Low Power Design And Verification of Module Solution

Project Report

Submitted in partial fulfillment of the requirements

for the degree of

Master of Technology

In

Electronics & Communication Engineering

(VLSI Design)

By

Vishal S. Bhimani

(12MECV04)

Department of Electronics & Communication Engineering

Institute of Technology

Nirma University

Ahmedabad-382 481

May 2014

Low Power Design And Verification of Module Solution

Major Project Report

Submitted in partial fulfillment of the requirements

for the degree of

Master of Technology

In

Electronics & Communication Engineering

(VLSI Design)

By

Vishal S. Bhimani

(12MECV04)

Under the Internal Guidance of

Dr. Usha Mehta

and

External Guidance of

Mr. Mohit Jain

Department of Electronics & Communication Engineering

Institute of Technology

Nirma University

Ahmedabad-382 481

Declaration

This is to certify that

1. The thesis comprises of my original work towards the degree of Master of Technology in VLSI Design at Nirma University and has not been submitted elsewhere for a degree.

2. Due acknowledgement has been made in the text to all other material used.

Vishal S. Bhimani

Certificate

This is to certify that the Major Project entitled "*Low Power Design and Verification of Module Solution*"submitted by *Vishal S. Bhimani* (*12MECV04*), towards the partial fulfillment of the requirements for the degree of Master of Technology in Electronics and Communication Engineering Branch of Institute of Technology, Nirma University, Ahmedabad is the record of work carried out by him under my supervision and guidance. In my opinion, the submitted work has reached a level required for being accepted for examination. The results embodied in this major project, to the best of my knowledge, haven't been submitted to any other university or institution for award of any degree or diploma.

----------	----------
Guide	
Dr. Usha Mehta	Dr. N.M.Devashrayee
Associate Professor	PG Co-ordinator (VLSI Design)
----------	----------
Dr. P.N Tekwani	Dr K Kotecha
Professor & Head, EE	Director, IT-NU
Date: ----------	**Place: Ahmedabad**

Acknowledgement

I would have never succeeded in completing my Report without the cooperation, Encouragement and help provided to me by various people.

Firstly, my sincere thanks to the MSM team,for their help during this training. Their wisdom, clarity of thought and support motivated me to bring this project to its present state.

I am highly indebted to my Group manager **Rashna Seli** and project manager **Mr. Mohit Jain** for providing necessary information regarding the project and also for their constant guidance, supervision, kind co-operation, and invaluable support in all aspects. My thanks and appreciations also go to my colleagues and team members in developing the project and for providing me with a lively and energetic work environment.

I would like to express my sincere gratitude to **Dr. Ketan Kotecha** (Director, Nirma University, Ahmedabad) for his continuous guidance, support and enthusiasm. I would take this opportunity to thank **Dr. P. N. Tekwani** (Head of Department, Electrical Engineering), **Dr. N. M. Devashrayee** (Professor and Program Coordinator, M.Tech - EC (VLSI Design)), Internal Guide **Dr. Usha Mehta** (Sr. Associate Professor, M.Tech - EC (VLSI Design)) and all the faculties at **Nirma University (VLSI Design)**, for their vision and relentless effort, support, and encouragement to provide me with this excellent opportunity to carry out my project work in such a highly renowned and esteemed organization, ST Microelectronics. I am equally thankful to **ST Microelectronics** for providing me the invaluable exposure to the industry and the current market trends.

Finally, I would like to express my heartfelt thanks to my parents and colleagues for their blessings and for their constant love and support.

Vishal S. Bhimani

(12MECV04)

Abstract

Low Power Design is driven by market pressures, regulatory aspects and technology migration to 65nm and lower. Multi-Voltage design is hence essential. Multi-Voltage design brings many new challenges to design as well as verification.

The power supplied to elements in an electronic design affects the way circuits operate. Although this is obvious when stated, todays set of high-level design languages have not had a consistent way to concisely represent the regions of a design with different power provisions, nor the states of those regions or domains. Unified Power Formate standard provides an HDL-independent way of annotating a design with power intent. In addition, the level-shifting and isolation between power domains may be described for a specific implementation, from high-level constraints to particular configurations. When the logic in a power domain receives different power supply levels, the logic state of portions of the design may be preserved with various state-retention strategies. This standard provides mechanisms for the refined and specific description of intent, effect, and implementation of various retention strategies. Incorporating components into designs is greatly assisted by the encapsulation and specification of the characteristics of the power environment of the design and the power requirements and capabilities of the components; this information encapsulation mechanism is also described in this standard. The analysis of the various power modes of a design is enabled with a combination of the description of the power modes and the collection, generation, and propagation of switching information.

STMicroelectronics At A Glance

- A world leader in providing the semiconductor solutions that help our customers improve quality of life for everyone, both today and in the future
- Among the world's largest semiconductor companies
- A leading Integrated Device Manufacturer serving all electronics segments
- A leading technology innovator (around 12,000 researchers approx. 21,500 patents)
- Key strengths in Multimedia Convergence, Power Applications and Sensors
- Rich, balanced portfolio (ASICs, Application-Specific Standard Products and Multi-Segment Products)
- A pioneer and visionary leader in sustainability
- President and CEO: Carlo Bozotti
- 2011 revenue $9.73 billion
- Approximately 50,000 employees including STEricssion at December 31, 2011
- Advanced research and development centers in 10 countries
- 12 main manufacturing sites
- Corporate Headquarters Geneva, Switzerland

- Global presence with sales offices all around the world
- Public since 1994 - shares traded on New York Stock Exchange (NYSE: STM), Euronext Paris, and Borsa Italiana
- Created as SGS-THOMSON Microelectronics in June 1987, from merger of SGS Microelettronica (Italy) and Thomson Semiconducteurs (France)
- Renamed STMicroelectronics in May 1998

Contents

List of Figures

Chapter 1

Introduction

Almost every engineer who designs and verifies integrated circuits (ICs) today is under tremendous pressure to reduce power. In a rather unprecedented way, we are faced with market pressures, regulatory pressures and process technology factors to reduce power, all at once. This multi-pronged pressure on power consumption has led many designs to adopt aggressive low power design techniques that involve the control of voltage. This book looks at the emerging voltage controlled techniques used today and how this new generation of power managed designs is verified. As is often the case with most design processes, verification is often a larger task than the design effort itself. The problem with power managed ICs is worse because this is not an area where traditional verification techniques can be easily applied. As we establish later in the chapter, an altogether new view of looking at boolean logic and the verification process is needed. With this book, we endeavor to develop a re-usable, rigorous and comprehensive methodology to verify power managed designs.

For over 25 years, the Electronic Design Application (EDA) industry did not have to deal with the various complexities of voltage controlled power managed designs. This was because most IC power management was handled at the system level with little impact to the specification and verification process and hence out of the scope of EDA tools used in the IC design process. Chip level aspects were restricted to clock gating, which caused a significant change to synthesis, placement

and routing portions of the design flow almost a decade ago. These however, did not deal (and did not need to deal) with the demands of voltage based control of ICs or IC components. Both automated analysis and implementation capabilities were based on hardware description languages (HDLs) that excluded a notion of voltage connections to the underlying logic. The exclusion of voltage from HDLs and the current adoption of voltage based control of ICs are at odds with each other, causing enormous disruption in the world of IC design and verification.

Many collective years of intellectual property (IP), existing code bases, EDA tools and flows have been built on the existing paradigm. Voltage control of CMOS devices now needs to be accounted for in these databases and processes: it is the desired output of a control system involving hardware, software, digital circuits and analog blocks. The mastery of this complexity is now the job of front end RTL design and verification engineers as well as back end implementation and signoff engineers.

In the following sections we look at what is driving power management first. Then we proceed to look at how designs are responding by adopting voltage control techniques. That leads us to the main problemverification of such design techniques. The chapter concludes with a section on the structure of the book and methodology adoption.

1.1 DRIVING FACTORS FOR POWER MANAGEMENT

Semiconductors are increasingly being used in mobile/consumer devices. Mobile devices are driven by battery life and form factor primarily, though performance is equally important, especially in the consumer market and for mobile applications involving multi-media. Even in the wired power segment, there is enormous pressure to reduce form factor and deliver better performance, as the availability of energy

supply and runtime electricity costs gains prominence. There is also increased awareness that the installed base of semiconductors contributes more to global warming than the entire airline industry. Hence, many governments have sought to regulate the electronics industry. Lastly, technology migration which typically helps achieve lower power consumption has complicated matters. This section looks at these three aspectsmarket, regulations and technology and their impact. However, before we dive into that, we need to take a deeper look at the word power itself.

DENSITY : Density refers to the amount of power consumed within an area, hence the heat dissipated in an area. Consider a 1cm x 1cm packaged IC dissipating an average of 1W. This alone amounts to a power density of 10GW per square kilometer, the power of many nuclear reactors combined! Needless to say, the heating produced by such a density is enormous: skin temperatures on the package often reach 100 degrees Celsius, with die/junction temperatures reaching 125 degrees. Heat dissipation leads to cost in terms of both components required to dissipate it such as a better package, heat sink, fan etc. and a run time cost in terms of operating the fan, cooling system etc. Often an overheated device collection such as a server farm incurs an equal or greater run time cost than the initial cost itself. In some extreme cases, the excess temperature being in proximity to a battery has caused fires and explosions in cell phones and laptops.

DELIVERY : Delivery is the problem relating to the amount of current that has to be delivered at progressively lower supply voltages and the ability to withstand fluctuations in current. Delivery is the most misunderstood of power requirements. However, it constitutes one of the most important aspects of power management. As process technology shrinks in its feature length, the current supplied is actually increasing; hence the fluctuations, especially the rapid ones in current requirements are especially troubling to design. Delivery is commonly classified as IR drop and di/dt problems. One of the nastiest aspects of power management is that density and leakage mitigation techniques often cause problems in delivery of current.

