Improvements in PHY Low Power Analysis and Verification

Major Project Report

Submitted in partial fulfillment of the requirements
for the degree of

Master of Technology
In
VLSI Design
By
Ravindra Kant
(16MECV19)

Under the Guidance of

Syed S.Thameem
Engineering Manager
Intel Technologies Pvt. Ltd.
Bengaluru

Prof. Usha Mehta
Professor Dept. of ECE
Institute of Technology
Nirma University

Department of Electronics & Communication Engineering
Institute of Technology, Nirma University
Ahmedabad - 382481
December, 2017

Declaration

This is to certify that:

i) The report comprises 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.

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

Ravindra Kant
16MECV19

NIRMA UNIVERSITY

INSTITUTE OF TECHNOLOGY

NAAC ACCREDITED 'A' GRADE

Certificate

This is to certify that the Major Project entitled **"Improvements in PHY Low Power Analysis and Verification"** submitted by **Ravindra Kant (16MECV19)**, towards the partial fulfillment of the requirements for the degree of **Master of Technology in VLSI Design, Nirma University, Ahmedabad** is the record of work carried out by him under our supervision and guidance. In our 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 our knowledge,haven't been submitted to any other university or institution for award of any degree or diploma.

External Guide:
Syed S.Thameem
Engineering Manager
Intel Technologies Pvt. Ltd.
Bengaluru

Internal Guide:
Prof. Usha Mehta
Professor Dept. of ECE
Institute of Technology
Nirma University

Dr. N.M. Devashrayee
Professor & PG Co-ordinator
M.Tech. (VLSI Design) Institute of Technology
Nirma University

Dr. D.K. Kothari
Head of Dept. Electronics and Communication Engineering
Institute of Technology
Nirma University

Date: Dec 04, 2017

Acknowledgment

It gives me a great sense of pleasure to present the Major project thesis. I owe special debt of gratitude to Syed S.Thameem for his constant support and guidance throughout the course of my work. His sincerity, thoroughness and perseverance have been a constant source of inspiration for me. It is only his cognizant efforts that our endeavors have seen light of the day.

I also take the opportunity to acknowledge the contribution of Dr. N.M. Devashrayee and Dr Usha Mehta for his full support and assistance during the development of the project.

- Ravindra Kant (16MECV19)

Certificate

This is to certify that the Project entitled ”**Improvements in PHY Low Power Analysis and Verification**” submitted by **Ravindra Kant (16MECV19)**, towards the submission of the Project for requirements for the degree of Master of Technology in VLSI Design, Nirma University, Ahmedabad is the record of work carried out by him under our supervision and guidance. In our opinion, the submitted work has reached a level required for being accepted for examination.

(External Guide)
Mr.Syed S Thameem
Engineering Manager
Intel Technology India
Bangalore

(Mentor)
Mr. Santhosh Vishnumurthy
Digital Design Engineer
Intel Technology India
Bangalore

Company Seal
Intel Technology India Pvt. Ltd.(Bangalore)

Date :

Place : Bangalore

Abstract

The voltages and currents within a chip's logic circuitry have been on a downward trajectory along with node sizes, but the dimensions of the outside world have not changed in the same manner. Wires outside the package are hundreds of times longer, and they have much higher capacitance and resistance. Inside the chip, processing has become faster, and that depends on getting enough memory throughput—or access to the raw data required—which in turn means that the external interfaces must operate at an ever-faster rate.Pushing data at ever-faster rates through boards and systems consumes increasing amounts of power, but the power budget for chips has not been increasing.PHYs need high drive capability and often contain circuits that try to compensate for the bad things that happen between the chip and the other end. That means they are also power hungry.A third to a half of the total power is consumed just for the PHYs.

This thesis report aims at providing the current improvements in the area of efficient low power design.As the technology size shrinks it brings a lot of complexity in the analysis and verification of the design process. The use of innovative design techniques to overcome leakage has created new challenges for verification that demand a creative response.No individual tool can sufficiently verify all the tricky issues engendered by today's low-power techniques.The report also deals with test case where a raw design is taken from staic verification and power exploration stage to the sign off stage employing the power checks and strategies defined to bring out a clean power efficient design.

Contents

List of Figures

Chapter 1

Introduction

With the escalation of multiple PHY in consumer mobile devices,the design teams look for key criteria when licensing IP, such as cost, system performance (interoperability), reliability, and power.So hard core technical evaluation is considered to be the key part for make versus buy decision for all but simplest IP cores.

As the advancement continues the design cycle is constrained with the competitive dynamics of global consumer markets including costs.Consequently the overall design productivity and cost of IP ownership must also be considered for e.g problem of reliability,manufacturing yield and interoperabilty can have profound effect on the total cost of ownership.Interoperabilty is the function of design specification and operating margin which in turn can impact the device yield. With the rise in demands for specification on power, driven by the need for longer battery life,a low power IP enables the SoC power budget to be maintained -a critical issue considered to be most important for battery powered devices like smart phones,digital cameras,flash drives etc. Reducing the power demands of the PHY by up to 50% not only extends battery life but also may mean that a lower cost power supply can be used. This is an important issue in portable, battery-operated products.The low-power architecture has other benefits.

By reducing the supply current requirement, the overall power consumption is lowered and enables the pin count to be minimized (by half) without sacrificing any functionality. The ultra-low pin count design is a major advantage in terms of enabling the use of lower-cost packaging. Alternatively, package pins can be made available for other signals. The need for fewer pins also reduces the cost of production test, as well as considerably easing SoC integration.

1.1 Low Power Design : An Overview

In the past, due to a high degree of process complexity and the exorbitant costs involved, low-power circuit design and applications involving CMOS technologies were used only in applications where very low power dissipation was absolutely essential, such as wrist watches, pocket calculators, pacemakers, and some integrated sensors. However, low-power design is becoming the norm for all high-performance applications, as power is the most important single design constraint. Depending upon the target application minimizing the overall power dissipation in a system has become a highest priority.

The consideration of portability is due to numerous factors

- First,the size and weight of the battery pack is fundamental. A portable system that has an unreasonably heavy battery pack is not practical and restricts the amount of battery power that can be loaded at any one time.

- Second, the convenience of using a portable system relies heavily on its recharging interval. A system that requires frequent recharging is inconvenient and hence limits the user's overall satisfaction in using the product.

Although the battery technology has improved over the years, its capacity has only managed to increase by a factor of two to four. the computational power of digital integrated circuits has increased by more than four orders of magnitude.For an e.g. case consider a multimedia terminal that supports high-bandwidth wireless communication; bi-directional motion video; high-quality audio, speech, and pen-based input; and full texts/graphics.The power of such a terminal—when implemented using off-the-shelf components not designed for low power—is projected to reach approximately 40 W. Based on the current Nickel-Cadmium (NiCd) battery technology, which offers a capacity of 20 W-hour/pound, a 20-pound battery pack is required to stretch the recharge interval to 10 hours. Even with new battery technologies, such as the rechargeable lithium or polymers, battery capacity is not expected to improve by more than 30 to 40% over the next 5 years. Hence, in the absence of low-power design techniques, future portable products will have either unreasonably heavy battery packs or a very short battery life.The rising power dissipation has been inevitably increased due to rapidly increasing packing density,clock frequency and computational power.The trends relating to the power consumption of microprocessors indicate that power has increased almost linearly with area-frequency product over the years.

For example, the DEC21164, which has a die area of 3 cm^2 and runs on a 300-MHz clock frequency, dissipates as much as 50 W of power. Such high power consumption requires expensive packaging and cooling techniques given that insufficient cooling leads to high operating temperatures, which tend to exacerbate several silicon failure mechanisms. To maintain the reliability of their products, and avoid expensive packaging and cooling techniques, manufacturers are now under strong pressure to control, if not reduce, the power dissipation of their products.

