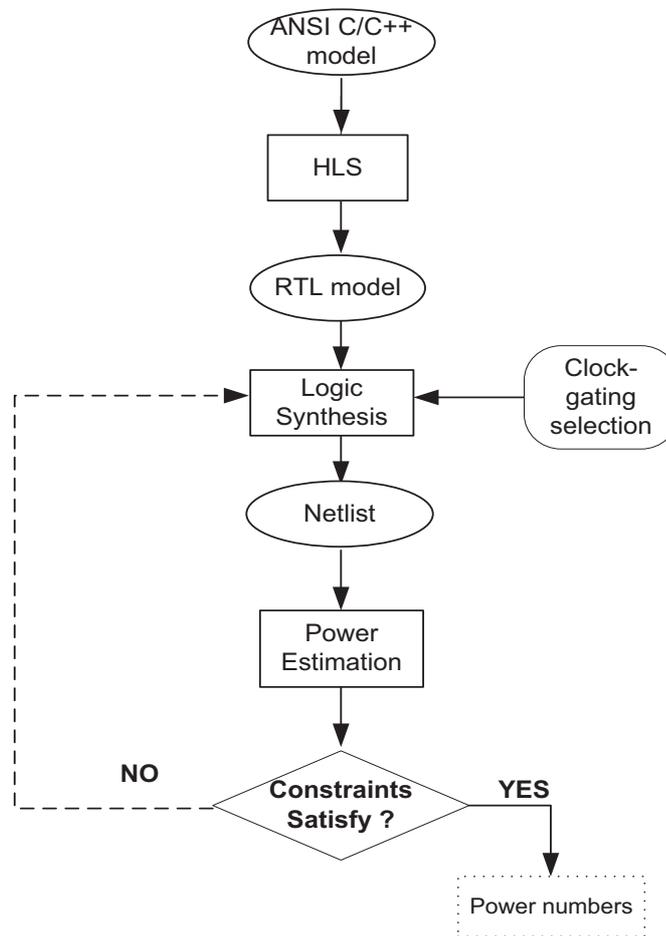


Figure 9.1: Traditional clock-gating methodology for HLS based design flow

We discuss in detail some of the distinct features of the methodology proposed in this chapter. These features include integration of tool generated SystemC model with un-timed TLM model to

achieve very high speed analysis. Second important part is the support of clock-gating at ANSI C based description itself. Section 9.4.1 provides a brief overview of our interconnection framework. Chapter 7 provide details on how to facilitate clock-gating from the C description itself. The advantage of providing the control from C description comes by controlling granularities of clock-gating, which RTL clock-gating tools can not find by netlist traversal.

In today's design flow, designer applies clock-gating based on their experience and there is no way in which he/she can measure individual register's power savings in the design. One possible solution is to generate power report for each register's power consumption before and after applying clock-gating and then make a decision on which register to clock-gate. At netlist level this is very detailed and time consuming. In combinational clock-gating approach, this decision is approximated by applying clock-gating on registers of size above particular bit-width (e.g. 4 or 8 bit banks). Apart for this approximation another important step is to find out the common enables for different registers. This helps in putting a clock-gating cell near to the root of clock tree. In an HLS flow, netlist traversal may not always find out the common enable. For example, the same conditional assignment may exist in two different processes at HLS, which at RTL may represent two different state machine. If this common case can be extracted during HLS then even better clock-gated RTL can be generated. The algorithm presented in Chapter 7 takes care of such considerations.

The methodology presented in this chapter eliminates iterations requiring measurements of power consumption of each register before and after clock-gating is applied. It also shows, how the selection of appropriate register(s) would be performed *automatically* using simulation performed at transaction level. The extraction of enable provides two distinct advantages with respect to existing methodologies 1-) un-timed simulation can help in guiding the reduction results, 2-) the reduction model specifically targeted for clock-gating gives an idea to the designer to find the possibilities of gating the logic. The proposed power model with technology specific parameters helps in finding this information. Another important point to note is that we can easily measure the impact of clock-gating, which helps us in getting the relative, yet accurate prediction on the savings opportunities.

Setup for our methodology includes C2R HLS tool [140], which takes ANSI C specification as input along with some constraints to generate RTL. These constraints help in finding an appropriate architecture for optimal timing and area. As illustrated in the Figure 9.2, our methodology utilizes the HLS tool in two passes. In the first pass, tool generates cycle accurate transaction level (CATL) model in SystemC. In the generated SystemC simulation model we embed the 'relative' power model as discussed in the Section 9.2. This SystemC CATL model is simulated and attached to the untimed transaction level test vectors using generic payload models and a report is generated. For every register as discussed in the Section 9.2, Equation 9.4 and Equation 9.9 are evaluated. If the inequality shown in these equations is not satisfied then a register is not a candidate for clock-gating. After this step, all the registers requiring the same enable signal are sorted.

Once the selection of registers is done, the HLS tool is utilized in the second pass to generate the RTL. For this purpose, we insert the clock-gating directives corresponding to the selected registers

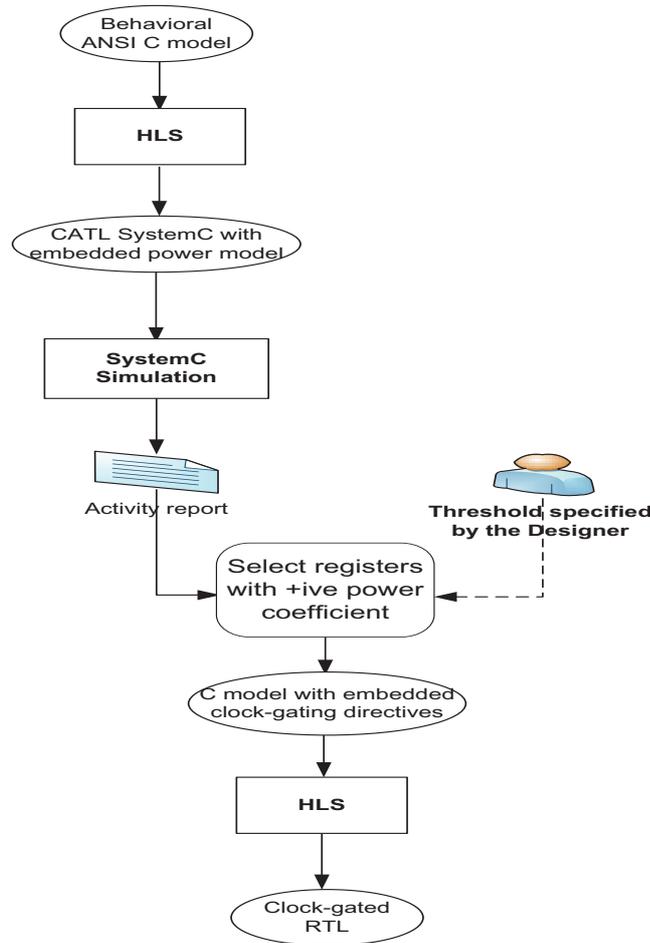


Figure 9.2: Pictorial Representation of Our Methodology

in the source code. These directives are macros such as `clock_gating(ON)` and not visible during the simulation performed in native C environment. Methodology described in the Chapter 7, shows how directives can be included in the C source code to handle various granularities of clock-gating. This auto-generated RTL contains clock-gating for the selected registers and will provide dynamic power savings as predicted. We have experimentally validated the accuracy of these predictions as reported later in this chapter. In our methodology there is an optional input from the designer. Designer can increase the threshold if he/she suspects that a typical application scenario would exhibit lesser ‘savings duty cycle’. Registers showing larger ‘savings duty cycle’ than the threshold value are clock-gated. By default, this threshold is set to zero.