LEAKAGE : Leakage is the current consumed by the chip even when there

is no activity. In prior generations of process technology, such idle current was a negligible entity. In deep submicron designs, this is no longer the case: causing immense problems for mobile devices, where it has a severe adverse effect on battery life. Leakage is not merely a battery powered device problem. With the advent of Green Design, there is an immense impact of leakage on all electronic systems. This is covered under the regulatory pressures section. Note that transistors have leakage current even when there is activity.

LIFETIME : Lifetime refers to the decreasing reliability of chips caused due to higher current densities. The cross section of wires in the ICs of today is quite narrow compared to that of yesteryears. Coupled with an increase in absolute current itself, the implication is that the material in the wires is likely to degrade much faster. A reduction in current, therefore benefits lifetime concerns tremendously. Electronics sold into various markets have requirements on lifetime, often accompanied by manufacturer warranties. Hence, ensuring reliable operation for that period is a key design constraint.

1.2 TECHNOLOGY MIGRATION AND POWER

Until 90nm technology became available, the cure for power ailments was simple: migrate the chip to a lower geometry, benefiting from the lower capacitance and the lower supply voltage. The onset of leakage has spoiled that easy route, and problems in the basic scaling for delivery and lifetime have started creeping up. For a quantitative illustration, Two chips A and B at 0.18um are integrated into a single chip AB at 0.13um. This reduces power overall, but notice that the current is now higher (at 0.45A), supplied at a lower voltage. When the next generation of integration with a similar Chip C, providing additional functionality, is done at 90nm, chip ABC consumes about 0.7W, as opposed to say, an expected 0.4-0.5W. This is problematic because scaling down process generations is broken at 90nm and beyond.

The current requirements at 90nm are quite high and the cost of such a current delivery mechanism is higher compared to previous generations, since the voltage is 6Verification Methodology Manual for Low Power also quite low. Further, the fact that almost 30

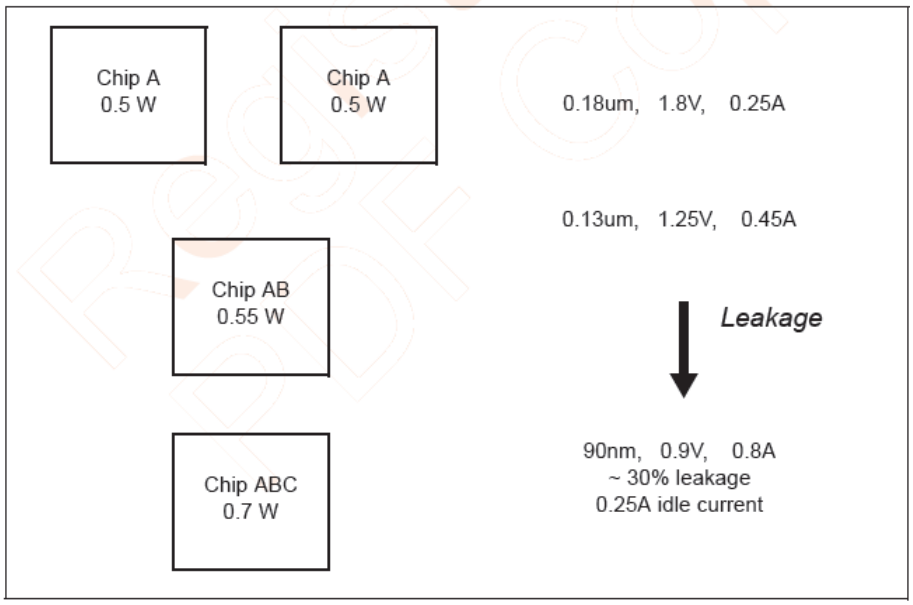

Figure 1.1: Process Migration

1.3 MULTI-VOLTAGE DESIGN IN PRACTICE

First, one must recognize the fact that once we resort to voltage control techniques, voltage is no longer a constant in space or time. The voltage applied to different transistors may be different even on the same die. The voltage applied to the same circuit may vary in time. Both of these effects can be termed as Multi-Voltage design.

Let us begin by looking at idle power. Fortunately, idle power is systemically easy to exploit for reduction. We never use all the systems or all the functions of a system all the time. The opportunity to reduce wasted energy is tremendous.

Hence, we must design electronics to idle efficiently. They must not consume energy when they are idling.

Idle mode efficiency can be envisaged at the design level to be an Off by Design architecture. Such a design must ideally satisfy two fundamental properties:

- Any part of the system must be powered up only when needed.
- Any part of the system that is not in use must be turned off.

Compliance with these two rules in the architecture implies that there is a power manager built into every IC and system that we build. Such a power manager would have to be both hardware and software savvy, as most modern systems and ICs are very heavily coupled with software for their active/idle profile.

The reader may point out a subtle corollary to the Off by Design concept. A system does not need to perform at the same level all the time. For example, one may buy a High Definition television, but may be watching a low resolution program or photo images on screen. Or a smartphone may be used as just a phone most of the time, a phone, as opposed to its use as a camera or a video player. So, while there may be some components that don't have enough opportunities to be turned off, there may be opportunities to run them at significantly lower power. A technique like this can be used to reduce the active power consumed by the device. In these lower performance modes, there is an opportunity to lower the supply voltage and hence the active power is reduced. Hence, a good corollary to Off by Design is as follows:

- A circuit must always be operated at the lowest possible voltage at which its per- formance goals are met.

DYNAMIC POWER REDUCTION: Dynamic Power reduction aims to control one of the following: a reduction in the amount of capacitance toggling, a reduction in the frequency of operation and/or a reduction in the supply voltage applied.

As such many designs see minimization of area (viz. capacitance) along high activity paths as well as reducing the overall capacitance to a minimum. One such effective method practiced widely is clock gating, which effectively reduces the capacitance of high activity networks. Commercial tools and solutions have delivered this feature in an automated way all the way from early synthesis to post layout optimization. Such designs do not suffer from the immense increase in verification as seen by voltage controlled designs.

Frequency reduction is architecturally driven and is controlled through hardware and software. Frequently, this is accompanied by dynamic voltage reduction as well, since we need only a lower supply voltage to meet the timing requirements of a slower Emergence of Voltage Control cycle time. In this book, we focus more on this latter method, since it represents the emerging challenge to traditional design and verification.

There are other methods to reduce voltage and frequency. If all parts of a chip do not need to be at high frequency, we can use a lower supply voltage for that part of the chip. Such a partition is typically called a multi-VDD architecture and is used in lieu of dynamic voltage scaling.

LEAKAGE POWER REDUCTION: Leakage current relies mostly on control of threshold voltage and/or supply voltage (since we really cannot control temperature variations). In the simplest of techniques, a static assignment of transistors to either a high Vt version, or a low Vt. Typically, only about 5-15

In more aggressive techniques, system or block level idleness can be exploited to further dramatically reduce leakage power and hence idle energy. One such technique is reverse-bias. As we describe further in Chapter 2, this is a technique where the Vt of transistors is increased by application of a voltage. This reduces leakage, although no operations can be performed.

Another way to exploit system idleness is to reduce the supply voltage to a point where the state is not corrupted. This technique known as low Vdd Standby, is quite effective at reducing leakage: just as with reverse bias, at speed operation is

not possible.

In both reverse bias and low Vdd Standby, there is no loss of state, unless something goes wrong. They are therefore quite useful for blocks with large amounts of memory. However, for many logic blocks, there is no need to hold state or perhaps there is a only combinatorial logic internally. This provides the opportunity to cut power off altogether to this block, a technique known as power gating which is again covered in further chapters.

Selection of the optimal architecture is alas, not as simple as this in real life. Every chip faces multiple constraints at the system level in terms of its average draw, battery life, bill of materials, latency and throughput in different modes. The optimization of such architecture is well beyond the scope of this book. Often, design teams do not spend time optimizing the architecture.

Chapter 2

Design Elements and Styles for Multi-Voltage

2.1 DESIGN ELEMENTS

2.1.1 RAIL/POWER NET

A virtual or physical network that is connected to the output of a voltage source, possibly controlled through some switching or value setting mechanism. Typically, this network does not have a logic value associated with it. However, it is common to find Logic 1 and Logic 0 assigned to their representations in HDLs, especially for wire constructs in Verilog HDL. A rail or a power net is a controlling entity for the transistors, but is often a controlled entity of the power management scheme.

2.1.2 VOLTAGE REGULATOR

This is a handily available component in the industry, both as a discrete component or a mixed-signal intellectual property block. However, its foray into the digital world is Multi-Voltage Basics new facet brought about by power management. A voltage regulator usually has a source power supply (connected to AC mains, bat-

tery, or another voltage regulator), an output regulated voltage, digital control bits to set the voltage or turn the regulator on/off, and some status/handshake signals to indicate stability of output. Obviously, many more sophisticated implementations are possible.

From the power management point of view, the voltage regulator squarely injects itself into the midst of the control scheme. The result is that a change in power management state must interact with the voltage regulation scheme in some fashion and then wait for a response.

2.1.3 PRIMARY RAILS

Primary rails actively drive logic values on signals and are responsible for supplying or charging/discharging current at the cells output. In current CMOS standard cell based implementations, the Vdd/Vss connections of a CMOS cell are often the functional rails. As illustrated in Figure 2-1, the Vdd rail supplies the current to charge the output node to Logic 1 when needed and the drive to maintain that level. Similarly, the Vss rail of Figure 2-1 discharges the output node as needed and maintains a Logic 0.

2.1.4 SECONDARY RAILS

Secondary rails control the behavior of the cell, but do not actively supply current for charge/discharge of the cells output. In Figure 2-1, the output of the charge pump, Vslp is connected to the gate of a transistor labeled Footer. Variations of Vslp influence the ability to functional rails to charge/discharge, but Vslp itself is not part of the charge/discharge path. Note that the Gate connection (node G in Figure 2-1) of the CMOS cell is also a secondary rail in that it influences the cell behavior but does not connect with the charge/discharge path. A secondary rail at one cell can be primary to another and vice versa. The interaction of the rail with CMOS logic elements determines its nature within the context of the logic element.