Finally, due to the increasing percentage of electrical energy usage for computing and communication in the modern workplace, low-power design is in line with the increasing global awareness of environmental concerns. As a result, power has emerged as one of the most important design and performance parameters for integrated circuits. Only a few years ago, the power dissipation of a circuit was of secondary importance to such design issues as performance and area. The performance of a digital system is usually measured only in terms of the number of instructions it can carry out in a given amount of time, that is, its throughput. The area required to implement a circuit is also important as it is directly related to the fabrication cost of the chip. Larger die areas lead to more expensive packaging and lower fabrication yield. Both effects translate to higher cost. Because the performance of a system is usually improved at the expense of silicon area, a major task for integrated chip (IC) designers in the past was to achieve an optimal balance between these two often-conflicting objectives. Now, with the rising importance of power, this balance is no longer sufficient. Today, IC designers must design circuits with low-power dissipation without severely compromising the circuits' performance.

Clearly, power has become a major consideration in VLSI and giga-scale-integration (GSI) engineering due to portability, reliability, cost, and environmental concerns. With the increasing integration levels, energy consumption has become one of the critical design parameters. Consequently, much effort has to be put in achieving lower dissipation at all levels of the design process. It was found that most low-power research is concentrated on components research: better batteries with more power per unit weight and volume; low-power CPUs; very low-power radio transceivers; low-power displays. We found that there is very little systems research on low-power systems. While these low-level circuit and logic techniques have been well established for improving energy efficiency, they do not hold promise for much additional gain. While low-power components and subsystems are essential building blocks for portable systems, a system-wide architecture that incorporates the low-power vision into all layers of the system is beneficial because there are dependencies between subsystems, e.g. optimization of one subsystem may have consequences for the energy consumption of other modules.The key to energy efficiency in future mobile systems will be designing higher layers of the mobile system, their system architecture, their functionality, their operating system, and indeed the entire network, with energy efficiency in mind. Furthermore, because the applications have direct knowledge of how the user is using the system, this knowledge must be penetrated into the power management of the system.

1.2 PHY layer

The physical layer or PHY layer is the heart of any interconnect solution.They have been derived into three high performance and cost optimized physical layers referred as D-PHY,C-PHY,M-PHY.MIPI D-PHY is used primarily to interconnect cameras and displays to an application processor. MIPI M-PHY supports multimedia and chip-to-chip/interprocessor communications. MIPI C-PHY supports cameras and displays. Each PHY meets very rigorous requirements for high performance, low-power operation and low electromagnetic interference (EMI) interfaces.for e.g D-PHY was developed primarily to support camera and display interconnections in mobile devices, and it has become the industry's primary high-speed PHY solution for these applications in smartphones today It meets the demanding requirements of low power, low noise generation, and high noise immunity that mobile phone designs demand.D-PHY is a flexible specification that gives designers a lot of choice when selecting specifications to meet their specific product requirements. Designers who are building devices to serve prevalent camera and design trends in the ecosystem will find that D-PHY is the preferred solution that meets most of the market's needs. The comparative simplicity of D-PHY compared to M-PHY can also reduce the complexities and costs of implementation

- Compatible with advanced CMOS process.
- clock forwarding topology along with separate clock lane alongside data lanes.
- Half duplex protocol that data flows only in one direction at a time.
- A typical application has data largely flowing in one direction only for e.g camera sensor to the application processor or from the application processor to the display panel.

1.3 Power Analysis

Power analysis is an estimation of power dissipation, both dynamic and static, of the chip in various operating modes. IR drop analysis deals with the chip's current draw and the associated voltage drop across the power grid, power switches, etc. Since gate delay depends greatly on the applied voltage, it is very important to make sure that a sudden current draw does not reduce the voltage and slow down the gate to the point of circuit failure.Static power analysis is the calculation of leakage power. A cell dissipates leakage power when voltage is applied even if it is not switching. As the process geometry shrinks, leakage power is becoming a greater percentage of a chip's overall power dissipation. It is something that we cannot ignore. Dynamic power consists of power dissipated inside a cell (mostly due to short-circuit current during switching) and power dissipated to charge/discharge net capacitance. Dynamic power is a function of voltage, toggle rate, and net loading.

For power analysis, each cell's power dissipation has been characterized in the library (.lib) file. For leakage power, the EDA tool simply adds up the leakage power of each cell. (Note: Leakage power is usually state dependent, so there is a bit of work here.) For dynamic power, the EDA tool either estimates net capacitance before P&R or calculates net capacitance after P&R. The designer has to provide the toggle rate. This can be based on educated guess, experience, simulation, or emulation. The accuracy of the power analysis depends directly on the accuracy of net capacitance and toggle rate.

Power analysis must be considered very early in the design cycle. Typically, 80% of a chip's power is determined at the RTL stage. After that, a design team can only impact 20% of the power. Here are some of the questions that a design team should answer at the architectural stage:

- Which voltage supply should we use?
- Can we achieve lower power with more than one voltage supply?
- Do we have inactive blocks that we can shut off to reduce leakage power?
- If we shut off blocks, are there registers that we have to retain the state?
- Do we have blocks that can run at slower rate in certain modes? Can we reduce the voltage during those modes?

Chapter 2

Power Modelling and Analysis

2.1 Design Flow

A system constitutes various level of abstraction.When a design is optimized for power it should embody all levels of abstraction.It is basically done at three levels,system ,logic and technological levels.or example, at the system level power management can be used to turn off inactive modules to save power, and parallel hardware may be used to reduce global interconnect and allow a reduction in supply voltage without degrading system throughput. At the logic level asynchronous design techniques can be used. At the technological level several optimization can be applied to chip layout, packaging and voltage reduction.Given a design specification, a designer is faced with several choices at different levels of abstraction. The designer has to select a particular algorithm, design or use an architecture that can be used for it, and determine various parameters such as supply voltage and clock frequency. This multidimensional design space offers a large range of possible trade-offs. At the highest level the design decisions have the most influence.Therefore,the most effective design decisions derive from choosing and optimizing architectures and algorithms at the highest levels.However, when designing a system it is a problem to predict the consequences and effectiveness of high level design decisions because implementation details can only be accurately modelled or estimated at the technological level and not at the higher levels of abstraction. Furthermore, the specific energy reduction techniques that are offered by the lower layers can be most effective only when the higher levels are aware of these techniques, know how to use them, and apply them.

2.2 CMOS Component Model

Most components are currently fabricated using CMOS technology. Main reasons for this bias is that CMOS technology is cost efficient and inherently lower power than other technologies. The sources of energy consumption on a CMOS chip can be classified as static and dynamic power dissipation. The main difference between them is that dynamic power is frequency dependent, while static is not.Bias Pb and leakage energy Pl cause static energy consumption.Short circuit currents Psc and dynamic energy consumption Pd is caused by the actual effort of the circuit to switch.

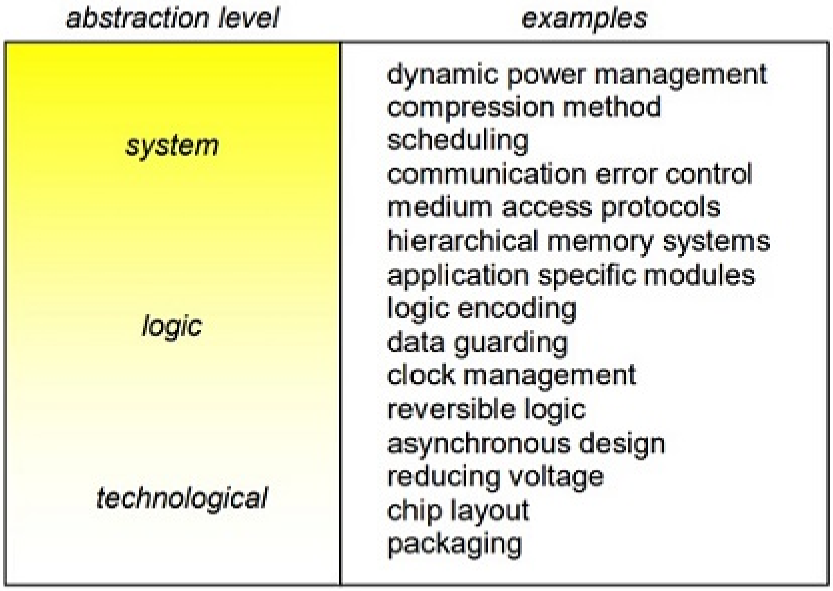

abstraction level	examples
	dynamic power management
	compression method
system	scheduling
	communication error control
	medium access protocols
	hierarchical memory systems
	application specific modules
logic	logic encoding
	data guarding
	clock management
	reversible logic
	asynchronous design
	reducing voltage
technological	chip layout
	packaging

Figure 2.1.1: General Design flow and related examples of energy reduction

$$P = Pb + Pl + Psc + Pd$$

The contributions of this static consumption are mostly determined at the circuit level.Leakage currents also dissipate static energy, but are also insignificant in most designs (less than 1%). In general we can say that careful design of gates generally makes their power dissipation typically a small fraction of the dynamic power dissipation, and hence will be omitted in further analysis.