One such scenario where designer may want to put a high threshold value is the presence of a lot of local enables in a block with very high duty cycle. It suggests that locally there might not be very high savings existing in the design by applying clock-gating. Alternatively, it gives a handle to the designer to perform a what-if analysis on power savings to choose a coarse or fine grain

clock-gating. Generally in block level clock gating the duty cycle of clock is reduced, now further inserting an icgc in clock tree path might not be desirable unless the local gating provides sufficient savings.

9.4.1 TLM interface for the generated CATL SystemC Model

Figure 9.3 shows how the CATL is attached to other TLM using simple sockets and generic payload. Generic payload is a data format defined in TLM2.0 standard for transport objects [141]. It contains data, address information, and typical attributes for memory, etc. It also has a response status field, so that recipient and caller can communicate status information. Our environment provides a communication medium between two parts, first is the test/application development platform where different application or test generators are running. Second is the CATL SystemC model of the design (DUT). Both test and design side works as recipient and caller while sending input and output for the design. For sockets, they are termed as initiator and target. Initiator initiates the transaction. We can see both the test and design having initiator and target sockets. Initiator from the test is connected to the target for the design (generated from the HLS tool). The ports on the cycle accurate model are scalar while the communication through the generic payload is using vectored data.

To ease the connection between the generic payload and the ports of CATL, we remodel the ports of the CATL as a part of a vectored representation of scalar ports. In generic payload models, the data is in the form of array of bytes, hence we used the port connections with a byte view of system-level interconnect for auto-generated CATL models. On the application/testbench side we initiate a transaction using generic payload data type. This interface helps in running the CATL model in the test environment to obtain the simulation statistics.

In the Snippet 12 `tlm_data` shows one such data structure. We provide a header file that can be used with any design. The purpose of this header is to connect the input and output transactions to the generated CATL model of the design. In this header we first create a struct data type for both input and output transactions (one can use array as well). In this struct we have the information for the input and output transactions. In generic payload, the data is in the form of array of bytes, hence we used the port connections with a byte granularity.

Here in this transaction the data pointer is set to the data which is used for transaction, length and streaming width of the transaction is also defined. With this information all the port connections can be automatically performed.

Snippet 13 shows the view from the user side. Some trivial details on the interface functions for `input_data` and `transaction` are not explained in detail because of space constraints.

Snippet 12 Details of our interface and data structure to facilitate vector to scalar mapping on the design side

```

#define DATASIZE 8096
// DATASIZE is set for vector mapping of TLM and
// struct data type, it also represents the size of the
// transaction
typedef struct
{
    uint8_t data[DATASIZE];
}tlm_data;

static tlm_data local_in_data;
static uint16_t local_in_len;
static tlm_data local_out_data;
static uint16_t local_out_len;

// interface function for input output connectivity
// with the test environment
void interface_function (tlm_data in_data, uint16_t
in_len, tlm_data out_data, uint16_t out_len)
{
    local_in_data = in_data;
    local_in_len = in_len;
    out_data = local_out_data;
    out_len = local_out_len;
}

```

Snippet 13 Setting up the transaction from test environment side

```

tlm :: tlm_generic_payload * test_trans;
trans-> set_data_ptr(test_out_data);
trans-> set_data_length(n);
trans-> set_streaming_width(n);

```

Snippet 14 Sending the transaction from the test side

```

local_out_data.data[0] = 6;
local_out_len = 1;

```

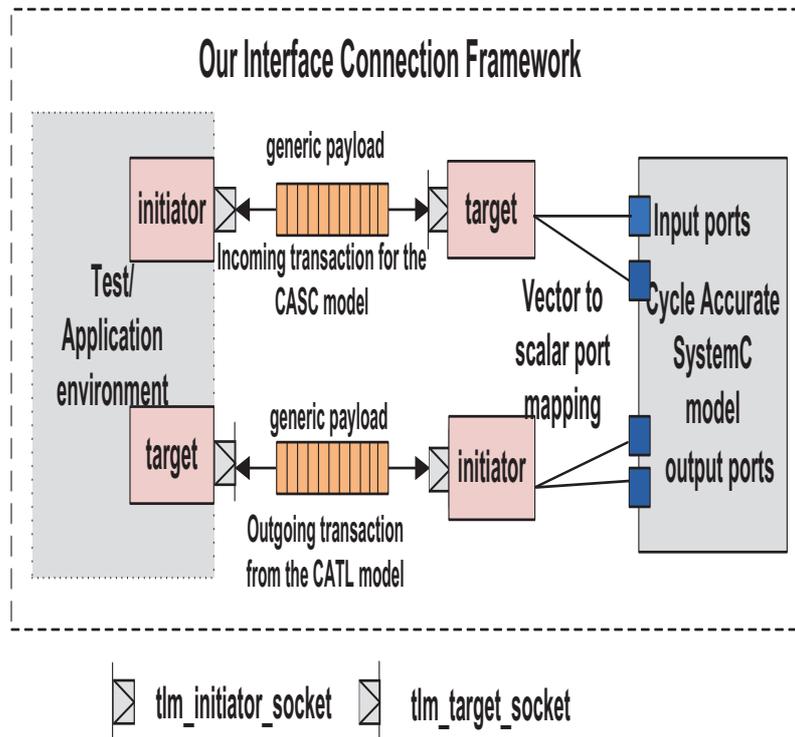


Figure 9.3: Detailed view of system-level interconnect for auto-generated CATL models with transaction level testbench

9.5 Results

In our experiments, we have used production quality industry strength benchmarks. Our benchmarks include various application domain such as compression, decompression, security, DSP related design blocks. RTL/CATL description of the designs are auto-generated and there is no manual modification involved in any part of the design flow. CATL is functionally equivalent to the RTL, cycle-by-cycle. Every part of the design flow is automated using makefile utility. In all experiments, RTL power estimation was performed using PowerTheater [84]. RTL simulation was performed using Modelsim [142], the high-level synthesis was performed using C2R HLS tool [22]. In our analysis, all the benchmarks were synthesized using standard cells hence register power is the biggest component. We did not replace memory with registers. In case memory is used, similar savings can be obtained for the total remaining registers in the design.

In the Table 9.1, column 1 represents the design name, column 2 represents the dynamic register power measured at RTL without clock-gating, column 3 represents the power measured after clock-gating, column 4 represents the % saving shown at the RTL, column 5 shows the % savings predicted by our relative power model without any technology specific information and column 6 (Err1) shows the % error in estimation and column 7 shows the error in measurement using our

power model with technology specific information. On these benchmarks, worst case error of our prediction is 7%, but in most cases we are very close. Even without the information on standard cells the accuracy is off by just 11% in the worst case. This presents a case that such a power reduction model can be used to find the effectiveness of power reduction transformations applied at RTL or higher-level of abstractions. In this chapter, we show the effectiveness on only one type of power reduction i.e. clock-gating but this framework can easily be extended for other power reduction techniques.