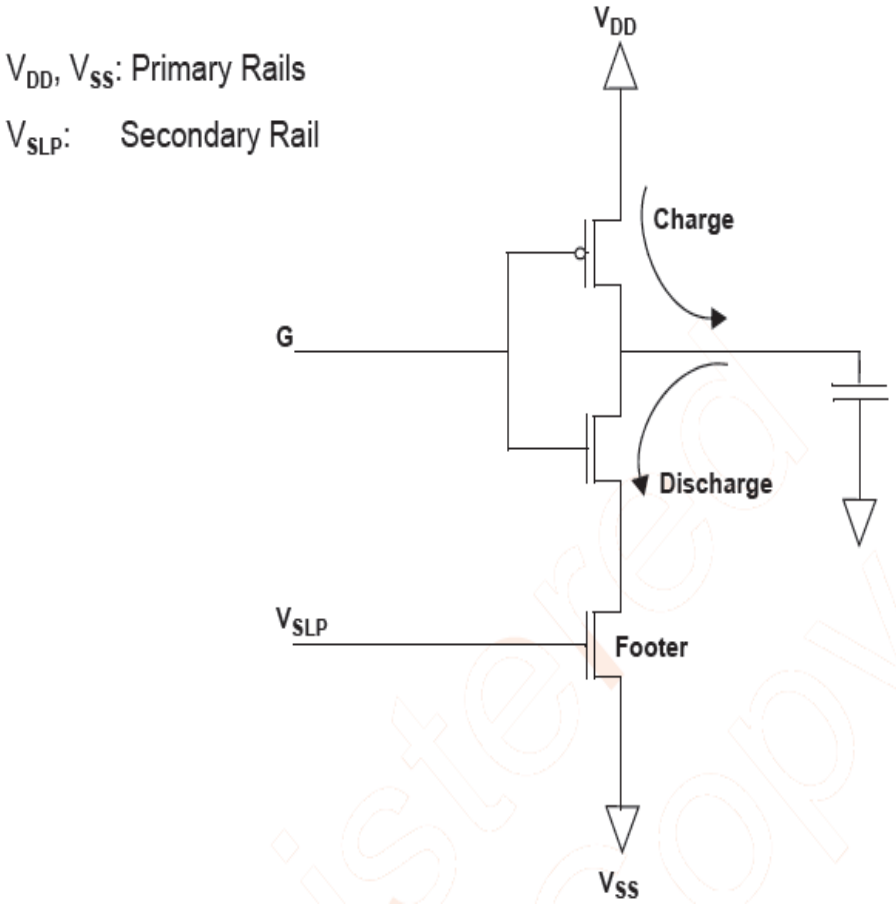

Figure 2.1: Rails

2.1.5 HEADER AND FOOTER CELLS

High threshold voltage (Vt) transistors are used to switch or gate the connection between the primary rails and the PMOS/NMOS elements that implement logic, as shown in Figure 2-2. (This leads to the term power gating, which will be discussed later.) These High Vt transistors are called Headers when inserted between Vdd and the PMOS elements; they are called Footers when they are in the path between Vss and NMOS elements. Except for some special cases, Headers are comprised of PMOS transistors, and Footers are comprised of NMOS transistors.

In theory, a single ideal Header or Footer cell is sufficient for any amount of logic that is fed off of the primary rails. However, in practice it is common to use multiple parallel transistors, which are often staged in time by a control mechanism. (Read: target of verification!). Many practical issues surround the use of Header/Footer cells,mostly tied to layout topology. The Low Power Methodology Manual [2] offers more information on this aspect. One important property to note at this time however is that based on the voltage applied to the gate of the Header/Footer (relative to the Vdd/Vss levels), the logic elements are either connected to Vdd/Vss or cut off from them. In the latter state, it amounts to shutting down these elements. Later in this chapter, in the Power Gating section, we will discuss the verification oriented aspects of using headers and footers.

Furthermore, the gate of the Header/Footer may be connected either to a rail or a logic signal. This distinction is really for the users design flows and a matter of how the on/off control to these gates is generated. The response of the Header/Footer is always to the relative voltage levels and must be considered mixed signal behavior. Note that Header/Footer cells are themselves simple voltage regulators with just an On/Off function.

2.1.6 ISOLATION

Isolation is a technique to protect a receiving island that is active from a signal originating in an island that is turned off. The temporal variation of islands from On to Off and vice versa implies this isolation is a gating element, controlled in conjunction with the state of the source island. As innocuous as it sounds, isolation is a common source of functional errors, thereby implying that it will be a significant target of our verification efforts. In terms of logic gates, AND/NAND/OR/NOR types are typical choices for isolation gates. Latches may also be used, but come with additional verification challenges, among other additional costs such as area/delay. Isolation gates come with a whole host of correctness requirements: from their

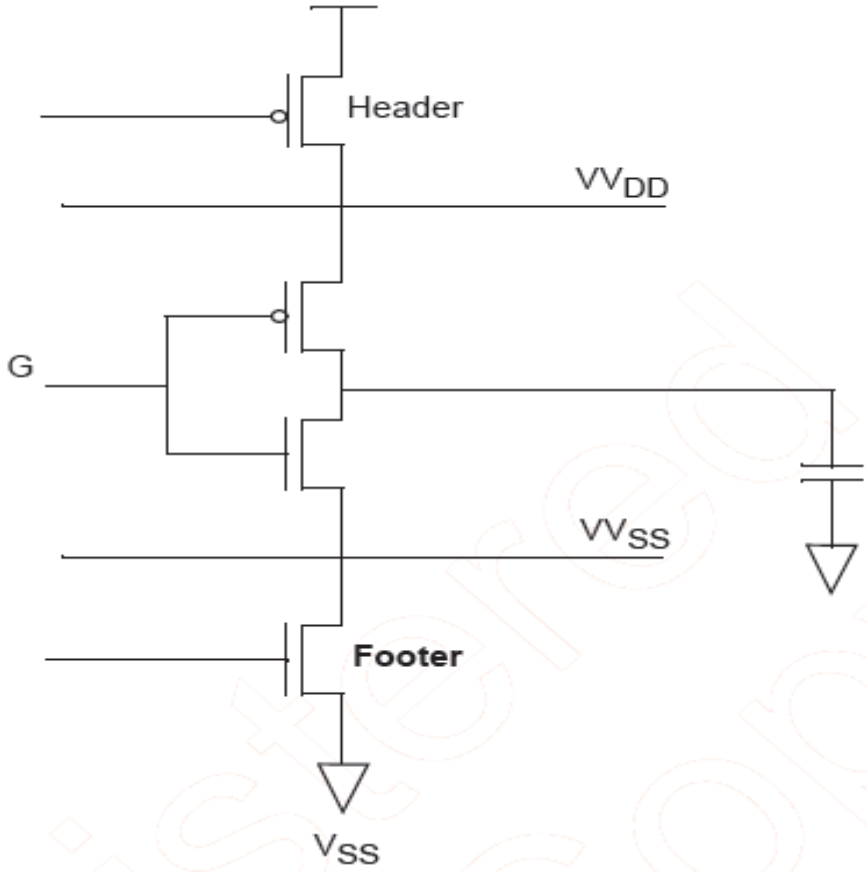

Figure 2.2: Header and Footer

placement spatially to the isolation control temporally.

2.1.7 LEVEL SHIFTING

Level Shifting is a technique to convert a signal driven by one set of primary rails (and hence referenced as Logic 0 and Logic 1 to them) to another set of primary rails to another set of primary rails on a spatial crossing. Typically, level shifters are either high voltage to low voltage or vice versa, although auto level shifters are also common: they shift both high to low and low to high. An example is shown in Figure 2-3, marked by the gate between domains V1 and V2 as Level Shifter.

An important class of level shifter is the Enabled Level Shifter, which is the

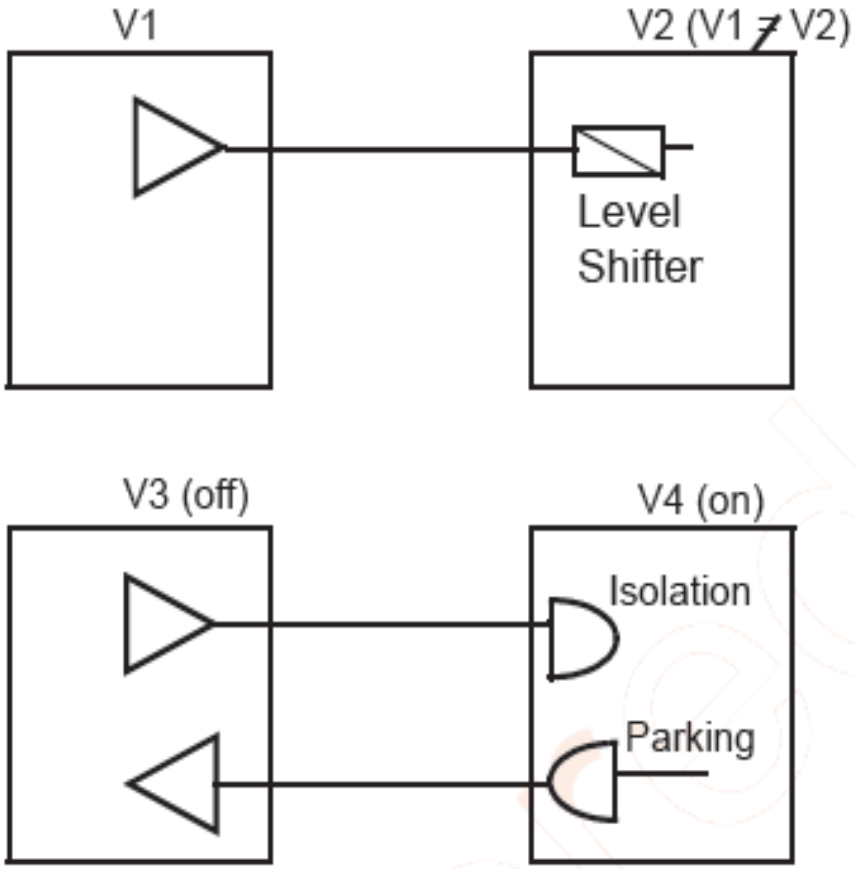

Figure 2.3: ISOLATION AND LEVEL-SHIFTER

combination of an isolation device and level shifting.