Dynamic power can be partitioned into power consumed internally by the cell and power consumed due to driving the load. Cell power is the power used internally by a cell or module primitive,for example a NAND gate or flip-flop. Load power is used in charging the external loads driven by the cell, including both wiring and fanout capacitance's.So the dynamic power for an entire chip is the sum of the power consumed by all the cells on the chip and the power consumed in driving all the load capacitance's.During the transition on the input of a CMOS gate both p and n channel devices may conduct simultaneously, briefly establishing a short from the supply voltage to ground (Icrowbar).This effect causes a power dissipation of approx. 10 to 15%.

The more dominant component of dynamic power is capacitive power. This component is the result of charging and discharging parasitic capacitances in the circuit. Every time a capacitive node switches from ground to Vdd an vice-versa energy is consumed.The dominant component of energy consumption (85 to 90%) in CMOS is therefore dynamic.A first order approximation of the dynamic energy consumption of CMOS circuitry is given by the formula:

$$Pd = CeffV^2 f$$

where Pd is the power in Watts, Ceff is the effective switch capacitance in Farads,V is the supply voltage in Volts,and f is the frequency of operations in Hertz.The power dissipation arises from the charging and discharging of the circuit node capacitance found on the output of every logic gate. Every low-to-high logic transition in a digital circuit incurs a voltage change

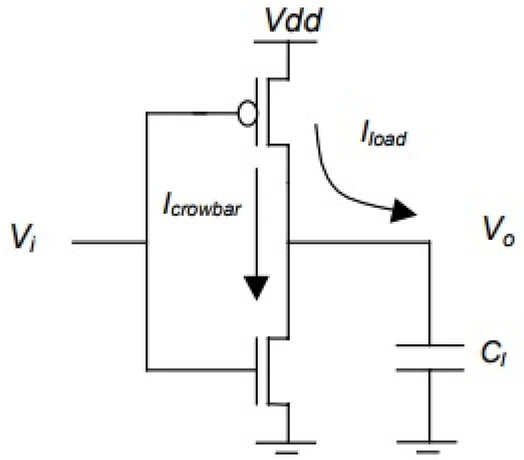

Figure 2.2.1: Dynamic power in CMOS inverter

ΔV, drawing energy from the power supply.Ceff combines two factors C,the capacitance being charged/discharged, and the activity weighting α, which is the probability that a transition occurs.

$$\text{Ceff} = \alpha\text{C}$$

The search for the optimal solution must include, at each level of abstraction, a design improvement loop. In such a loop a power analyzer/estimator ranks the various design, synthesis, and optimization options, and thus helps in selecting the one that is potentially more effective from the energy consumption standpoint. Obviously, collecting the feedback on the impact of the different choices on a level-by-level basis, instead of just at the very end of the flow (i.e. at the gate level), enables a shorter development time. On the other hand, this paradigm requires the availability of power estimators, as well as synthesis and optimization tools, that provide accurate and reliable results at various levels of abstraction.Power analysis tools are available primarily at the gate and circuit levels, and not at the architecture and algorithm levels where they could really make an impact.

The technological level comprises the technology level, dealing with packaging and process technologies, the layout level that deals with strategies for low-power placement and routing, and the circuit level that incorporates topics like asynchronous logic and dynamic logic.

2.2.1 Technology scaling

The process technology has been improved continuously, and as the SIA roadmap indicates, the trend is expected to continue for years Scaling of the physical dimension involves reducing all dimensions: thus transistor widths and lengths are reduced, interconnection length is reduced, etc. Consequently, the delay, capacitance and energy consumption will decrease substantially. For example, MIPS Technologies attributed a 25% reduction in power consumption for their new processor solely to migration from 0.8 µm to 0.64 µm process.Another way to reduce capacitance at the technology level is thus to reduce chip area. However, note that a sole reduction in chip area at architectural level could lead to an energy-inefficient design. For example, an energy efficient architecture that occupies a larger area can reduce the overall energy consumption, e.g. by exploiting locality in a parallel implementation.

2.2.2 Chip Layout

There are a number of layout-level techniques that can be applied. Since the physical capacitance of the higher metal layers are smaller, there is some advantage to select upper level metals to route high-activity signals. Furthermore, traditional placement involves minimizing area and delay, which in turn translates to minimizing the physical capacitance (or length) of wires. Placement that incorporates energy consumption, concentrates on minimizing the activity-capacitance product rather than capacitance alone. In general, high-activity wires should be kept short and local. Tools have been developed that use this basic strategy to achieve about 18% reduction in energy consumption.

2.2.3 Capacitance reduction

The capacitance is an important factor for the energy consumption of a system. However, reducing the capacity is not the distinctive feature of low-power design, since in CMOS technology energy is consumed only when the capacitance is switched. It is more important to concentrate on the switching activity and the number of signals that need to be switched. Architectural design decisions have more impact than solely reducing the capacitance.

2.2.4 Reduce voltage and frequency

One of the most effective ways of energy reduction of a circuit at the technological level is to reduce the supply voltage, because the energy consumption drops quadratically with the supply voltage. For example, reducing a supply voltage from 5.0 to 3.3 Volts (a 44% reduction) reduces power consumption by about 56%. As a result, most processor vendors now have low voltage versions. The problem that then arises is that lower supply voltages will cause a reduction in performance. In some cases, low voltage versions are actually 5 Volt parts that happen to run at the lower voltage. In such cases the system clock must typically be reduced to ensure correct operation. Therefore, any such voltage reduction must be balanced against any performance drop. To compensate and maintain the same throughput, extra hardware can be added. This is successful up to the point where the extra control, clocking and routing circuitry adds too much overhead . In other cases, vendors have introduced true low voltage versions of their processors that run at the same speed as their 5 Volt counterparts. The majority of the techniques employing concurrency or redundancy incur an inherent penalty in area, as well as in capacitance and switching activity. If the voltage is allowed to vary, then it is typically worthwhile to sacrifice increased capacitance and switching activity for the quadratic power improvement offered by reduced voltage.