Table 9.1: Percentage Power reduction summary on the tested Designs

Design	P_{nocg} (uW)	P_{cg} (uW)	R_{save} %	P_{pm} %	Err1. %	Err2. %
Fibo.	785	467	46.75	47	.25	5.03
Caesar	1630	191	88.3	87.71	.6	5.85
FIR	2090	768	63.3	72.67	9.37	5.07
DES	36500	2350	93.4	96.5	3.1	0.1
Volt.	3240	391	98.8	97.06	1.74	5
Gzip	174000	30100	82.7	94	11.3	7.2
Gunzip	85200	18000	78.9	85.85	5.95	.61
Vit.	48200	20800	56.84	50.66	6.2	7.01

Table 9.2 shows the speed up of the register selection process using our method over methods involving RTL power estimation tools. Here column 2 shows the CPU time spent in our analysis, column 3 shows the CPU time spent in RTL analysis, column 4 shows the speedup obtained and column 5 represent the area savings obtained after applying the clock-gating as compared to non clock-gated designs. Note that TS represents cases where the time taken (too small to measure) were in the range of a few seconds and hence not significant. We only show the cases where the RTL estimation took at least 1000 seconds and then compared against our analysis timings. Note that in the case of the Gunzip benchmark, speedup reached around 3000 times. This can be attributed to two factors. First, the system level analysis is much faster than RTL power analysis (expected in the range of 100-1000 times). Second, the size of the vcd was quite large and it was provided as zipped vcd. It required additional processing to unzip the vcd.

One may argue that a relative power model can be created at the RTL as well, but the majority of power estimation time was spent at RTL simulation and simulation dump analysis. This step can not be avoided for RTL activity/power estimation and hence a power model at the RT level will never be as fast as one at the transaction level.

As mentioned earlier, our default threshold is 0%, which means, that any register that is not 100% active will be considered a clock gating candidate. However, clock gating all or close to 100% registers may incur area overhead and timing closure issues. Also an icgc will be fed with clock port with 100 % activity. Only the generated clock in the clock-tree would be having a lesser activity but the cost of including icgc may overcome the savings. We experimentally verified this claim. We ran some experiment on impacts of the choice of the threshold value such as 0%, 15%

Table 9.2: Speedup with Our Approach

Design	Our Ana. (seconds)	RTL (seconds)	Spd.up	Area sav. %
Fibo.	TS	TS	TS	24
Caesar	TS	TS	TS	9.3
FIR	TS	TS	TS	8.3
DES	TS	TS	TS	.7
Volt.	TS	TS	TS	.01
GZIP	17.27	5153.37	298	0
Gunzip	25.4	77922	3067	5.1
Viterbi	2.3	1886	820	.01

and 25%. In the following, we present results for distinct choices for thresholds at 0%, 15% and 25%. Then we calculated how much power can really be saved in each scenario.

Table 9.3 shows how many registers in the each design show negative values for the Equation 9.4 discussed in the Section 4.1, with each of the threshold settings. First, we saw that algorithm by default tries to find out quite a few registers that should not be clock-gated. Once, we increase the threshold, number of registers needs to be clock-gated decreases. Earlier, from the Table 9.1, we can see that on an average every case design is saving power but finding out the registers not saving power is extremely difficult. Our methodology very efficiently solves this problem.

Our experiments show that additional 2-10% of total register power could have been saved by applying clock-gating selectively. The fourth column of the Table 9.3 shows total number of registers in the designs. The table shows that that up to 45% of the registers would not need clock-gating to save the power for Gzip design. Also for the same designs the number of icgc's that can be saved to insert in the clock-tree ranges upto 260 cells. This is a huge saving in terms of optimizing for the clock tree. This saving comes from the fact that our methodology can effectively find out the places to insert clock-gates only when it is needed.

Interestingly, we did not find much area overhead due to clock gating of selective registers. This can be attributed to two reasons: a-) we do a common enable analysis (by a traversal through CDFG of the design) to make sure that clock-gating logic is introduced for more than one register, which helps in reducing more multiplexers, b-) including a threshold helps in selectively introducing clock-gating. We found that a lot of single bit registers are not a good candidates for clock-gating; some times even 8 bit registers might not require clock-gating. However, we found that register banks of size 16 bits or above show a lot of power savings using clock-gating without exceptions.

Table 9.3: Number of Registers not to clock gate under various thresholds

Design \ Threshold	Auto	15%	25%	Total regs
Fibo.	25	59	61	91
Caesar	8	10	12	194
FIR	0	88	90	243
DES	32	51	125	4326
Volt.	32	32	34	2048
GZIP	208	382	459	17909
Gunzip	280	745	1077	2266

Chapter 10

Conclusion and Future Work

Power estimation and reduction have become primary concerns for hardware design. Increasing adoption of HLS tools in design flow is necessitating the proposition of power estimation and reduction methodologies at higher levels of abstraction. This thesis developed various techniques in these two areas. We observed that for power estimation accuracy and efficiency are the two areas of concern. In most of our approaches, we put an effort in providing a rationale behind the idea and then backed it with experimental results. We also proposed power reduction at high abstraction levels using various granularities of control for clock-gating based power reduction. Finally, we show how to utilize power estimation to guide power reduction. On the power estimation side, we present three different approaches a-) reusing the RTL power estimation frameworks at higher level, b-) providing characterization based power models to facilitate the power estimation at high-level, c-) utilize various verification collaterals such as assertions and properties to speedup power estimation and thereby the design time.

Chapter 4 presented characterization based power estimation methodology, utilizing GEZEL co-simulation environment. In this environment, hardware modeled as FSMMD in GEZEL language can be co-simulated with existing processor models such as ARM. The power model proposed in this case study is a linear regression model, which is trained through various statistical test data. This regression based power model utilizes various states of FSMMD as regression variables compared to input and output based characterization used in previous approaches. It gives a better visibility of power consumption because this approach helps in relating activity of the states of FSMMD to power consumption. This power model is verified on a variety of benchmarks modeled in GEZEL and co-simulated with arm processor model. Further, the work is aided by working on different technology nodes (90 and 130nm) and an elaborate discussion is provided on how to improve the estimation process when lower level power reduction techniques such as clock-gating, power-gating etc. are applied to the design. These reductions are introduced after the RTL design stage, so it is important to know the impact of such techniques on power consumption at high-level. Error in the analysis is found to be less than 10% for the benchmarks used in our experiments.

Chapter 5 presented the idea of reusing RTL power estimation frameworks in HLS based design

methodologies. This idea comes from the fact that in an HLS enabled design flow, RTL is tool generated and the biggest bottleneck in power estimation is the processing of simulation dump (vcd/fsdb). To reduce the time spent by the tool in performing power estimation, we showed how at higher levels speedup could be achieved by using system-level simulation and RT level design details. We experimentally show on a variety of benchmarks that the range of speedup in power estimation process might be as much as 15 times over RTL power estimation techniques. The error or loss of accuracy for the proposed methodology was less than 10% with respect to RTL. The benchmarks experimented for this methodology include a processor model (VeSPA), DSP algorithms such as (FFT, FIR) etc.