The reader may wonder what really is behind low to high and high to low. The hidden implication here is that we are referring to the Vdd levels of the driver and receiver. This is a convenience that we can get away with: most ICs we implement today are common ground and have a common Vss level. Therefore, the level shifting of Logic 0 is not needed, except in some special situations.

Why do we need level shifters in the first place? Because Logic 1 is charged up to the level of the Vdd of the driver. This may be either inadequate to represent a Logic 1 at the receiver based in its Vdd level or even if it is, it may cause excess current consumption. (Remember the CMOS transfer curve?)

One of the other common questions in many engineers minds is why then, is high to low shifting needed? The drivers Vdd, surely, is sufficient to represent Logic 1. The answer is that over driving the gate of the receiver can result in quite a bit of additional current consumption and in some extreme cases even cause a Logic 0 at the driver to be perceived as a Logic 1 on the receiving end. Sometimes, the transistors of the receiving domain may not be even rated for the higher voltage requiring a high to low conversion, or risking damage to the oxide layer in the receiving gates. As we migrate to lower process geometries, gate leakage is an increasing component. Hence, the cost of not converting high to low rises with smaller geometries

2.1.8 PMU

The Power Management Unit (PMU) is the functional block that makes the temporal variation of voltages occur. Its function is to sense the state of the system or IC and ensure that the device makes a transition to the appropriate system. Power Management Units dont just control voltages, they make sure that clocks, resets, retention controls etc. all act in conjunction. They monitor the various blocks and assert the appropriate controls. They are often a combination of hardware and software based control and observation.

2.2 DESIGN STYLES FOR MULTI-VOLTAGE

The previous section described the basic elements of multi-voltage control. Various design styles are possible by using the basic elements and controlling them appropriately with a power management unit.

2.2.1 SHUTDOWN

This is a state in which the Vdd to a domain or an island is turned Off, not necessarily to zero volts. The voltage regulator typically resides off-chip, but it is increasingly common to achieve this with an on-die regulator or a power switch. (Power gating will be discussed later in this chapter.) Although in current usage, shutdown has become synonymous with power gating, it can be accomplished by direct control of the voltage source as well, not just power switches.

Shutdown works primarily as a leakage reduction technique. As we discussed before, a shutdown of a block implies that all state is lost and needs to be either reset or restored upon wakeup. Further, shutting down a domain also implies that certain amount of pre- and post shutdown control of clocks, isolation control signals etc., is required Design Styles for Multi-Voltage.

2.2.2 STANDBY

In general, Standby is a low power state, in which a quick wakeup is expected. State retention in the memory elements is essential. Typically, clocks are gated, but PLLs are not disabled. However, there could be multiple grades of Standby, turning off more circuitry as time passes. The overall Standby scheme is one that involves a process of gradually turning off (on) clocks, PLLs, and voltages. In this context, we focus on the use of voltage control. Various techniques are used for Standby.

2.2.2.1 CLOCK GATED STANDBY

Clock Gated Standby is a mode in which no Vdd control exercised. The clock network is gated, and that in itself saves a lot of dynamic power. The PLL may be On or Off, depending on the desired resume time. While there is no Vdd control, this mode serves as a preparation for it.

Clock gating can be combined with two voltage control techniques to reduce leakage. Low Vdd standby and Back Bias, which are described below.

2.2.2.2 BACK BIAS

Back Bias, also called Reverse Bias as shown in Figure 2-4, raises the threshold voltage of the island, which in turn reduces the leakage. State in registers is not lost. The effective frequency of operation is reduced; this technique is typically applied for standby mode. Blocks with large RAMs and register files are especially amenable to this technique.

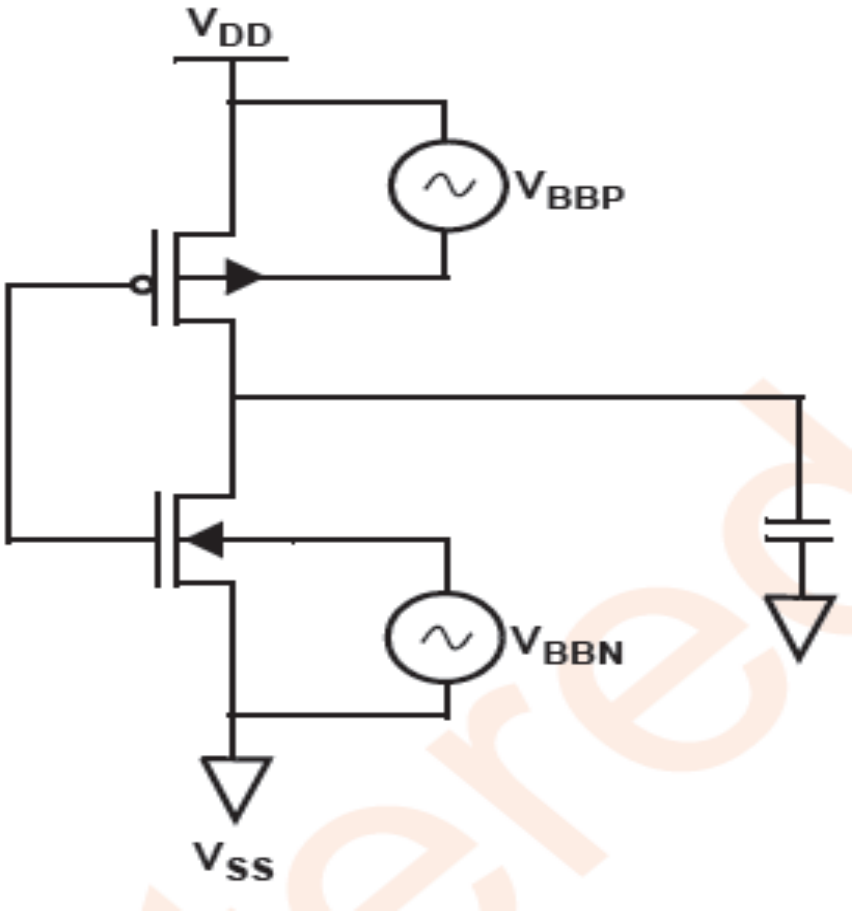

Figure 2.4: Back Bias

2.2.2.3 LOW VDD STANDBY

Low Vdd Standby is a state in which the Vdd is lowered to just enough value for the memory elements to not lose state. Inevitably, the clock must be gated down for

this situation, because there (typically) is not sufficient drive to do write operations at this voltage level. Notethe outputs of a block in low Vdd Standby need level shifting, which they do not need in the case of back bias. However, there is no need to layout another grid to the bulk terminals.

2.2.3 SLEEP/POWER GATING

Power Gating is a form of Shutdown and is increasingly referred to as Power Shut Off or MTCMOS. A Header or Footer are applied to the block of logic. The Header and/or Footer are controlled with either a zero or negative Vgs to cut off the cells from Vdd and Vss. The Header and Footer may be shared between cells, rows, or even whole areas of the die. Often, multiple Header/ Footer transistors are used in parallel, to ensure there is sufficient current delivery strength and capacity to withstand fluctuations in current.

A common approach is using logic signals to control the Header/Footer transistors, although charge pumps provide better leakage reduction. However, using a logic signal relieves the need to use both charge pumps on die as well as route a separate rail.

2.2.4 RETENTION

Retention is a combined variation of Standby and Shutdown: the block is shutdown achieving low leakage, but state is not lost. There are many ways to accomplish this, all the way from using software based schemes to leaf cell implementations in the library. The generic aspect in every scheme however is utilizing a save operation prior to shutdown and a restore after wake-up, both of which are carefully orchestrated in control sequence. Design for Retention: Strategies and Case Studies [5], a paper co-authored by one of the authors, David Flynn, is an excellent source of information on this topic.

2.2.5 DYNAMIC VOLTAGE SCALING

Dynamic Voltage Scaling is a technique in which the rails of an island are varied, especially, the Vdd rail, to achieve various power/energy targets. Strictly speaking, this should be Dynamic Voltage Frequency Scaling (DVFS), since frequency is usually scaled with Voltage.

Chapter 3

Power Management Bugs

3.1 STRUCTURAL ERRORS

These are errors caused by a design structure. This class of errors is mostly detectable by pure static analysis,assuming there is a specification to check against . In general, these errors can be rectified by changing the design structure. While the emphasis on the existence of a specification seems trivial, one must remember that at the time of this books initial edition, power intent specification languages were in their infancy of discussion and standardization. It is not necessary that structural errors occur in only the RTL or the netlist stages. They could occur in any phase of the design. However, the end result is often the same. A structural error results in electrically bad connections, leading to functional bugs, device breakdown or excessive power consumption.

3.1.1 ISOLATION AND RELATED BUGS

Power gating and external VDD shutdown are the most commonly used low power techniques today. Therefore, isolation devices are used quite often. It is therefore, no surprise that many bugs related to isolation devices occur in the design. This subsection gives a few examples.