The variables voltage and frequency have a trade-off in delay and energy consumption. Reducing clock frequency f alone does not reduce energy, since to do the same work the system must run longer. As the voltage is reduced, the delay increases. A common approach to power

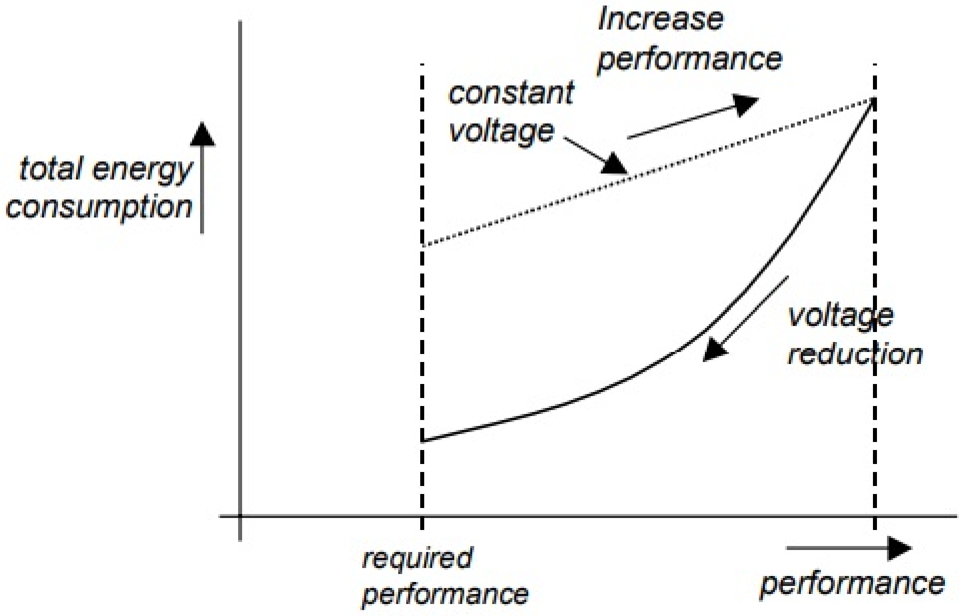

Figure 2.2.2: Impact of voltage scaling and performance to energy consumption

reduction is to first increase the performance of the module – for example by adding parallel hardware – , and then reduce the voltage as much as possible so that the required performance is still reached.Therefore, major themes in many power optimization techniques are to optimize the speed and shorten the critical path, so that the voltage can be reduced. These techniques often translate in larger area requirements, hence there is a new trade-off between area and power.The main limitation of all voltage scaling approaches is that they assume the designer has the freedom of choosing the voltage supply for the design. Unfortunately, for many real-life systems, the power supply is not a variable to be optimized.

2.2.5 Avoid Unnecessary Activity

The capacitance can only marginally be changed and is only important if switched, the voltage is usually not under designer's control, and the clock frequency, or more generally, the system throughput is rather a constraint than a design variable. The most important factor contributing to the energy consumption is the switching activity. Actually, once the technology and supply voltage have been set, major energy savings come from the careful minimization of the switching activity α While some switching activity is functional, i.e. it is required to propagate and manipulate information, there is a substantial amount of unnecessary activity in virtually any digital circuit Unnecessary switching activity is due

- spurious transitions due to unequal propagation delays. transitions occurring within units that are not participating in a computation .
- whose computation is redundant.

Reordering of logic inputs to circuits can have significant energy consumption consequences.For e.g two functional identical circuits, but with a different energy consumption due to the differ-

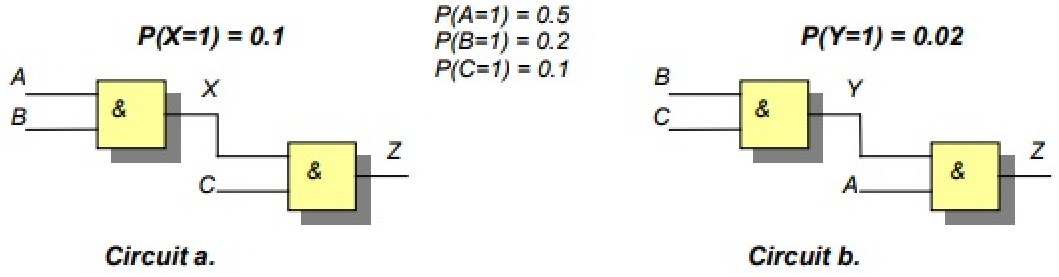

Figure 2.2.3: Reordering Logic Inputs

ent signalling activity. The normalized energy consumption equals 0.11 of circuit a, and 0.021 for circuit b.Thus, much energy can be saved by minimizing the amount of switching activity needed to carry out a given task within its performance constraints.

2.3 Low-power logic-level design

At the logic level, opportunities to economize on power exist in both the capacitance and frequency spaces. The most prevalent theme in logic-level optimization techniques is the reduction of switching activities.

2.3.1 Cell Library

The choice of the cell library to use for a chip design provides the first obvious opportunity to save energy. Standard cells have lower input capacitances than gate arrays because they use a variety of transistor sizes. For the same reason, the cells themselves consume less power when switching. Using libraries designed for low power can also reduce capacitance. These libraries contain cells that have low-power micro architectures or operate at very low voltages. Some of the leading application-specific IC (ASIC) vendors are providing such libraries today, and many captive design groups are producing specialized libraries for low-power applications. But no matter which type of library is utilized, the logic designer can minimize the power used by each cell instance by paying careful attention to the transition times of input and output signals. Long rise and fall times should be avoided in order to minimize the crowbar current component of the cell power.

2.3.2 Clock Gating

Several power minimization techniques work especially well at the logic level. Most of them rely on switching frequency. The best example of which is the use of clock gating. Because CMOS power consumption is proportional to the clock frequency, dynamically turning off the clock to unused logic or peripherals is an obvious way to reduce power consumption. In clock gating, a control signal enables a clock signal so that the clock toggles only when the enable signal is true, and is held steady when the enable signal is false. Gated clocks are used, in power management, to shut down portions of the chip, large and small, that are inactive. This saves on

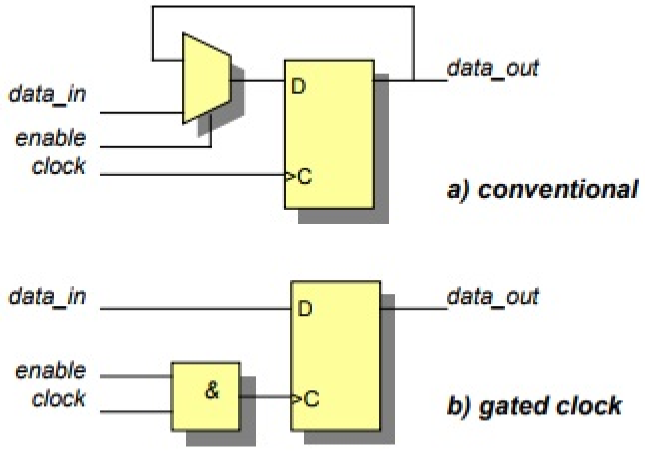

Figure 2.3.1: Clock gating

clock power, because the local clock line is not toggling all the time. Consider the case of a data bus input register as depicted in Figure 7. With the conventional scheme, the register is clocked all the time, whether new data is to be captured or not. If the register must hold the old state, its output is fed back into the data input through a multiplexer whose enable line controls whether the register clocks in new data or recycles the existing data. With a gated clock, the signal that would otherwise control the select line on the multiplexer now controls the gate. The result is that the energy consumed in driving the register's clock input is reduced in proportion to the decrease in average local clock frequency. The two circuits function identically, but utilization of the gated clock reduces the power consumption.Clock gating can be implemented locally by gating the clocks to individual registers, or globally, by building the gating structures into the overall architecture to turn off large functional modules. While both techniques are effective at reducing energy, global gating results in much larger energy reductions and is often used in implementing power-down and power-management modes. Some processors and hardware devices have sleep or idle modes. Typically they turn off the clock to all but certain sections to reduce power consumption. While asleep, the device does no work. Control can be done at the hardware level or the operating system or the application can manage it.

Chapter 3

Power Management and Verification Considerations

Design having special cells like isolation,level shifter,retention cells and power switches used to control the logic's of a particular power domain in IP block should be verified ,that whether they are instrumented correctly by the synthesis tool and correctly controlled by the power management block.

3.0.1 Isolation Cells

Isolation cells that are placed between an off domain and an on power domain.The main advantage of placing this cell in the design is make the signal predictable as the signal travelling from off to on domain might carry and unpredictable value 'X' which could cause unnecessary failure in the design and that can be avoided by using this special cells.The cells provided in the standard cell library has an special library attribute related to it specifying it as an isolation cell.In general AND and OR type isolation cells are used which can clamp the value to zero or one depending upon the condition.