Power estimation methodologies/tools may be of no use if it is not aided with representative simulation scenarios. Chapter 6 presented a methodology to utilize verification collaterals to enhance the accuracy of power estimation at higher levels. We used an Esterel based high level modeling framework to model control intensive designs such as a power state machine controller. We presented a methodology to verify various properties extracted from the specification using the formal verification framework inbuilt into the Esterel Studio. We provide guidance on how to write negative properties to help in creating counter cases. This guidance helps us in providing cases for reachability of a state, transitions from a state and staying in particular state of the controller using invariants. These counter cases can be synthesized to test-vectors, and data vectors are used as random test vectors while control test cases are created by utilizing properties. These comprehensive sets of test vectors provide the behavior of the controller of the design for various different states of the design at higher level. The model used in our study has four different states of a processor i.e. active, idle, sleep and deep sleep. We measure the power for each state and valid transitions and finally utilize these numbers during the simulation of the model.

On power reduction side, we proposed various flavors of clock-gating based power reduction. Such a reduction techniques is extensively used during logic synthesis. Overarching idea for our approach was to facilitate these reductions at higher-level and check its efficacy at that level itself. In that direction our proposal contains the following: a-) Enabling clock-gating from ANSI C description itself, b-) Enabling sequential clock-gating and finding a relationship of various registers at high-level, c-) Enabling selective clock-gating, where guidance is provided by high-level power reduction model.

Chapter 7 presented a power reduction technique based on clock-gating from high level model description. We use a C2R HLS tool for our cases studies and show how from ANSI C description a clock-gated RTL can be auto-generated. We show, how various granularities of clock-gating can be affected from the high level description itself. This helps in finding out the opportunities, which lower level tools cannot achieve such as block-level clock-gating. In an HLS context, this is important because a lot of the block-level clock-gating decisions are done at the system-level. However, the design complexity is increasing and a lot of designers have started using the HLS to design the hardware blocks. We showed how various pragmas related to clock-gating can be introduced for an HLS tool to recognize them. We also show how to correctly capture the designer's intent on power management using clock-gating with these pragmas. Our studies also show the utilization of Integrated Clock-Gating Cell (ICGC) in generated RTL. ICGC is generally provided

by the library vendors. Our proposed algorithm considers finding out the common enables for the registers so that across the blocks much lesser number of ICGC's are inserted on the clock-path. The other indirect benefit of such an approach is the introduction of ICGC closer to the root of the clock, otherwise there would be lot more ICGC's towards the leaf of the clock-tree, causing more power wastage in the clock tree.

To extend the clock-gating based power reduction, Chapter 8 showed how sequential clock-gating opportunities can be identified at higher abstraction levels. Generally the difficulty to implement sequential clock-gating comes from two facts 1-) designers may inadvertently change the intended behavior of a design while trying out such aggressive reductions, 2-) If tools are used to facilitate such an optimization, they generally work at netlist level and try to find out the observability don't care or stability conditions. These conditions suggest that a change in the value of register may not show impact on other register's output. Such an exploration at high-level improves the design cycle greatly. In this regard, we proposed a model checking based technique, which utilizes temporal properties to facilitate such a power reduction. It is shown that dependence of registers with each other can be captured in a property and then supplied to a model checker. If the model-checker passes that property then those registers become the possible candidates for sequential clock-gating. These conditions are put as "always true" properties, which is a strong condition but this helps in reducing the chances of facing bugs later in the design stage. This approach addresses the two problems mentioned earlier 1-) enabling ease of verification, 2-) since the relationship between registers is obtained for a behavioral model, there is no need to go to the netlist level. This work is further extended on an HLS tool. Our experimental results show more power savings than conventional clock-gating based power reduction techniques.

Chapter 9 presented a methodology to automatically find the clock-gating opportunities by performing a system-level simulation of a design. To find such an opportunity, a system level power estimation model is used. It utilizes certain properties of clock-gating logic. It shows how technology dependent and independent models help in measuring the suitability of clock-gating applied at higher-level. These models help in applying the clock-gating at appropriate places in the design. The power model is integrated with untimed transaction level test/application environment. Our experimental results show 2-3 orders of magnitude faster analysis can be performed as compared to the state of the art RTL power estimation techniques. This facilitates analysis of more power reduction opportunities and helps in putting the appropriate granularity of clock-gating. This can also serve as a basis for applying estimation guided reduction for other optimizations such as sequential clock-gating, operand isolation, memory gating etc. Currently, most of the approaches rely on putting the reduction algorithm based on architecture or designer knowledge, while this approach distinguishes itself by finding possible power savings by utilizing fast and accurate power reduction models.

Future Work

This thesis presented power estimation and reduction techniques for high level synthesis based design frameworks. There are many possible future directions, where the work presented here will be useful:

1. Power estimation framework presented in Chapter 4 can be extended in many ways. One of the possible extensions is to provide model for leakage power. Regression model presented in this thesis considers constant leakage power. However, for lower technology nodes this relationship is not constant, thus an effort is required to improve the developed model. Second possibility is to extend the regression model for multiple IPs, that is one model to measure power consumption of multiple IPs. Such a model will make the estimation task very easy and requires minimal changes for new architectures and technology nodes.
2. Estimation framework utilizing verification collaterals presented in Chapter 6 can be improved. In our framework, we used random vectors to capture various simulation scenarios in a particular state of the design. This method can be augmented with test compaction techniques. Currently, exhaustive set of vectors are required to capture possible power consumption scenarios. A good compaction mechanism can help in capturing various power consumption patterns by using lesser simulation vectors.
3. We presented dynamic power reduction techniques at high-level (mainly clock-gating and its extensions). One of our approaches has shown enabling clock-gating from ANSI C description. This approach can be extended for some of the dynamic and static power reduction techniques such as operand-isolation, sequential clock-gating, power-gating. This will help in reducing the power consumption of design using multiple techniques while remaining at high-level.
4. In this thesis, we presented how a power reduction model guides HLS to reduce power consumption of a design. In the recent past some trends suggest power wastage caused by overdoing the optimization. Our experimental results also confirm such trends. There are several possible extensions that can be applied to this power model, such as trade-offs for using operand-isolation, sequential clock-gating, memory-gating, power-gating etc.
5. Formalization of the high-level synthesis framework helps us in controlling and getting better view of how design will look like. Existing frameworks such as C2R, used in this thesis, require manual refinements on the functional C model of the target design. These manual refinements are applied to convert the pure sequential code into a structural C code, which contains the information of the macro-architecture of the design. In the macro-architecture, generally designers aim to include necessary parallelism using threads or fork to explicitly bring out the concurrency opportunities in the sequential C-code. During the conversion process various minute details are inserted in the design description, which includes pipelining

stages to achieve the desired throughput or necessary parallelism needed to meet the performance requirement, etc. A study on formalizing various intermediate stages to generate the control-data-flow graph (CDFG) of the input macro-architecture will be very useful. It will help in checking if the refinements are preserving essential properties. Also the formalism of the refinements may help in finding more power saving opportunities during HLS.