3.1.1.1 MISSING ISOLATION

When a source domain such as the On/Off domain in Figure 3-1 goes into the off state, an isolation device must be placed in the path between any source signals from the off domain to any fanouts in the on or standby state domains. In the absence of isolation, a floating net will directly feed the CMOS gate or diffusion terminal in the on domain, causing the following effects: logic corruption in the subsequent stages of the on domain, excessive power consumption in the on domain and in extreme cases, a breakdown of the device itself. Missing Isolation (or level shifter) is one of the most common bugs encountered in low power design. Fortunately, it often takes only a structural check to detect this. Avoidance of this bug has also become easier as design teams transition from manual insertion at RTL to automated insertion in netlist stage through the use of policies in the power intent description.

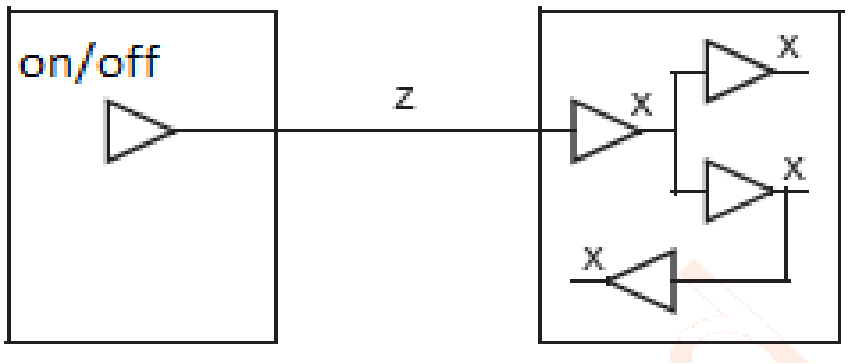

Figure 3.1: Missing Isolation

3.1.1.2 INCORRECT ISOLATION POLARITY

By default, most signals are isolated to Logic 0. However, if the signal being isolated is an active low signal, then the isolation polarity must be Logic 1. As shown in Figure 3-2, it would be incorrect to isolate an active low signal named reset_N to Logic 0, since it would place all of its fanout in reset state for the duration of isolation. Similarly, many protocols, such as PCI and USB, are implemented with

active low request lines for bus access. Isolating them to Logic 0 causes spurious grants, degradation in bus performance, and frequently deadlock conditions.

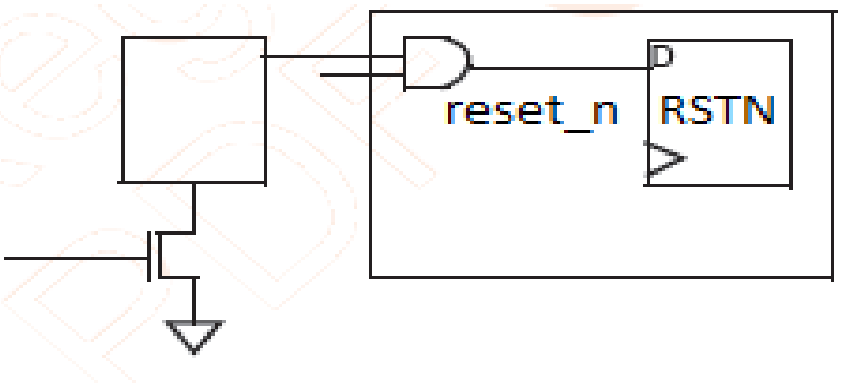

Figure 3.2: Isolation Wrong Polarity

3.1.1.3 INCORRECT ISOLATION ENABLE POLARITY

In this case as shown in Figure 3-3, the isolation enable polarity is incorrect. This can be caused either by an incorrect connection to the Power Management Unit (PMU) or by implementation error in RTL/Netlist. The effect is that the isolation device is not functional when it is neededin fact, it may turn the other way around. Isolation may kick in when not needed and be dysfunctional when actually needed.

3.1.1.4 INCORRECT ISOLATION GATE TYPE

In this case as shown in Figure 3-3, the isolation enable polarity is incorrect. This can be caused either by an incorrect connection to the Power Management Unit (PMU) or by implementation error in RTL/Netlist. The effect is that the isolation device is not functional when it is neededin fact, it may turn the other way around. Isolation may kick in when not needed and be dysfunctional when actually needed.

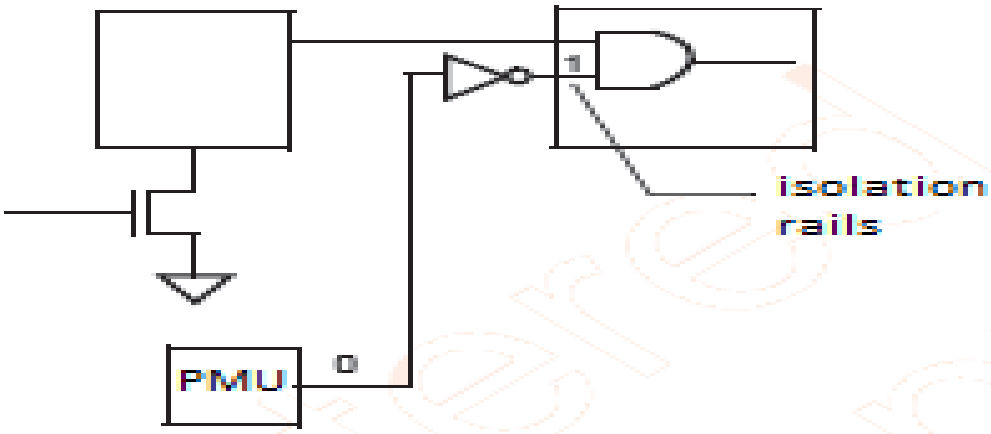

Figure 3.3: Isolation Enable Wrong Polarity

3.1.1.5 REDUNDANT ISOLATION

Often, the problem with isolation is that it is applied where it is NOT needed. Such isolation is redundant and is extremely detrimental to functionality. There are various flavors of redundant isolation: two dangerous ones are described below.

Consider the spatial crossing (crossover) between Domain 1 and Domain 2. No isolation device is needed: inserting one is a waste of area and causes a timing penalty. Even worse, when Domain 3 is switched off, Signal A is needlessly tied to an isolated value of Logic 0, thereby freezing Signal A as long as Domain 3 is off. Any changes on Signal A are lost in this period. Redundant isolation devices can hence cause very serious functional errors by killing or freezing logic signals inappropriately. This bug can be caught statically. Note that only the isolation device on Signal A is redundant in this case. A variation of this bug can be caused by control error as we see in the section Control/Sequence Errors- this will be a situation where the required isolation gate on Signal B is rendered redundant

3.1.2 LEVEL SHIFTING AND RELATED BUGS

A level shifter must exist on a spatial crossing when there is a disparity in voltage levels between its source and destination. As the user might recall from Chapter 2,

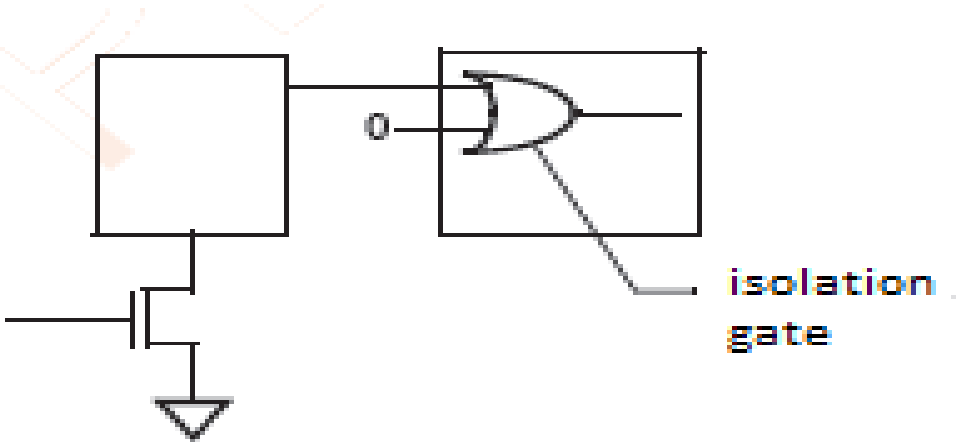

Figure 3.4: Incorrect Isolation Gate Type

a Logic 1 is referenced with respect to the Vdd rail of the domain from which the signal originates. The driven net is charged up to a voltage equal to the Vdd. This works well within the domain. However, when crossed into another domain, the voltage level on the net may not be adequate to be recognized as a Logic 1. Even if it is, it may place the receiving device(s) in a region of the CMOS curve that is consuming enormous current.

3.1.2.1 LEVEL SHIFTER OUT OF RANGE

As shown in Figure 3-5, it is not sufficient to connect a level shifter. One must pick a level shifter that can handle the voltages applied to it. In this figure, the level shifter specification requires Vs (the source voltage) to be between 0.8V and 1.2V. The applied source voltage of 0.7V is lower than the minimum required specification of 0.8V as source voltage. Note that library formats such as Liberty now have attributes that have specifications for voltage ranges, which can hence be checked statically by tools. The error is mostly caused by manual insertion. Automated tools that synthesize power intent often do not have this bug.