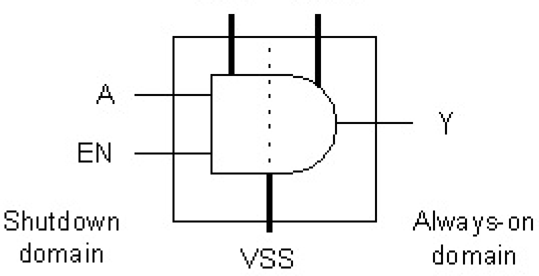

Figure 3.0.1: Isolation Cell

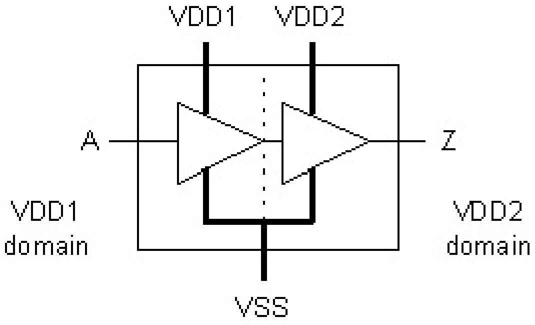

Figure 3.0.2: Level Shifter

3.0.2 Level Shifters

In a multi voltage design, a level shifter is required where each signal crosses from one power domain to another. The level shifter operates as a buffer with one supply voltage at the input and a different supply voltage at the output. Thus, a level shifter converts a logic signal from one voltage swing to another, with a goal of having the smallest possible delay from input to output.The library description of a level-shifter cell must have information about the type of conversion performed (high-to-low,low-to-high, or both), the supported voltage levels, and the identities of the respective power pins that must be connected to each power supply.

3.0.3 Retention Registers

In a design with power switching, there are several different ways to save register states before power-down and restore them upon power-up in the power-down domain. One method is to use retention registers, which are registers that can maintain their state during power-down by means of a low-leakage register network and an always-on power supply.The library description of a retention register specifies the power pins and the input signals that control the saving and restoring of data. It also specifies which power pins are normal and can be powered down and which ones are the always-on pins used to maintain the data during power-down.

3.0.4 Power Switches

Power switches control the flow of electrical current from the supply to the power domains. They are managed through signals sourced from a power management authority within the SoC. It should be noted that power switches establish the flow of current almost immediately after receiving a request, but the beneficiary power domain might not become operational so quickly. Depending on the electrical characteristics of the beneficiary power domain, there is a finite amount of time before the input supply voltage reaches an optimum operational threshold. Care must be taken when accessing the newly powered-on power domain.n a design with power switching, either header or footer type power switch cells are required to supply power for cells that can be powered down. A header type power switch connects the power rail to the power

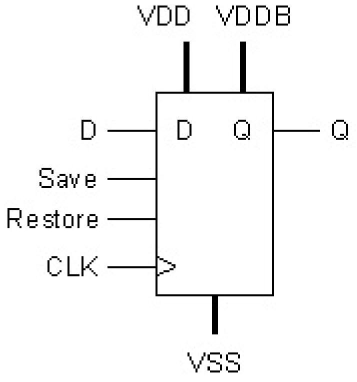

Figure 3.0.3: Retention Cell

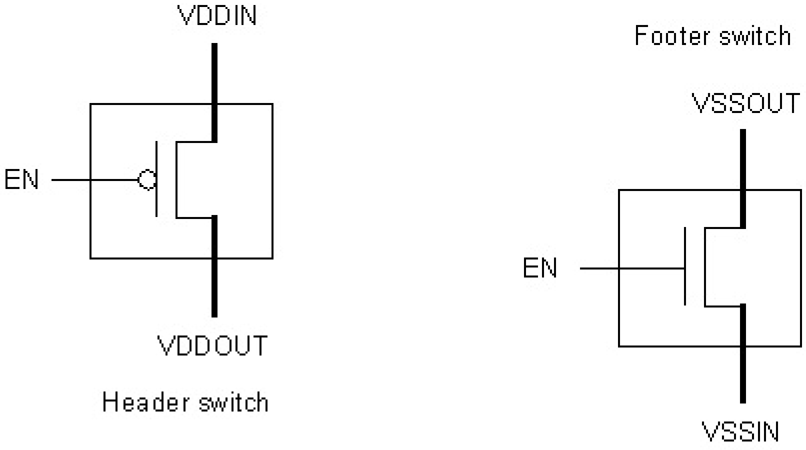

Figure 3.0.4: Power Switch

supply pins of the cells in the power-down domain. A footer type power switch connects the ground rail to the ground supply pins of the cells in the power-down domain. An input logic signal to the power switch controls the connection or disconnection state of the switch

3.0.5 Power supply network

The concept of power domains and multi-voltage creates a need for a network and hierarchy of power lines, switches, and on-chip and off-chip power regulators. Since some of the power lines can be controlled to supply different voltage levels at different times, care must be taken when coordinating the voltage levels with clock frequencies for a given power state

3.1 Power Estimation and Exploration

The goal here is to determine the power estimation at a very early RTL stage to reduce power at the subsequent stages.Estimation and exploration of power results are very closely linked and calibrated with the tool to give excellent correlation with gate level and silicon estimates when setup correctly.

Currently at an very early stage three classes of power estimations are defined :

Vector less This method can be used in very first phase of design when all the meaningful simulation vectors are not available for estimation.Probabilistic activity on input ports and definition of clocks through SDC constraints.The activity is propagated to get the accurate power estimate.This approach may not be fully appropriate for power reduction but it provides the earliest power trends in the design flow.

Exploratory It gives much better accurate with the activity and hence the power estimate.For e.g comparing power reduction opportunities with the what-if analysis if design changes.This stage do not care about the exact correlation with the gate level or silicon power results.

Absolute Power Analysis This method relies on simulation results and aims to provide accurate power correlation with gate level analysis and silicon estimates.More attention is paid towards clock tree modelling and input drive and output load modelling.

Also there are two classes of power reduction

reductions associated with the register gating

reduction associated with memory gating using both of them leads to significant power savings.

3.1.1 Perquisites for power estimation and exploration

Correctly setting up design and ensuring the analyzed design contains minimum unintended black boxes.

Include appropriate technology library (.lib) files, characterized for power, in your analysis. Read Specifying Functionality Information through .lib File.Access to all .lib files for PVT corners and Vt options should be included in the analysis.

If you plan to estimate power based on simulation activity, you should have access to simulation activity files (VCD, FSDB or SAIF) corresponding to the same RTL you plan to analyze.

If you are using power format files (UPF or CPF) to specify power management, you should have access to those files corresponding to the same RTL you plan to analyze.

3.1.2 Horizontal Simulation Flow

a simulation file may not be available for the complete top-level design. However, it may be available for each block in the design. In such cases, estimate the power for the complete design by reading the power of individual blocks from their respective simulation files. This flow is also known as Horizontal Simulation File Flow.

For example, the following can be the simulation files for the three block instances of the design, TOP:

B1.vcd for instance B1

B2.vcd for instance B2

B3.vcd for instance B3

```
]current_design TOP
activity_data -format vcd -file B1.vcd  -instname Top.B1
activity_data -format vcd -file B2.vcd  -instname Top.B2
activity_data -format vcd -file B3.vcd  -instname Top.B3
```

Here, -instname field in the activity_data constraint refers to the design hierarchy for which the simulation file is specified. If you do not specify this field, the simulation file is specified for the complete design.In Horizontal Simulation File Flow, you can also specify the start time and end time of the simulation file analysis. The recommendation is to use the same time window for all the simulation files in this flow. Then, the average activity of signals in each simulation file is calculated independently and scaled uniformly.

3.2 Verification of Power Intent

3.2.1 UPF

The IEEE 1801 Standard for Design and Verification of Low Power Integrated Circuits, also known as the Unified Power Format (UPF), consists of a set of Tcl-like commands used to specify the design intent for multi voltage electronic systems. Using UPF commands, you can specify the supply network, switches, isolation, retention, and other aspects of power management for a chip design. A single set of low-power design specification commands can be used throughout the design, analysis, verification, and implementation flow.