6. Finally, our approaches were not extensively targeted for hardware software co-design purposes. It will be interesting to see if these power reduction approaches can be extended in performing hardware software partitioning to avail power aware designing. A robust power model may help in finding out opportunities to save power for FPGA based SoC designs requiring run-time or partial reconfiguration.

Bibliography

- [1] A. Raghunathan, N. K.Jha, and S. Dey, *High-Level Power Analysis And Optimization*. Kluwer Academic Publishers, 1998.
- [2] L. Cai and D. Gajski, "Transaction level modeling: an overview," *IEEE/ACM/IFIP International Conference on Hardware/Software Codesign and System Synthesis*, 2003.
- [3] F. Ghenassia, *Transaction-Level Modeling with Systemc: Tlm Concepts and Applications for Embedded Systems*. Secaucus, NJ, USA: Springer-Verlag New York, Inc., 2006.
- [4] Jairam Ramesh et al., "Clock Gating for Power Optimization in ASIC," www.islped.org/X2008/Jairam.pdf.
- [5] "A Practical Guide to Low Power Design," http://www.powerforward.org/lp_guide/.
- [6] G. E. Moore, "Cramming more components onto integrated circuits," *Electronics*, vol. 38, no. 8, April 1965.
- [7] Mentor Graphics Corp, "Catapult Synthesis," www.mentor.com/products/esl/high_level_synthesis/catapult_synthesis.
- [8] Bluespec, "Bluespec," <http://bluespec.com/>.
- [9] "Celoxica Limited - <http://www.celoxica.com/>," Celoxica Limited, agility Compiler - Advanced Synthesis Technology For SystemC.
- [10] J.-M. Chang and M. Pedram, *Power Optimization And Synthesis At Behavioral And System Levels Using Formal Methods*. Kluwer Academic Publishers, 1999.
- [11] G. Lakshminarayana, A. Raghunathan, N. K. Jha, and S. Dey, "A Power Management Methodology for High-Level Synthesis," *International Conference on VLSI Design*, pp. 24–29, January 1998.
- [12] R. Namballa, N. Ranganathan, and A. Ejnioui, "Control and Data Flow Graph Extraction for High-Level Synthesis," *VLSI, IEEE Computer Society Annual Symposium on*, vol. 0, p. 187, 2004.

- [13] S. Gupta, R. Gupta, N. Dutt, and A. Nicolau, *SPARK: A Parallelizing Approach to the High-Level Synthesis of Digital Circuits*. Kluwer Academic Publisher.
- [14] C. A. R. Hoare, "Communicating sequential processes," *Commun. ACM*, vol. 21, no. 8, pp. 666–677, 1978.
- [15] E. A. Lee and D. G. Messerschmitt, "Static scheduling of synchronous data flow programs for digital signal processing," pp. 237–248, 1989.
- [16] OSCI Group, "SystemC Website," <http://www.systemc.org/>.
- [17] Celoxica, "Handel-C Language Reference Manual RM-1003-4.0, 2003," <http://www.celoxica.com>.
- [18] ESTEREL Technologies, "Esterel Studio," <http://www.esterel-technologies.com/products/esterel-studio/>.
- [19] N. H. E. Weste, D. Harris, and A. Banerjee, *CMOS VLSI DESIGN: A Circuits and Systems Perspective*. Pearson Education, 2006.
- [20] G. Berry, "The foundations of ESTEREL," *Proof, Language and Interaction: Essays in Honour of Robin Milner*, pp. 425–454, 2000.
- [21] Forte Design Systems, "Low Power Design flow," <http://www.forteds.com/solutions/lowpower.asp>.
- [22] Cebatech Inc., "C2R Compiler," <http://www.cebatech.com>.
- [23] Sequence Design Inc., "RTL power management," <http://sequencedesign.com/solutions/powertheater.php>.
- [24] Synopsys Inc., "Design for Low Power," <http://www.synopsys.com/products/solutions/galaxy/power/power.html>.
- [25] GEZEL reference manual, "GEZEL Language Information," http://rijndael.ece.vt.edu/gezel2/index.php/Main_Page.
- [26] S. Ahuja, D. A. Mathaikutty, A. Lakshminarayana, and S. Shukla, "Statistical regression based power models for co-processors for faster and accurate power estimation," *22nd IEEE International SOC Conference*, pp. 399–402, 2009.
- [27] S. Ahuja, D. A. Mathaikutty, A. Lakshminarayana, and S. K. Shukla, "SCoPE: Statistical Regression Based Power Models for Co-Processors Power Estimation," *Journal of Low Power Electronics*, vol. 5, no. 4, pp. 407–415, 2009.
- [28] S. Ahuja, D. A. Mathaikutty, G. Singh, J. Stetzer, S. Shukla, and A. Dingankar, "Power estimation methodology for a high-level synthesis framework," *10th International Symposium on Quality Electronics Design (ISQED)*, pp. 541–546, 2009.

- [29] S. Ahuja, D. A. Mathaikutty, S. Shukla, and A. Dingankar, "Assertion-Based Modal Power Estimation," *8th International workshop on Microprocessor test and Verification (MTV)*, pp. 3–7, 2007.
- [30] S. Ahuja, D. A. Mathaikutty, and S. Shukla, "Applying verification collaterals for accurate power estimation," *9th International workshop on Microprocessor test and Verification (MTV)*, pp. 61–66, 2008.
- [31] S. Ahuja, W. Zhang, A. Lakshminarayana, and S. K. Shukla, "A Methodology for Power Aware High-Level Synthesis of Co-Processors from Software Algorithms," *proceedings of International VLSI design Conference, India, 2010*, pp. 282–287, January, 2010.
- [32] S. Ahuja and S. K. Shukla, "MCBCG: Model Checking Based Sequential Clock-Gating," *IEEE International workshop on High Level Design Validation and Test*, pp. 20–25, November 2009.
- [33] S. Ahuja, W. Zhang, and S. K. Shukla, "System level simulation guided approach to improve the efficacy of clock-gating," *To appear in the IEEE International workshop on High Level Design Validation and Test*, June 2010.
- [34] W. Wolf, "Household hints for embedded systems designers," *IEEE Computer Society Press*, 2002.
- [35] "Unified modeling language," <http://www.uml.org/>.
- [36] "XPower Estimator (XPE), Xilinx Inc." http://www.xilinx.com/products/design_resources/power_central/.
- [37] "PowerPlay Early Power Estimator, Altera Inc." http://www.altera.com/literature/hb/qts/qts_qii53006.pdf.
- [38] V. Tiwari, S. Malik, and A. Wolfe, "Power analysis of embedded software: a first step towards software power minimization," in *ICCAD '94: Proceedings of the 1994 IEEE/ACM international conference on Computer-aided design*. Los Alamitos, CA, USA: IEEE Computer Society Press, 1994, pp. 384–390.
- [39] A. Bona, M. Sami, D. Sciuto, C. Silvano, V. Zaccaria, and R. Zafalon, "Reducing the complexity of instruction-level power models for vliw processors," *Design Automation for Embedded Systems*, 2005.
- [40] J. Russell and M. Jacome, "Software power estimation and optimization for high performance, 32-bit embedded processors," *Proc. of ICCD*, 98.
- [41] V. Tiwari, S. Malik, A. Wolfe, and M. Lee, "Instruction level power analysis and optimization of software," *Journal of VLSI Signal Processing*, 1996.