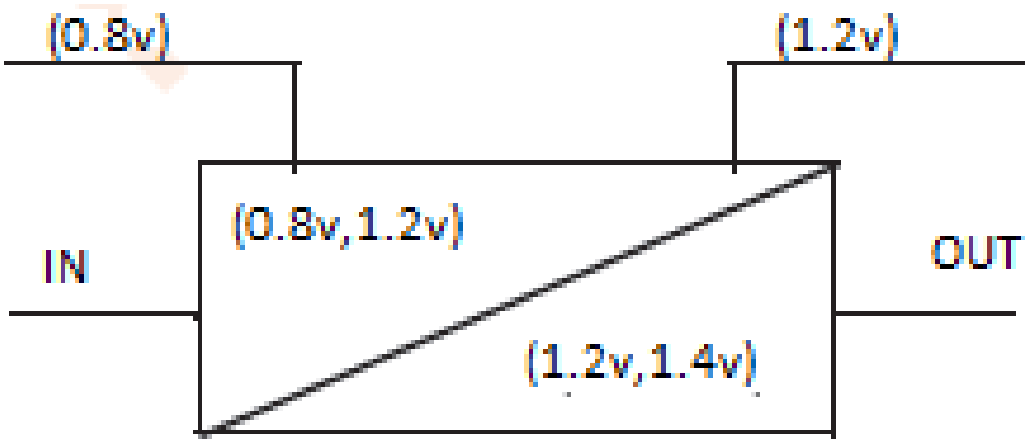

Figure 3.5: Level Shifting device out of range

3.1.2.2 INCORRECT DOMAIN OF LEVEL SHIFTER

A level shifter must be connected to the appropriate rails on the source and destination side, so as not to violate its purpose itself, as shown in Figure 3-6. In this example, the level shifter is placed in intermediate domain V1 which has a level of 1.0V. This means that another level shifter is required when the output of this level shifter is connected to fanout in the power domain supplied by V3. Thus two level shifters are required instead of one, increasing area, power and delay. The placement of level shifters in intermediate islands is fine as long as it does not require another level shifter and isolation device by doing so.

3.1.3 OTHER STRUCTURAL ERRORS

We certainly do not want to leave the user with the impression that structual errors are all related to Isolation and Level Shifting! Other structual errors are possible but not often within the purvey of front end logical verification.

For example, consider a power switch structure that causes excessive in-rush current (di/dt effects). This error needs to be detected by a post layout electrical integrity analysis. Unchecked, this leads to a collapse of the power rails and logic levels.

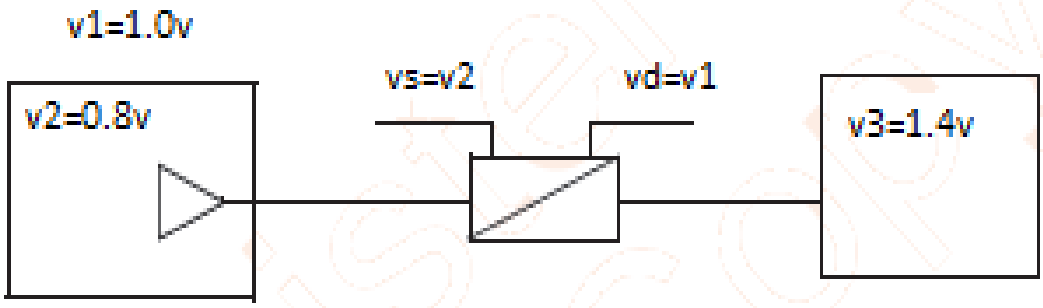

Figure 3.6: Level Shifter in Wrong Domain

Errors are also possible within individual standard cells by erroneous transistor structures. These require static transistor level analysis. Increasingly these tools are available in the market for use at the netlist level.

Hard macro internals form another source of structural errors. In our experience, some of these are caused solely due to the fact that there was no formal specification method to communicate the internal power domain partitions of the hard macro or the pin level partitions. For example, consider an IO-pad cell that has 1.8V and 1.2V domains. If the 1.2V domain can be internally shutoff, but has built in isolation, then it is an error to connect isolation devices externally as well. Fortunately, library specification formats have evolved to represent all these attributes, leading to automated tool support for static detection of such errors.

3.2 CONTROL/SEQUENCE ERRORS

This category is by far the most frequent source of functional errors. Errors often occur because the control of power management events is not exercised appropriately. Power management events include both movements in voltage as well as changes in logical signals such as isolation enables, resets, clock gating signals etc.

In many cases, these bugs rise from the design of the power management unit itself. Most of todays designs combine IP blocks from various sources, including inherited blocks from previous generations and third parties. The power management unit is therefore a major aspect of the integration, and its thorough verification is central to avoiding bugs.

3.2.1 ISOLATION CONTROL ERRORS

The mere presence of isolation devices is not sufficient. The isolation device must be exercised on its control input to be actively isolating when there is a shutdown condition. This condition cannot be verified structurally: it needs temporal verification either by dynamic simulation or other techniques, such as property checking.

3.2.1.1 INCORRECT ISOLATION ENABLE TIMING

In this bug, which is extremely common, the source island is shutdown, but the receivers isolation gate is not enabled. This causes an unknown value at the output of the isolation gate. Refer to Figure 3-7 below. When Signal A transitions from Logic 1 to Logic 0, the footer transistor is cut off. However, Signal B, the isolation enable, arrives much later. This causes an unknown logic value and excess current consumption in the isolation gate. Worse yet, the unknown logic value at the output of the isolation gate may not be observed in the logic regression and end up detected only on silicon, as high power consumption. The methodology recommended in this book seeks to avoid this first by requiring coverage of such elements and second by recommending assertions around each isolation device to test for such a condition.

3.2.1.2 MEMORY CORRUPTION IN STANDBY (UNSAFE WRITE)

Memory elements such as registers, latches and RAMs in the design can be taken to an optimum lower level of voltage (Low Vdd Standby) and hence significantly reduce leakage. This way, they retain the value but will not be able to drive or

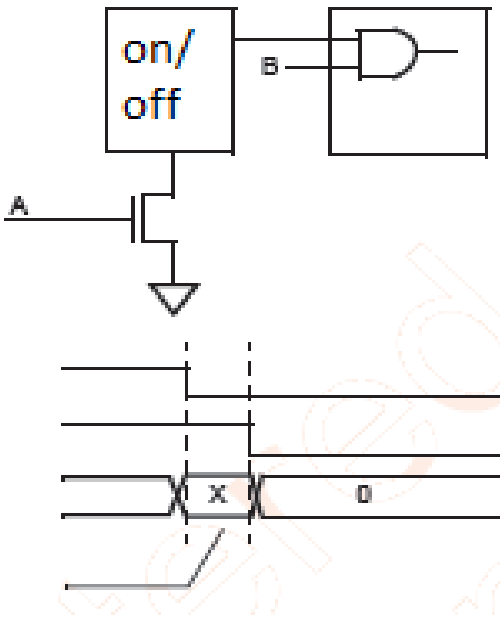

Figure 3.7: Incorrect Isolation Enable Timing

accept any value to be written to them. If a clock wiggles at the input of standby element, it may so happen that the data it is retaining gets corrupted. This is a functional bug. As the memory element wakes up, it may contain corrupted bits that can cause a design malfunction

3.2.1.3 SAVE AND RESTORE SEQUENCE

Retention methodology typically requires special control signals, such as save/restore, to be generated by a power management unit. The sequence for this power save mode should be typically, save-power gate-wakeup-restore. Though the sequence is followed, premature scheduling of these signals will corrupt the restored value and design will wake up in unknown state. This will cause the design to malfunction

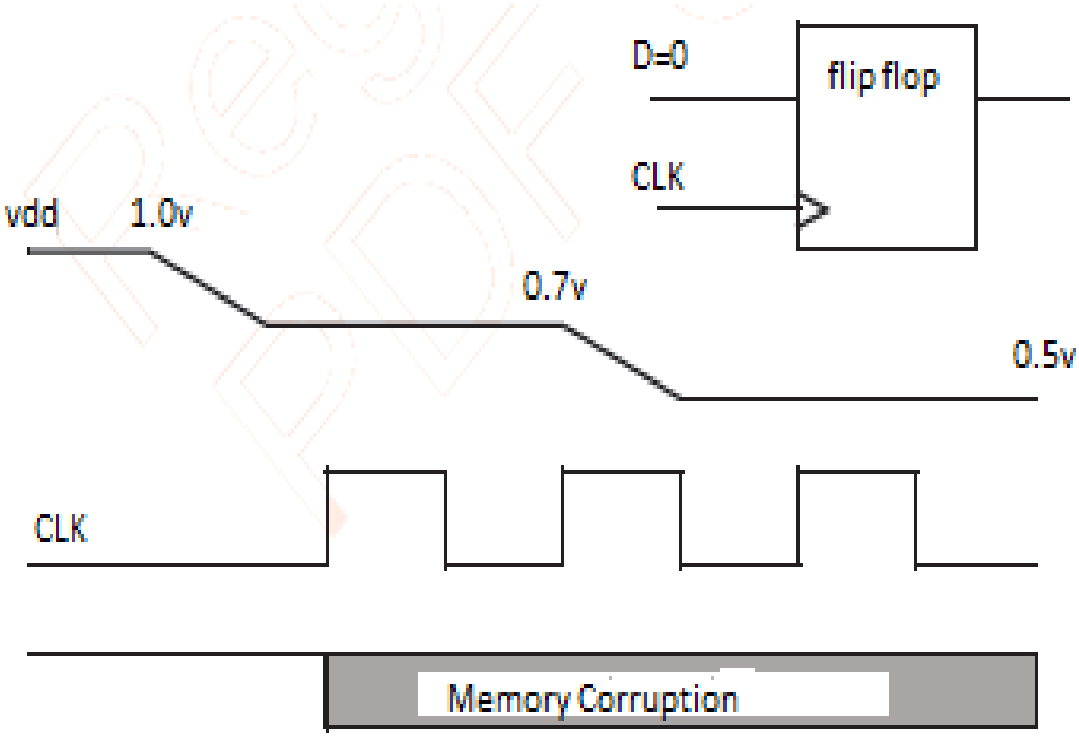

Figure 3.8: Unsafe Write

3.2.2 LOGIC CORRUPTION

Most physical libraries come up with ranges of voltages where a level shifter is not required. In this case, the specification is such that any difference less than 150mV of Source-Receiver driving voltages does not need a level shifter.

Many designers look at the power state table combinations to determine whether a level shifter is needed.The two states possible for the domains are (1.2V, 1.1V) and (0.9, 0.8V) as power states. If one were to statically analyze this, it can be concluded that a level shifter is not needed between these domains. However, when the design transitions from one power state to another, the voltages may not maintain the difference to be under 150mV, hence violating the level shifter rule. This kind of error is purely dynamic and can be caught by only by targeted testing or exhaustive random testing. Some amount of static analysis of ramp rates is possible, but that may not include the dynamic events that influence the voltage ramp.