3.2.2 Verification

While functional intent is expressed in RTL and defines the architecture, application, and usage of standard interfaces for the design, power intent is captured in IEEE 1801 and defines the power domains, supply rails, power state strategy, operating voltages, isolation cells, level shifters, power switches, and retention registers. The power intent has a great influence on the functionality of the SoC and as such should be defined alongside the functional intent. Perhaps most, if not all, of the functional-verification tests should become power aware. Without incorporating IEEE 1801 into the flow, a great many features of power-management measures will remain un-tested and potentially leave the resultant silicon SoC with sometimes disastrous bugs.

- An complex SoC works on multiple power domain and if the domain are working on multiple votages than .it often happens a signal changes it domain and may propagate an unpredictable value.
- Moreover if during the power down of domain left the retention pin disabled will cause hard to find system bugs.

3.2.3 Role of Application Software in Verification

It is impractical to manually create test cases for all possible corner cases identifiable for all of the SoC's features in functional and power intents. The impracticality is partially due to the fact that most often in some corner cases and during the occurrence of certain events, the use of a feature is not obvious and remains unknown until the silicon SoC is under the control of the application software.However, a few problematic factors might prevent the verification engineers from using or delay the use of the actual application software for verification:

- Due to the simulation speed, the execution of the application software in simulation is impractical. However,it becomes much more practical at emulation speeds. Especially if the application software is stripped of some of its features that might not have an influence on the usage of hardware. For example, the decompression of software or data before loading into memory can be avoided by storing it as decompressed image.

- The actual application software might not be available at the time of RTL verification, usually due to unavailability of drivers for the new hardware. Sometimes, the basic framework of an older version of the application software combined with the drivers for the new hardware will be a very close fit.

- Unavailability of peripheral models that are capable of interacting with the application software. Fortunately, this big task can be overcome because it is not a technical issue. The good news is that, when available, the models can be reused for other projects.

3.3 Verification

The verification of the power intent specified in IEEE 1801 UPF format is checked at various stages.

- *Power Intent Consistency Checks* Checking of syntax and semantics checks at the first stage of implementation.

- *Signal Corruption Checks*

- *Structural Checks* Verification of special cells used in design like isolation cells,level shifter cells,retention cells and power switches.

- *Power and Ground Checks* Checking the PG consistency against the power network routing on physical netlsts.

- *Functional Checks* Verifies the functionality of Isolation cells and power switches.

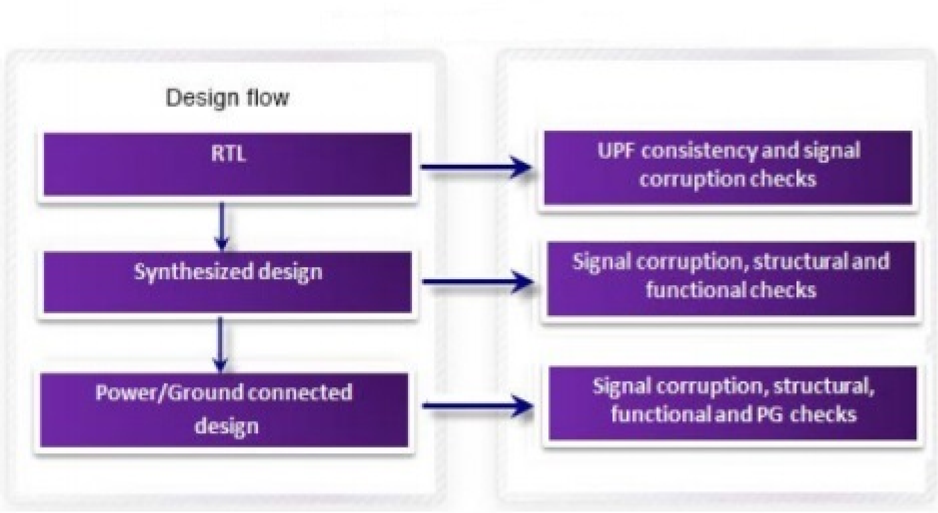
Design flow
RTL
Synthesized design
Power/Ground connected design
UPF consistency and signal corruption checks
Signal corruption, structural and functional checks
Signal corruption, structural, functional and PG checks

Figure 3.3.1: Low power checks flow

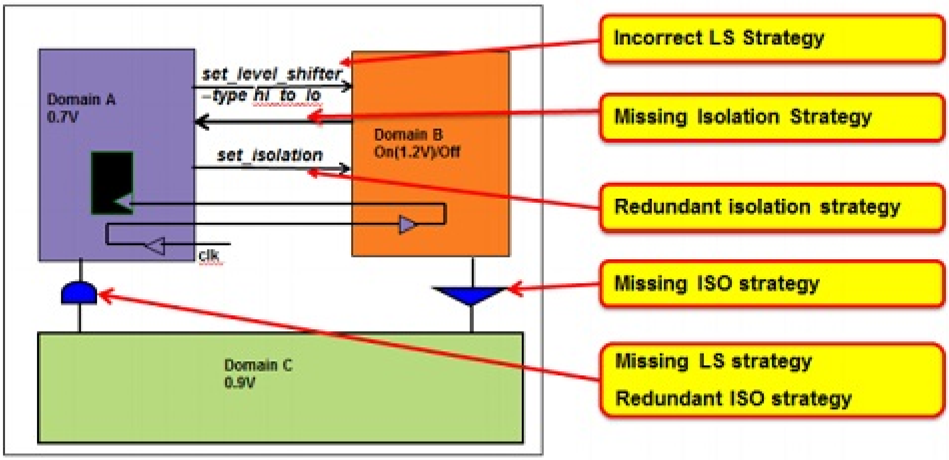

Figure 3.3.2: power intent checks

Chapter 4

Power Optimization and Analysis

4.1 Toggle Activity Estimation

4.1.1 Overview

In CMOS technologies, the chip components draw power supply current only during a logic transition if we ignore the small leakage current. The current is also proportional to the supply voltage value seen by the cell or macro. While this is considered an attractive low-power feature of these technologies, it makes the power estimation and voltage drop highly dependent on the switching activity inside circuits. It means, a more active circuit will consume more current and hence will contribute higher Voltage drop. The activity of circuit is known by running simulation patterns and analyzing the data. The pattern-dependence problem is serious. Often, the power of a functional block needs to be estimated when the rest of the chip has not yet been designed, or even completely specified. In such a case, very little may be known about the inputs to this functional block, and complete and specific information about its inputs would be impossible to obtain. This drives pattern independent toggle activity estimation problem, often referred as vector less approach. Since vector less approach does not require patterns, it is also called 'static' whereas vector based approach is called 'dynamic.

This work describes the approach used for toggle frequency estimation and its limitations.Further it proposes solution to handle these limitations which makes the approach usable for big designs.

Static	Dynamic
Uses probabilistic approach or zero delay simulation based approach.	Uses Logic simulation to generate switching activity or SPICE simulation to calculate power.
Vector-less approach	Vector based approach. Hence quality is as good as input vectors. Imagine number of patterns possible for 100 inputs block.
Modeling of certain element (hard macro/complex block) is difficult.	Since it is vector based, functional models can be used during simulation.
Lot of research into products for average power estimation	Can give instantaneous power.

Transition Density : If a logic signal x(t) makes n(T) transitions in a time internal of length T, then the transition density of x(t) is defined as

D(x)= n(T)/T , where T is very huge time (infinite ideally).

For large T, D(x) becomes time invariant function and hence there is no need to account for temporal correlation.
Toggle Frequency: If a node x is toggling n(T) times over a time interval of length T, then the toggle frequency F(x) is defined as F(x) = n(T)/(2*T) where T is very huge time (infinite ideally)
Example- if the node is switching at 20 MHz, it is expected that the node will switch 2 times in 50 ns. As it can be seen, the toggle frequency can be converted to transition density or switching activity by the following equation, Toggle density = No of transitions/Period = Switching Activity

It should be noted that toggle frequency of a node has no direct relation with the clock domain(s) in which node (or logic) exists. We have used the clock domain frequency to upper bound the toggle frequency calculated by our approach.

Signal Probability: Signal probability P(x) at a node x is defined as the average fraction of clock period in which the stead state value of x is logic high.