- [42] L. Benini, D. Bruni, M. Chinosi, C. Silvano, V. Zaccaria, and R. Zafalon, "A power modeling and estimation framework for vliw-based embedded systems," *Proceedings of International Workshop-Power And Timing Modeling, Optimization and Simulation, PATMOS'01*, 2001.
- [43] S. Gupta and F. N. Najm, "Power macromodeling for high level power estimation," in *Design Automation Conference*, 1997.
- [44] A. Bogliolo, L. Benini, and G. D. Micheli, "Regression-based rtl power modeling," *ACM Trans. Des. Autom. Electron. Syst.*, vol. 5, no. 3, pp. 337–372, 2000.
- [45] N. R. Potlapally, A. Raghunathan, G. Lakshminarayana, M. Hsiao, and S. T. Chakradhar, "Accurate power macro-modeling techniques for complex rtl components," in *Proceedings of International Conference on VLSI Design Bangalore, India, Jan. 2001*, p. 235, 2001.
- [46] L. Negri and A. Chiarini, "Power simulation of communication protocols with statec," p. 277-294, *Applications of Specification and Design Languages for SoCs*, Springer Netherlands, 2006.
- [47] L. Zhong, S. Ravi, A. Raghunathan, and N.K.Jha, "Power estimation for cycle-accurate functional descriptions of hardware," in *International Conference on Computer Aided Design, 2004. ICCAD-2004*, 2004.
- [48] M. Caldari, M. Conti, M. Coppola, P. Crippa, S. Orcioni, L. Pieralisi, and C. Turchetti, "System-level power analysis methodology applied to the amba ahb bus," in *Proceedings of the conference on Design, Automation and Test in Europe (DATE)*, 2003.
- [49] N. Bansal, K. Lahiri, A. Raghunathan, and S. T. Chakradhar, "Power monitors: A framework for system-level power estimation using heterogeneous power models," in *Proceedings of the 18th International Conference on VLSI Design (VLSID'05)*, 2005.
- [50] M. Lajolo, A. Raghunathan, S. Dey, and L. Lavagno, "Efficient power co-estimation techniques for system-on-chip design," in *Design Automation and Test in Europe (DATE)*, 2000.
- [51] R. A. Bergamaschi, Y. Shin, N. Dhanwada, S. Bhattacharya, W. E. Dougherty, I. Nair, J. Darringer, and S. Paliwal, "Seas: a system for early analysis of socs," in *CODES+ISSS '03: Proceedings of the 1st IEEE/ACM/IFIP international conference on Hardware/software codesign and system synthesis*. New York, NY, USA: ACM, 2003, pp. 150–155.
- [52] H. Shin and C. Lee, "Operation mode based high-level switching activity analysis for power estimation of digital circuits," in *IEICE transactions on Communications, E90-B(7):1826-1834*, 2007.
- [53] A. Bogliolo, L. Benini, and G. D. Micheli, "Adaptive least mean square behavioral power modeling," *Design Automation and Test in Europe*, 1997.
- [54] ———, "Characterization-free behavioral power modeling," *Design Automation and Test in Europe*, 1998.

- [55] A. Bogliolo, I. Colonescu, E. Macii, and M. Poncino, “An rtl power estimation tool with on-line model building capabilities,” *Int. Workshop Power Timing Modeling, Optimization Simulation Yverdon-les-Bains, Switzerland*, 2001.
- [56] S. Gupta and F. N. Najm, “Energy and peak-current per-cycle estimation at rtl,” *IEEE Transactions on VLSI Systems*, 2003.
- [57] H. Mehta, R. M. Owens, and M. J. Irwin, “Energy characterization based on clustering,” *Design Automation Conference*, 1996.
- [58] M. Nemani and F. N. Najm, “Towards a high-level power estimation capability,” *IEEE Trans. Comput.-Aided Des. Integr. Circuits Syst.*, 1996.
- [59] “Synfora - <http://www.synfora.com/>,” Synfora, pICO Express.
- [60] GAUT, “High-Level Synthesis tool from C to RTL,” <http://wwwlabsticc.univ-ubs.fr/www-gaut/>.
- [61] “Bluespec inc. - <http://www.bluespec.com/>,” bluespec Compiler.
- [62] “Forte Design Systems - <http://www.forteds.com/>,” Forte, cythesizer.
- [63] G. D. Micheli, “Hardware synthesis from C/C++ models,” in *Design Automation and Test in Europe Conference and Exhibition 1999 (DATE'99)*, March 1999, pp. 382–383.
- [64] S. Edwards, “The Challenges of Hardware Synthesis from C-like Languages,” in *International Workshop on Logic Synthesis, Temecula, California (IWLS'04)*, June 2004, pp. 509–516.
- [65] C. Stroud, R. Munoz, and D. Pierce, “Behavioral model synthesis with Cones,” *IEEE Design and Test of Computers*, vol. 5, pp. 22–30, June 1988.
- [66] D. Ku and G. D. Micheli, “HardwareC – A Language for Hardware Design (Version 2.0),” *Technical Report: CSL-TR-90-419, Stanford University*.
- [67] G. D. Micheli, D. Ku, F. Mailhot, and T. Truong, “The Olympus Synthesis System,” vol. 7, pp. 37–53, October 1990.
- [68] R. D. omer, A. Gerstlauer, and D. Gajski, “SpecC Language Reference Manual, Version 2.0,” March 2001.
- [69] D. D. Gajski, J. Zhu, R. D. omer, A. Gerstlauer, and S. Zhao, *SpecC: Specification Language and Methodology*. Kluwer Publications, 2000.
- [70] K. Wakabayashi, “C-based Synthesis Experiences with a Behavior Synthesizer, “Cyber”,” in *Design Automation and Test in Europe Conference and Exhibition 1999 (DATE'99)*, March 1999, pp. 390–393.