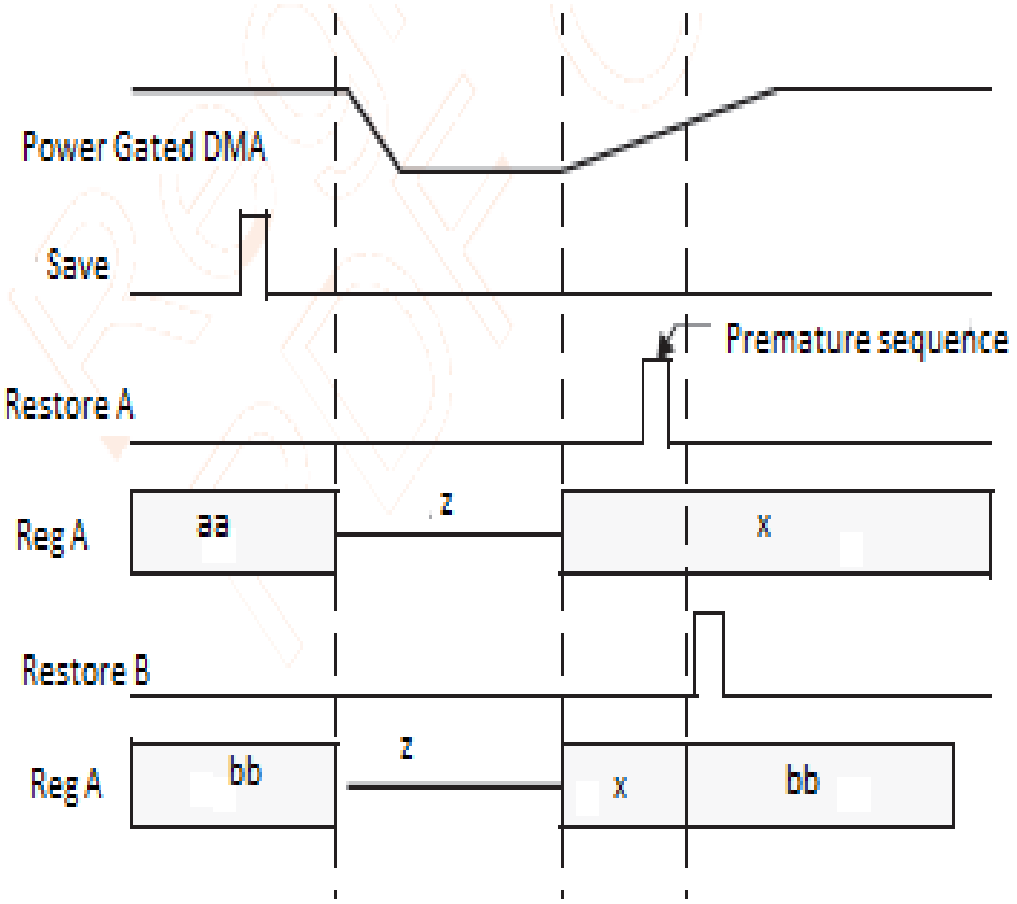

Figure 3.9: Premature Restore Error

The design may or may not display functional failure; however, electrical failure, timing variations, and excess power consumption are likely. The damage entirely depends on the difference in voltages and whether any transistor structures are exposed to unsafe conditions. The source of error is obviously in control: multiple rails moving in value at the same time cause a lot of electrical disruption. A CMOS device with its source/destination Vdd rails moving in value at the same time could behave quite unpredictably. This can be avoided by scheduling the voltage changes in a safe manner.

8 bit data needs to be transferred to the Logic Wrapper which includes 4 bits for the ring selected, 2 bits for divide speed and 2 bits for divide leak.

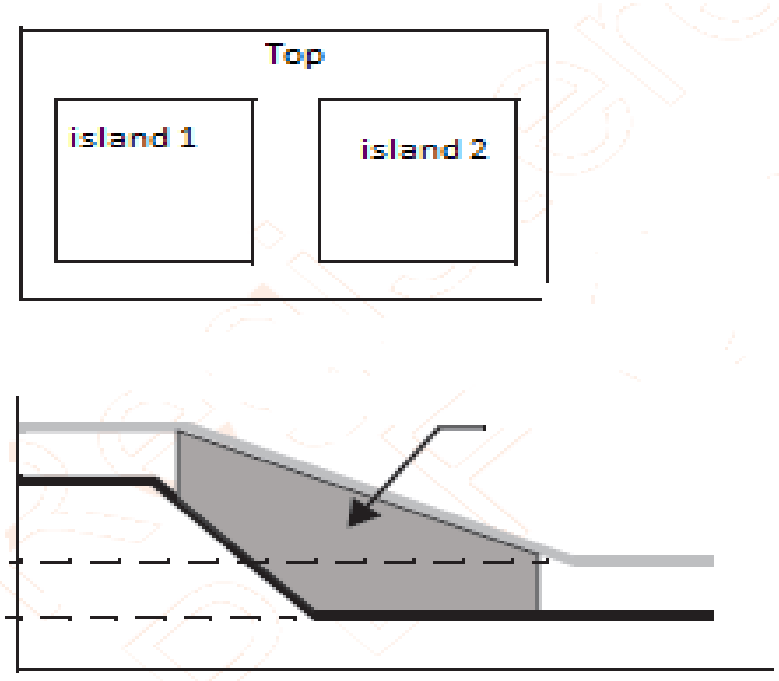

Figure 3.10: Dynamic Level Shifting Error

Program the REFERENCE counter before sensor control register, reference counter defines the time window during which the period of the oscillator will be measured. This time window is defined as a number of reference clock cycles. The counted number should be loaded into the reference counter register. Two values of REFERENCE COUNTER length are supported, 16 and 32.

SENSOR CONTROL REGISTER is programmed for parallel run of all sensors or single sensor run. The values of SENSOR CONTROL REGISTER length are supported 16 bits. The 10th bit of sensor control register decides the valid run/read operations. For run operation it should be 0 and for read operation it should be 1.

After programming the ref counter and sensor control register wait for the SIGNAL to go to 0 for starting the run operation. That's called start of operation.

Now for the MODE selection, that's also selected by sensor control register bit, for parallel, the data available on the input port is continuously latched on the sensor control register in every cycle when SIGNAL is 1.

After completion of operation, wait for SIGNAL to go 0 , that indicates the end of run operation. Now the sensor data is available on the counters.

Chapter 4

VRAM - Video Random Access Memory

4.1 REQUIREMENTS

- 4M Byte total memory size.
- 250MHz clock frequency
- Write masking per byte.
- Two fixed read latency settings, 3 and 5 cycles.
- Repairable

4.2 SPECIFICATION

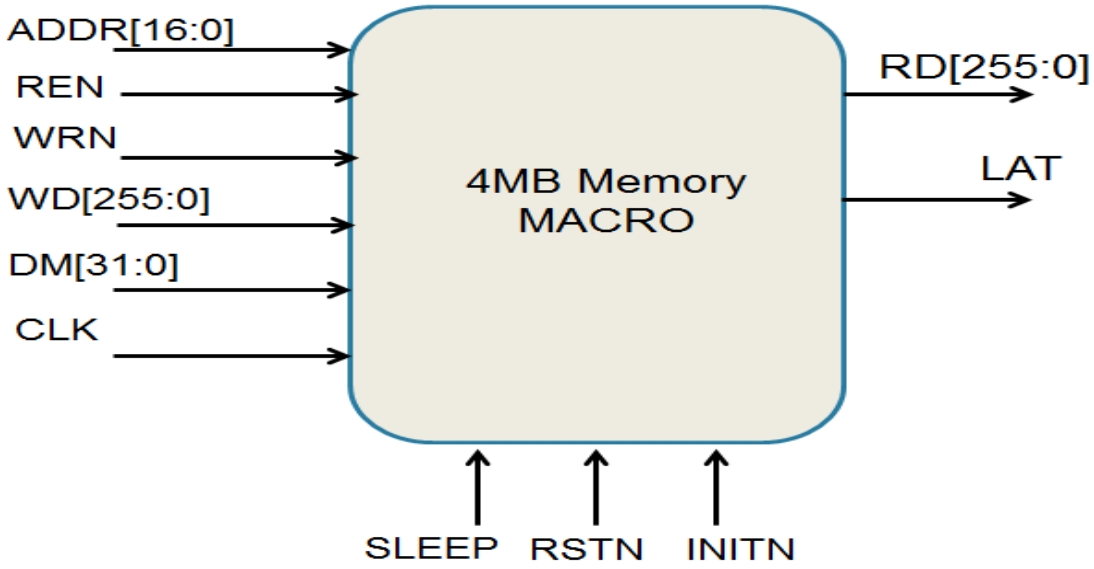

Figure 4.1: Memory Macro

Table 4.1: Memory functional signals

Signal	IN/OUT	Bus Width	Description
CLK	I	1	Max target 250MHZ
LAT	O	1	0=Short latency(3 cycle) , 1=Long latency(5 cycle)
RD	O	256	Read Bus Width
WD	I	256	Write Bus Width
ADDR	I	17	Memory Address Bus
REn	I	1	Read Enable
WRn	I	1	Write Enable
DM	I	32	Data Mask-Use in write operation mask out 8-bit

4.2.1 OPERATING MODE

CK	REN	WRN	DM	SLEEP	RSTN	INITN	RD	LAT	Comments
-	-	-	-	¯\	-	-	X	0	Power up duration
-	-	-	-	0	-	_/	No change	0	Memory Initialisation
-	-	-	-	-	0	-	0	0	Reset mode
_/	-	0	-	0	1	1	No change	0	Write mode no mask
_/	1	0	1	0	1	1	No change	0	Write mode all masked
_/	0	1	-	0	1	1	Mem[Addr]	0	Read mode
_/	1	1	-	0	1	1	No change	0	No-op
_/	0	0	-	0	1	1	No change	0	No-op with error message