4.1.2 Toggle Activity

Boolean difference of output is computed with respect to each input pin. Boolean difference of function y (output) depends on x(each of the input). It is defined as

$$\frac{dy}{dx} = y/x_1 \oplus y/x_0$$

(4.1)

if the inputs xi to boolean logic are (spatially) independent, then the density of its output y is given by

$$D(y) = \sum_{i=1}^{n} P\frac{dy}{dx_i}Dx_i \tag{4.2}$$

In (2), it is assumed that all inputs are independent. This can lead to inaccuracy where primary inputs will be diverging and than re converging to primary outputs – they are not really spatially independent. However, at a block, the primary inputs can be considered pretty much independent and hence the above approach can be modeled more accurately if the whole block's Boolean difference is computed. Given the signal probability and toggle density values at the

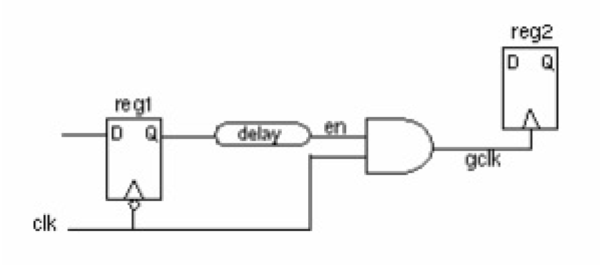

Figure 4.1.1: Schematic of logic 1

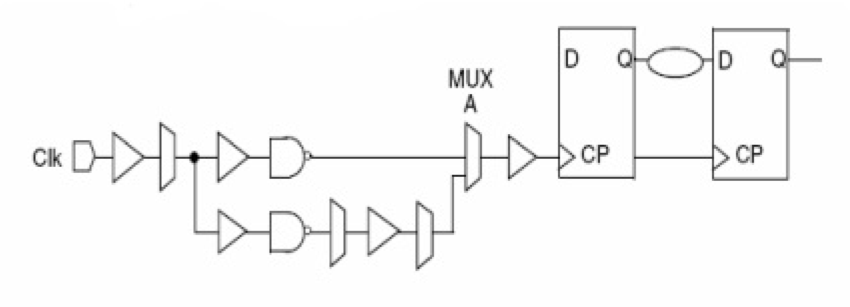

Figure 4.1.2: Schematic of logic 2

primary inputs of a logic circuit, a single pass over the circuit, using (2), gives the density at every node. Note that apart from estimating toggle densities at the output node, we also need to calculate output signal probabilities to do toggle density estimation of subsequent circuit logic. This is simple for two input AND gate.

$$P(Y)=P(A)*P(B)$$
$$P(Y)=1-P(A)P(B)$$

When we apply the above approach, it gives good results for designs which are small and can be analyzed flat and dominated by combinational logic. Beside, it is always not possible to run flat due to other logistic concerns like blocks are designed first or rest of the design is being done hierarchically or there is reusable IPs in design which do not have net list. The approach described in previous section was extended to handle such requirements.

Deriving automatic toggle frequency values

Primary Input Handling The toggle rate at Primary Input is not known. Since they are driven externally, there is no easy way to predict toggle rate for the same. The same is true for primary input signal probability.
Input Delay Specification A constraint that specifies the minimum or maximum amount of delay from a clock edge to the arrival of a signal at a specified input port. Input delay specification is with respect to a clock that triggers events on that signal
Clock specification specifies the characteristics of a clock, including the clock name, source

period and waveform.
Mode Specifications specifies the constant values applied on certain port or pins to drive timing analysis in a specific mode. This means that these pins or ports are not toggling during the analysis. It also specifies the constant value to which the port or pin is tied to.

Sequential element modeling
Sequential elements do not directly switch arbitrarily when the input switches. Hence, we can not apply the formula as mentioned in equation (1,2). We used following formula to compute toggle frequency at the output of sequential cells. Note that we are referring latches and basic flip-flops as part of sequential cells and not the complex macros.
Qout = min(DataInput, clock/2)
The upper bounding of clock/2 is required since we identified certain cases where Data Input toggles more than clock/2. This is explained below. For the cases, where data input is not toggling more than clock/2, output can not toggle more than Data Input

Unconnected inputs going into logic
This was handled by reverse tracking the first sequential cell encountered in the transitive fan out of unconnected inputs. This algorithm gives the clock controlling the toggle rate down the line. If the unconnected inputs are clocks, we assigned the worst toggle rate of the block itself.

4.1.3 Annotating the Switching Activity

Types of Switching Activity to Annotate

- Simple switching activity on design nets, ports, and cell pins. Simple switching activity consists of the static probability and the toggle rate. The static probability is the fraction of the time that the object is at logic 1. The toggle rate is the rate at which the design object switches between logic 0 and logic 1

- State-dependent toggle rates on input pins of leaf cells. The internal power characterization of an input pin of a library cell can be state dependent. The input pins of instances of such cells can be annotated with state dependent toggle rates.

- State-dependent and path-dependent toggle rates on output pins of leaf cells.The internal power characterization of output pins can be state dependent and path dependent. Output pins of cells with state and path-dependent characterization can be annotated with state- and path-dependent toggle rates.

- State-dependent static probability on leaf cells. Cell leakage power can be characterized using state dependent leakage power tables.

Annotating Switching Activity Using RTL SAIF Files

Optimal power analysis and optimization results occur when switching activities reported in the RTL SAIF file are accurately associated with the correct design objects in the gate-level net list. For this to occur, the RTL names must map correctly to their gate-level counterparts. During synthesis, however, mapping inaccuracies can occur that can affect your annotation.

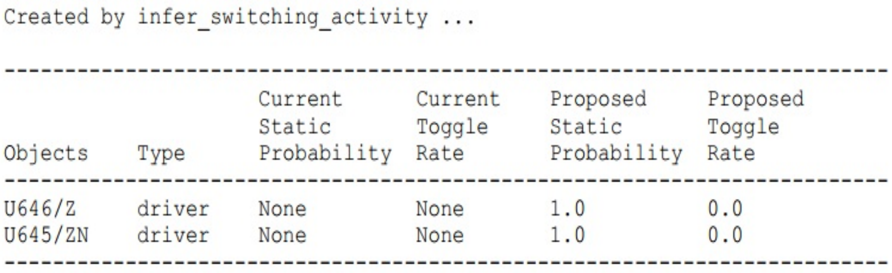

```
Created by infer_switching_activity ...

------------------------------------------------------------------------
                    Current      Current    Proposed     Proposed
                    Static       Toggle     Static       Toggle
Objects    Type     Probability  Rate       Probability  Rate
------------------------------------------------------------------------
U646/Z     driver   None         None       1.0          0.0
U645/ZN    driver   None         None       1.0          0.0
------------------------------------------------------------------------
```

Figure 4.1.3: Infer Switching Activity

Creating a name mapping database that will be further used to perform power analysis.

Before compiling specify name mapping to perform RTL SAIF annotation using the name-mapping database

Gate Level SAIF files

This section reads a SAIF file and annotates switching activity information on the nets, pins, and ports of the design.A SAIF file is usually generated in the HDL simulation flow, where a simulation testbench instantiates the design being simulated and provides simulation vectors. The generated SAIF file contains the switching activity information organized in a hierarchical fashion, where the hierarchy of the SAIF file reflects the hierarchy of the simulation testbench. If a design is instantiated in the testbench (tb) as the instance i, then the SAIF file contains the switching activity information for the design under the hierarchy tb/i. The SAIF file contains time duration values and specifies a time unit which is usually the time unit used during simulation.The synthesis time units are obtained from the time units of the target or link library.

Infer Switching Activity

This task detects the drivers of special pins such as asynchronous set, asynchronous clear, synchronous set, and synchronous clear, and suggests values for toggle rate and static probability. It reports the current and proposed static probability and toggle rate.

4.1.4 Estimating Non-annotated Switching Activity

During power analysis, the PX tool requires switching activity information on all design nets and state-dependent and path-dependent information on all design cells and pins. The tool, by default, estimates switching activity that is not user-annotated before calculating power.