- [71] G. Berry and G. Gonthier, “The Esterel Synchronous Programming Language: Design, Semantics, Implementation.,” *Science of Computer Programming*, vol. 19, no. 2, pp. 87–152, November 1992.
- [72] G. Berry, “Esterel on hardware,” *Philosophical Transactions of the Royal Society of London (Series A, 339)*, pp. 87–104, April 1992.
- [73] S. Edwards, “High-level Synthesis from the Synchronous Language Esterel,” in *International Workshop of Logic and Synthesis, New Orleans, Louisiana (IWLS’02)*, 2002.
- [74] B. A. Jose, B. Xue, S. K. Shukla, and J.-P. Talpin, “Programming Models for Multi-Core Embedded Software,” *Multi-Core embedded systems*, 2010.
- [75] P. L. Guernic, M. Borgue, T. Gautier, and C. Marie, “Programming real time applications with SIGNAL,” *Proceedings of the IEEE*, vol. 79, no. 9, pp. 1321–1335, 1991.
- [76] B. A. Jose and S. K. Shukla, “MRICDF : A polychronous Model for Embedded Software Synthesis,” *Synthesis of embedded software - frameworks and methodologies for correctness by construction software design*, 2010.
- [77] ———, “An Alternative Polychronous Model and Synthesis Methodology for Model-Driven Embedded Software,” *Proc. of IEEE Asia and South Pacific Design Automation Conf. (ASP-DAC 2010)*, pp. 13–18, January 2010.
- [78] ESPRESSO Project, IRISA, “The Polychrony Toolset,” <http://www.irisa.fr/espresso/Polychrony>.
- [79] B. A. Jose, J. Pribble, L. Stewart, and S. K. Shukla, “EmCodeSyn: A Visual Framework for Multi-Rate Data flow Specifications and Code Synthesis for Embedded Application,” *12th IEEE Forum on specification and Design Languages (FDL’09)*, pp. 1–6, September 2009.
- [80] B. A. Jose, J. Pribble, and S. K. Shukla, “Faster Embedded Software Synthesis using Actor Elimination Techniques for Polychronous Formalism,” *10th International Conference on Application of Concurrency to System Design (ACSD2010)*, June 2010.
- [81] B. A. Jose, S. K. Shukla, H. D. Patel, and J.-P. Talpin, “On the deterministic multi-threaded software synthesis from polychronous specifications,” in *Formal Methods and Models for Co-Design, 2008. MEMOCODE 2008. 6th ACM/IEEE International Conference on*, June 2008, pp. 129–138.
- [82] B. A. Jose, H. D. Patel, S. K. Shukla, and J.-P. Talpin, “Generating multi-threaded code from polychronous specifications,” *Electron. Notes Theor. Comput. Sci.*, vol. 238, no. 1, pp. 57–69, 2009.
- [83] B. A. Jose, B. Xue, and S. K. Shukla, “An analysis of the composition of synchronous systems,” *Electron. Notes Theor. Comput. Sci.*, vol. 245, pp. 69–84, 2009.

- [84] Sequence Design Inc., “RTL power estimation,” in *PowerTheater User guide, release 2007.2*, 2007.
- [85] “Power Compiler, Synopsys Inc.” <http://www.synopsys.com>.
- [86] Frank Emmett and Mark Biegel, “Power Reduction Through RTL Clock Gating,” presented at the Synopsys Users Group (SNUG), San Jose, CA, 2000.
- [87] W.-T. Shiue and C. Chakrabarti, “Low-Power Scheduling with Resources Operating at Multiple Voltages,” *IEEE Transactions On Circuits and Systems - II: Analog and Digital Signal Processing*, vol. 47, no. 6, pp. 536–543, June 2000.
- [88] A. Kumar and M. Bayoumi, “Multiple Voltage-based Scheduling Methodology for Low-power in the High-level Synthesis,” *IEEE International Symposium on Circuits and Systems*, vol. 1, pp. 371–374, June 1999.
- [89] A. Raghunathan and N. K. Jha, “Behavioral Synthesis For Low Power,” *ICCAD’94*, pp. 318–322, 1994.
- [90] A. Murugavel and N. Ranganathan, “A Game-Theoretic Approach for Binding in Behavioral Synthesis,” in *16th International Conference on VLSI Design, 2003*, January 2003, pp. 452–458.
- [91] D. Chen, J. Cong, and Y. Fan, “Low-Power High-Level Synthesis for FPGA Architectures,” in *IEEE/ACM International Conference on Low Power Electronics and Design (ISLPED’03)*, August 2003, pp. 134–139.
- [92] P. Babighian, L. Benini, and E. Macii, “A Scalable ODC-Based Algorithm for RTL Insertion of Gated Clocks ,” *Design Automation and Test in Europe, (DATE’04)*, 2004.
- [93] R. Fraer, G. Kamhi, and M. K. Mhameed, “A new paradigm for synthesis and propagation of clock gating conditions,” in *DAC ’08: Proceedings of the 45th annual Design Automation Conference*. New York, NY, USA: ACM, 2008, pp. 658–663.
- [94] J. Cong, B. Liu, and Z. Zhang, “Behavior-level observability don’t-cares and application to low-power behavioral synthesis,” in *ISLPED ’09: Proceedings of the 14th ACM/IEEE international symposium on Low power electronics and design*. New York, NY, USA: ACM, 2009, pp. 139–144.
- [95] M. Munch, B. Wurth, R. Mehra, J. Sproch, and N. Wehn, “Automating RT-Level Operand Isolation to Minimize Power Consumption in Datapaths,” *Design Automation and Test in Europe (DATE’00)*, 2000.
- [96] N. Agarwal and N. Dimopoulos, “High-level fsm design and automated clock gating with codel,” *Electrical and Computer Engineering, Canadian Journal of*, pp. 31–38, 2008.

- [97] Mentor Graphics Inc., “Advanced Clock Gating Techniques in Catapult C Synthesis,” <http://www.techonline.com/learning/techpaper/218400046>.
- [98] G. Singh, J. B. Schwartz, S. Ahuja, and S. K. Shukla, “Techniques for Power-aware Hardware Synthesis from Concurrent Action Oriented Specifications,” *Journal of Low Power Electronics (JOLPE)*, vol. 3, no. 2, pp. 156–166, August 2007.
- [99] G. Singh and S. K. Shukla, “Algorithms for Low Power Hardware Synthesis from CAOS - Concurrent Action Oriented Specifications,” *Special Issue of International Journal of Embedded Systems on Power/Energy/Thermal topics (IJES’07)*, 2007.
- [100] —, “Low-Power Hardware Synthesis from TRS-based Specifications,” *Fourth ACM and IEEE International Conference on Formal Methods and Models for Codesign (MEM-OCODE’06)*, Napa Valley, CA, USA, pp. 49–58, July 2006.
- [101] —, “Model Checking Bluespec Specified Hardware Designs,” *8th International Workshop on Microprocessor Test and Verification (MTV)*, Austin, TX, USA, December 2007.
- [102] —, “Verifying Compiler Based Refinement of Bluespec Specifications using the SPIN Model Checker,” *15th International SPIN Workshop on Model Checking of Software (SPIN)*, Los Angeles, CA, USA, pp. 250–269, August 2008.
- [103] G. Singh, S.S.Ravi, S. Ahuja, and S. K. Shukla, “Complexity of Scheduling in Synthesizing Hardware from Concurrent Action Oriented Specifications,” *Power-aware Computing Systems, Dagstuhl Seminar Proceedings 07041*, Dagstuhl, Germany, 2007. [Online]. Available: {<http://drops.dagstuhl.de/opus/volltexte/2007/1105>}
- [104] G. Singh, J. B. Schwartz, and S. K. Shukla, “A Formally Verified Peak-Power Reduction Technique for Hardware Synthesis from Concurrent Action-Oriented Specifications,” *Journal of Low Power Electronics (JOLPE)*, vol. 5, no. 2, pp. 135–144, August 2009.
- [105] S. Irani, G. Singh, S. K. Shukla, and R. K. Gupta, “An Overview of the Competitive and Adversarial Approaches to Designing Dynamic Power Management Strategies,” *Journal on IEEE Transactions on VLSI Systems (TVLSI)*, vol. 13, no. 12, pp. 1349–1361, December 2005.
- [106] G. Singh, S. Gupta, S. K. Shukla, and R. Gupta, *Chapter 11 of CRC Handbook on EDA for IC System Design, Verification and Testing*. CRC Press, Taylor and Francis Group, 2006, ch. Parallelizing High-Level Synthesis: A Code Transformational Approach to High-Level Synthesis.
- [107] J. Bingham, J. Erickson, G. Singh, and F. Andersen, “Industrial Strength Refinement Checking,” *Formal Methods in Computer Aided Design (FMCAD)*, Austin, Texas, USA, November 2009.