Figure 4.2: Operating Mode

- Timescale is 1ns/1ns
- Default latency value is 4. Corresponding LAT output is 0
  - Latency value can be programmed to 5. Corresponding LAT output is 1
  - In the model this can be changed by parameter named VRAM latency
- Caution
  - As mentioned in spec there is no write latency which is also represented in system model. However in the RTL model which will have actual implementation of pipelining and physical memories, if memory array is probed, the contents will be written in next cycle i.e latency of 1. From the VRAM interface Write and Read, this latency will not be visible

- SLEEP, INITN, RSTN are the new pins added
  - For Functional operation it is mandatory that SLEEP=0; INITN=1;RSTN=1
  - If SLEEP goes to 1, RD and memory contents will be corrupted immediately
  - If RSTN goes to 0, output will be 0 immediately. No change in memory contents
  - INITN going low does not impact the memory contents and outputs
- If ADDR = X
  - During valid write, memory content will be corrupted
  - During valid read, RD and memory content both are corrupted
- During valid write If DM[i]= X, and
  - WD[i*8-1:i] = Mem [A] , no change in memory contents
  - WD[i*8-1:i] Mem[A] , memory contents corrupted
- If CK=X, RD is corrupted and memory contents corrupted depending on Read/Write operation
  - After valid Read Cycle, if CK is X in any of next 3 cycles, RD will remain X
- If Read and Write are performed simultaneously (REN=WRN), error message will appear to warn the system
- LAT output is stuck to 0 at t=0 ns
- Default timings between SLEEP, RSTN and INITN (Please refer to power-up sequence on next figure for definition of these parameters)
  - tpw sleep ¿= trampup = 1us

  - tpw reset ¿=1 cycle = 4ns@250MHz
  - tpw initn ¿= 3cycles=12ns @250Mhz
  - tsetup initn ¿= 3 cycles=12 ns@250MHz
  - All these parameters can be overriden from command line to speed up the simulation

- For violation of any timing between SLEEP, INITN and RSTN, a message is displayed with ERROR severity

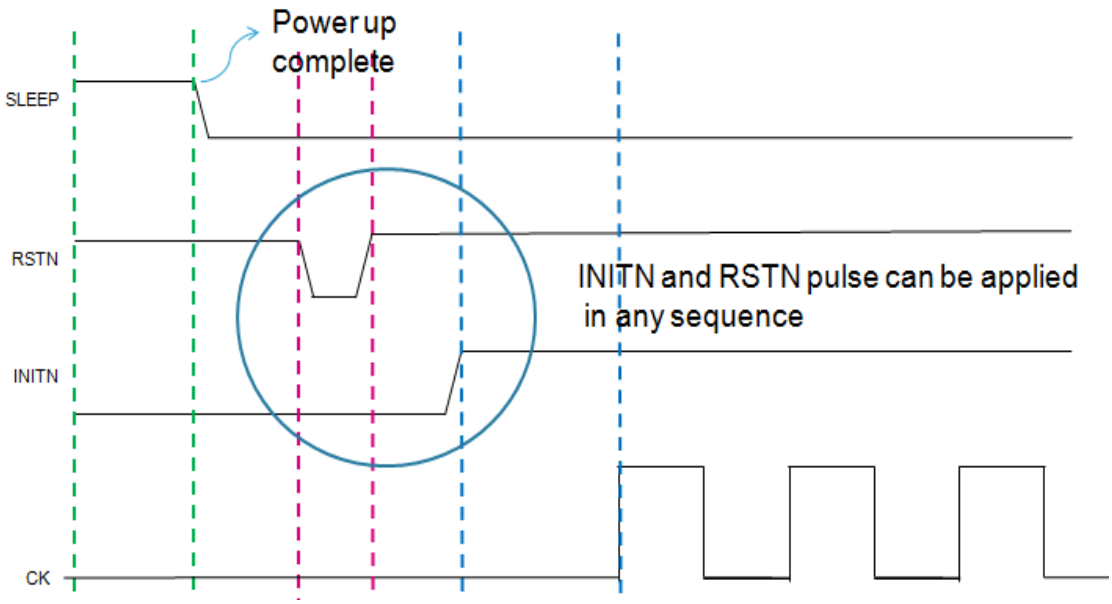

Figure 4.3: Power Up Sequence

Chapter 5

Test Power-Design

Test Power Design is a generic architecture used to check low power design of different IPs.

- It contains a top level switchable power domain (pd testpd) which itself contains.
  - an embedded switchable power domain
  - switchable power domain
  - an always on active area
  - two memory groups. (sram0 and sram1) including each 4 SPHD and a top level BIST
  - AXI bridge to which are connected all the peripherals

POWER INTENT COMMANDS

- create supply port

  port name - A simple name.

  domain domain name-The domain where this port defines a supply net connection point.

  direction in out inout-The direction of the port. The default is in.

- create supply net

  net name - A simple name.

  domain domain name-The domain in whose scope the supply net is to be created.

  reuse-Extend availability of a supply net previously defined for another domain.

- connect supply net

  - ports ports list
  - pg type pg type list element list
  - vct vct name
  - cells cells list
  - domain domain name
  - pins pins list
  - rail connection rail type

  net name-A simple name.

  ports ports list-A list of rooted port names.

  vct vct name-A value conversion table VCT defining how values are mapped from UPF to an HDL model or from the HDL model to UPF.

  domain domain name The domain indicates the scope to use for -pg type or -rail connection.

- create supply set
  - connect supply function pg type list
  - elements element list
  - exclude elements exclude list

  supply set ref-The rooted name of the supply set

  connectsupply function pg type list-Define automatic connectivity for a supply function of the supply set ref as ports having the specified pg type list attributes

  elements element list-The list of instance names to add

  exclude elements exclude-list The list of instances to exclude from the effective element list.

- create power domain
  - simulation only
  - atomic
  - elements element list
  - exclude elements exclude list
  - supply supply set handle [supply set ref]
  - available supplies supply set ref list
  - define func type supply function pg type list
  - update
  - include scope
  - scope instance name

  simulation only-Define a power domain for simulation purposes only.

  atomic-Define the minimum extent of the power domain.

  elements element list-The list of instances to add.

exclude elements exclude list-The list of instances to exclude from the effective element list.

supply-The supply option specifies the supply set handle for domain name.

- create power switch
  - output supply port port name [supply net name]
  - input supply port port name [supply net name]
  - control port port name [net name]
  - on state state name input supply port boolean expression
  - off state state name boolean expression
  - supply set supply set ref
  - on partial state state name input supply port boolean expression

output supply port-The output supply port of the switch and optionally the net where this port connects

input supply port-An input supply port of the switch and optionally the net where this port is connected

control port-A control port on the switch and optionally the net where this control port connects

on state-A named on state, the input supply port for which this is defined, and its corresponding Boolean expression

- set isolation
  - domain domain name
  - elements element list
  - exclude elements exclude list
  - source source domain name — source supply ref

- sink sink domain name — sink supply ref
- diff supply only TRUE — FALSE
- applies to inputs — outputs — both
- applies to clamp 0 — 1 — any — Z — latch — value
- applies to sink off clamp 0 — 1 — any — Z — latch — value
- applies to source off clamp 0 — 1 — any — Z — latch — value
- no isolation
- force isolation
- location self — other — parent — automatic — fanout — fanin — fani-nout — sibling
- clamp value 0 — 1 — any — Z — latch — value
- isolation signal signal list isolation sense high — low

domain domain name-The domain for which this strategy is defined.

elements element list-A list of instances or ports to which the strategy potentially applies

exclude elements exclude list-A list of instances or ports to which the strategy does not apply

source-The rooted name of a supply set or power domain. When a domain name is used, it represents the primary supply of that domain

sink-The rooted name of a supply set or power domain. When a domain name is used, it represents the primary supply of that domain

diff supply only-Indicates whether ports connected to other ports with the same supply should be isolated

applies to inputs — outputs — both-A filter that restricts the strategy to apply only to ports of a given direction

Chapter 6

Work Summary

I have done Unified Power Formate design for different IPs. During that i get to know the internal of different IPs.Like Singe Port High Density,Single Port Registive,Dual Port Registive,Read only Memories and mainly power aware design of different memories.I have run design with cadence simulator-ncsim,mentor simulator-questasim,synopsys simulator-VCS.During that i got information about different tool interpretation regarding UPF and I have done modification in UPF according to Tools and i learn basic concept use in low power and different strategies use in UPF design.During this i also worked on one Soft-macro development and in that my role is system model design and Verification.and also help in periphery development.

Chapter 7

Conclusion

- Multiple techniques are available to the user for effective power management using voltage variation in spatial domains and time or both. However, each technique as adopted and implemented brings forth new challenges in achieving functional correctness

- The path to low power design is paved with many errors that are not conventional logic failures. They are caused by a variety of reasonsby design structure, by faulty control, or faulty architecture. Debug can be hard, especially since a majority of the bugs just end up consuming a lot of power. A combination of static checks, test bench components and focused test vectors is the best way to detect these errors early in the design

- In a nutshell, multi-voltage design brings about significant changes to the way libraries are represented modeled and used. New library attributes and standards,already reflect these changes the key however, is what we mentioned earlier. Consistent interpretation and comprehensive testing of these new attributes is essential for both design and verification processes

Bibliography

[1] http://www.american.edu/cas/hpc/upload/LSF-commands.pdf

[2] http://acms.ucsd.edu/files/slpfug.pdf

[3] http://www.systemverilog.in/lowdesign.php

[4] http://venividiwiki.ee.virginia.edu/mediawiki/index.php/Multi-VDDDesignFlow

[5] EEE Std 1801-2009, IEEE Standard for Design and Verification of Low Power Integrated Circuits.