- [108] G. Singh, "Optimization and Verification Techniques for Hardware Synthesis from Concurrent Action-Oriented Specifications," 2008.
- [109] J. Brandt, K. Schneider, S. Ahuja, and S. K. Shukla, "The model checking view to clock gating and operand isolation," June 2010.
- [110] J. Monteiro, S. Devadas, P. Ashar, and A. Mauskar, "Scheduling Techniques to enable Power Management," *Proceedings of Design Automatin Conference*, pp. 349–352, June 1996.
- [111] A. Chattopadhyay, B. Geukes, D. Kammler, E. M. Witte, O. Schliebusch, H. Ishebabi, R. Leupers, G. Ascheid, and H. Meyr, "Automatic ADL-based Operand Isolation for Embedded Processors," *Design Automation and Test in Europe (DATE'06)*, 2006.
- [112] D. Dal, D. Kutagulla, A. Nunez, and N. Mansouri, "Power Islands: A High-Level Synthesis Technique for Reducing Spurious Switching Activity and Leakage," in *48th Midwest Symposium on Circuits and Systems, 2005*, vol. 2, August 2005, pp. 1875–1879.
- [113] M. Banga, M. Chandrasekar, L. Fang, and M. S. Hsiao, "Guided test generation for isolation and detection of embedded trojans in ics," pp. 363–366, 2008.
- [114] M. Banga and M. S. Hsiao, "A region based approach for the identification of hardware trojans," pp. 40–47, 2008.
- [115] —, "A novel sustained vector technique for the detection of hardware trojans," pp. 327–332, 2009.
- [116] —, "Vitamin: Voltage inversion technique to ascertain malicious insertions in ics," pp. 104–107, 2009.
- [117] K. S. Khouri, G. Lakshminarayana, and N. K. Jha, "High-Level Synthesis of Low-Power Control-Flow Intensive Circuits," *IEEE Transactions on Computer-Aided Design (TCAD'99)*, vol. 18, December 1999.
- [118] A. Chandrakasan, M. Potkonjak, R. Mehra, J. Rabaey, and R. Brodersen, "Optimizing Power Using Transformations," vol. 14, pp. 12–31, January 1995.
- [119] J. Uchida, N. Togawa, M. Yanagisawa, and T. Ohtsuki, "A Thread Partitioning Algorithm in Low Power High-Level Synthesis," in *Asia and South Pacific Design Automation Conference (ASP-DAC'04)*, January 2004, pp. 74–79.
- [120] V. Raghunathan, S. Ravi, A. Raghunathan, and G. Lakshminarayana, "Transient Power Management Through High Level Synthesis," In *Proceedings of the ICCAD*, pp. 545–552, 2001.
- [121] W.-T. Shiue and C. Chakrabarti, "ILP-based Scheme for Low Power Scheduling and Resource Binding," in *IEEE International Symposium on Circuits and Systems, Geneva, Switzerland (ISCAS'00)*, vol. 3, May 2000, pp. 279–282.

- [122] W. T. Shiue, “High level synthesis for peak power minimization using ilp,” *In Proceedings of the IEEE International Conference on ASSAP*, pp. 103–112, 2000.
- [123] S. P. Mohanty and N. Ranganathan, “A framework for energy and transient power reduction during behavioral synthesis,” *In Proceedings of the International Conference on VLSI Design*, pp. 539–545, 2003.
- [124] S. P. Mohanty, N. Ranganathan, and S. K. Chappidi, “ILP Models for Energy and Transient Power Minimization During Behavioral Synthesis,” *Proceedings of the International Conference on VLSI Design*, p. 745, 2004.
- [125] A. Lakshminarayana, S. Ahuja, and S. Shukla, “Coprocessor design space exploration using high level synthesis,” *10th International Symposium on Quality Electronics Design (ISQED)*, pp. 879–884, March 2010.
- [126] H. Li, S. Bhunia, Y. Chen, T. N. Vijaykumar, and K. Roy, “Deterministic clock gating for microprocessor power reduction,” in *HPCA '03: Proceedings of the 9th International Symposium on High-Performance Computer Architecture*. Washington, DC, USA: IEEE Computer Society, 2003, p. 113.
- [127] Mitch Dale, “The Power of RTL Clock-gating,” <http://www.chipdesignmag.com/display.php?articleId=915>.
- [128] JMP Description, “JMP Data Analysis,” <http://www.jmp.com/software/jmp7/>.
- [129] X. Liu and M. C. Papaefthymiou, “A markov chain sequence generator for power macro-modeling,” in *ICCAD '02: Proceedings of the 2002 IEEE/ACM international conference on Computer-aided design*. New York, NY, USA: ACM, 2002, pp. 404–411.
- [130] F. Vahid, *Digital Design*. John Wiley and Sons Publishers, 2006.
- [131] Synopsys Inc., “VCS Comprehensive RTL Verification Solution,” <http://www.synopsys.com/vcs/>.
- [132] David J. Lilja, Sachin S. Saptnekar, “Designing digital computer systems with verilog,” <http://www.arctic.umn.edu/vespa/>.
- [133] H. Patel and S. Shukla, “Power state machine implementation,” http://fermat.ece.vt.edu/EWD/Power_State_Machine.htm.
- [134] M. Ohashi and et al., “A 27mhz 11.1mw mpeg-4 video decoder lsi for mobile application,” *ISSCC digest of technical papers*, pp. 366–367, 2002.
- [135] J. E. Stine, J. Grad, I. Castellanos, J. Blank, V. Dave, M. Prakash, N. Iliev, and N. Jachimiec, “A framework for high-level synthesis of system-on-chip designs,” in *MSE '05: Proceedings of the 2005 IEEE International Conference on Microelectronic Systems Education (MSE'05)*. Washington, DC, USA: IEEE Computer Society, 2005, pp. 67–68.

- [136] Cadence Inc., “The Cadence SMV model checker,” <http://www.kenmcmil.com/smv.html>.
- [137] “Verification Interacting with Synthesis (VIS),” <http://vlsi.colorado.edu/vis/>.
- [138] “A System for Sequential Synthesis and Verification,” <http://www.eecs.berkeley.edu/~alanmi/abc/>.
- [139] Jerry Frenkil, “Too Much of a Good Thing?” <http://chipdesignmag.com/display.php?articleId=3316>.
- [140] S. Ahuja, S. T. Gurumani, C. Spackman, and S. K. Shukla, “Hardware coprocessor synthesis from an ansi c specification,” *IEEE Design and Test of Computers*, vol. 26, pp. 58–67, 2009.
- [141] OSCI Group, “TLM 2.0,” <http://www.systemc.org/downloads/standards/tlm20/>.
- [142] Mentor Graphics Inc., “RTL Simulation ,” <http://www.model.com/>.