- The design nets where the switching activity cannot be derived through propagation are assigned with default activity.
- The switching activity is implied on the non-annotated pins of buffers and inverters or register outputs, from existing annotation. Based on the functionality of the cell, the tool

derives the switching activity of non-annotated pins without simulation of the annotated pins.

- To determine the activity on all other nets in the design, the tool propagates the annotated activity with zero-delay simulation.

- The simple switching activity on all the nets (either user-specified annotation, default, implied or propagated) is used to derive the state-dependent and path-dependent switching activity for the power arcs of all the cells.

4.1.5 Performing Power Analysis

The tool considers the type and amount of switching activity annotated on your design and chooses the most accurate method to compute your design's power. The method used depends on whether you annotate some or all of the elements in your design.

Factors Affecting the Accuracy of Power Analysis

Annotation
Annotating switching activity relies on the ability to map the names of the synthesis invariant objects in the RTL source to the equivalent object names in the gate-level net list. Mapping inconsistencies can cause the SAIF file to be incorrectly or incompletely annotated, which can affect the power analysis results. In turn, the quality of these results affects the results of power optimization's that rely on the annotation.

Clock Frequency Scaling
If a design is synthesized at a frequency that is different from the frequency the simulation is run, the SAIF file generated from the simulation reflects this difference. This causes a mismatch in timing and affects dynamic power analysis.

Delay Model
Tool uses a zero-delay model for internal simulation and for propagation of switching activity during power analysis. This zero-delay model assumes that the signal propagates instantly through a gate with no elapsed time. The zero-delay model has the advantage of enabling fast and relatively accurate estimation of power dissipation. The zero-delay model does not include the power dissipated due to glitching. If your power analysis must consider glitching, use power analysis after annotating switching activity from full-timing gate-level simulation. As mentioned previously, the internal simulation is used only for nodes that do not have user-annotated switching activity.

Switching Activity Propagation and Accuracy
While propagating switching activity through the design, the logic states of inputs of the gates' can have inter dependencies that affect the accuracy of any statistical model. Such inter dependency of inputs is called correlation. Correlation affects the accuracy of propagation of toggle rates. Because accurate analysis depends on accurate toggle rates, correlation also affects the accuracy of power analysis. It considers correlation within combinational and sequential logic, resulting in more accurate analysis of switching activity for many types of designs. The types of

circuits that exhibit high internal correlation are designs with re convergent fan outs, multipliers, and parity trees. However, Power Compiler has no access to information about correlation external to the design. If correlation exists between the primary inputs of the design, Power Compiler does not recognize the correlation. Power Compiler considers correlation only within certain memory and CPU thresholds, beyond which correlation is ignored. As the design size increases, Power Compiler reaches its memory limit and is not able to fully consider all internal correlation. As an example of correlation, consider a 4-bit arithmetic logic unit (ALU) that performs five instructions. The data bus is 4-bits wide, and the instruction opcode lines are 3-bits wide. The assumption of uncorrelated inputs holds up well for the data bus lines inputs but fails for the opcode inputs if some instructions are used more often. If your design has black boxes, such as complex cells, RAM, ROM, or macro cells you can annotate switching activity at the outputs of these elements.

4.1.6 Analysis

Performing Vector Analysis

During gate-level analysis, use the vector analysis feature to identify the simulation windows with the highest activity, which are likely to produce the peak power. Vector analysis provides information about the vectors by plotting the average toggle rate per interval, over the specified time period. The tool reads the activity file and writes an activity waveform for the top-level module of the activity file. The tool partitions the total simulation time into intervals. For each interval, the average toggle rate is calculated using the following formula:

$$AverageToggleRate = \frac{NumberofTogglesonAllSignalsPerInterval}{NumberofSignals} \tag{4.3}$$

View the activity wave forms to determine if the test bench is simulated as expected and if the vectors have sufficiently covered the design to use them as inputs for power analysis. The coverage per interval is calculated using the following formula:

$$Coverage = \frac{NumberofSignalsWithatLeastOneToggle}{NumberofSignals} \tag{4.4}$$

The output generated from vector analysis has the same names as the instances in the original activity file.

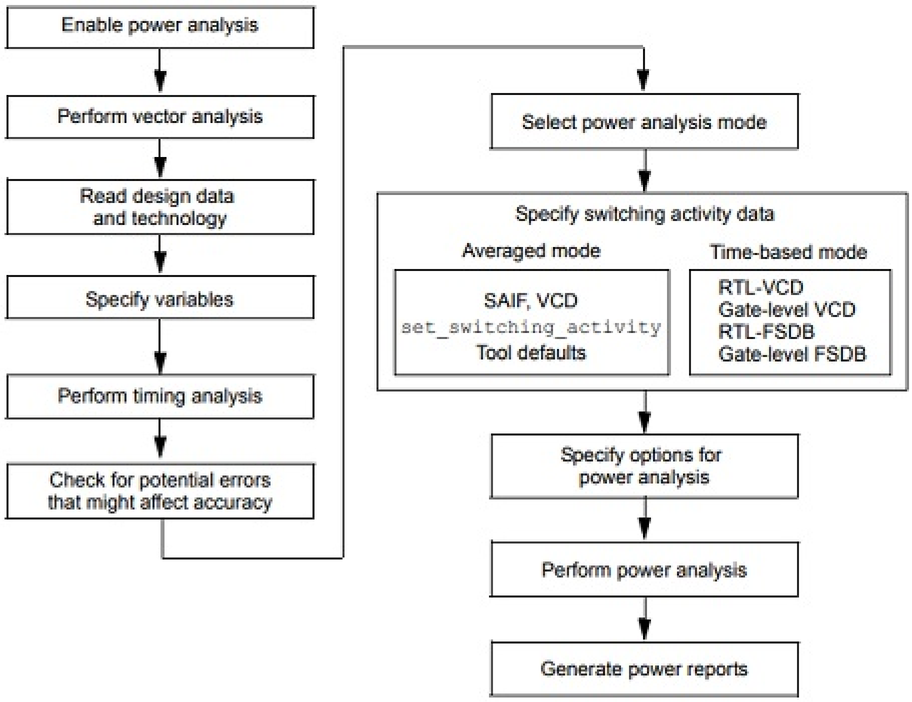

Figure 4.1.4: Power Analysis Flow

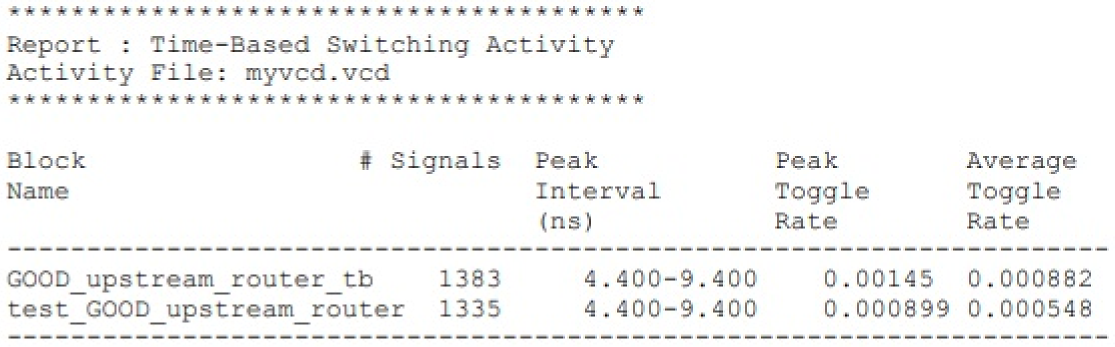

```
****************************************
Report : Time-Based Switching Activity
Activity File: myvcd.vcd
****************************************

Block                   # Signals  Peak          Peak        Average
Name                               Interval      Toggle      Toggle
                                   (ns)          Rate        Rate
---------------------------------------------------------------------
GOOD_upstream_router_tb     1383     4.400-9.400    0.00145   0.000882
test_GOOD_upstream_router  1335     4.400-9.400    0.000899 0.000548
---------------------------------------------------------------------
```

Figure 4.1.5: Time based activity report

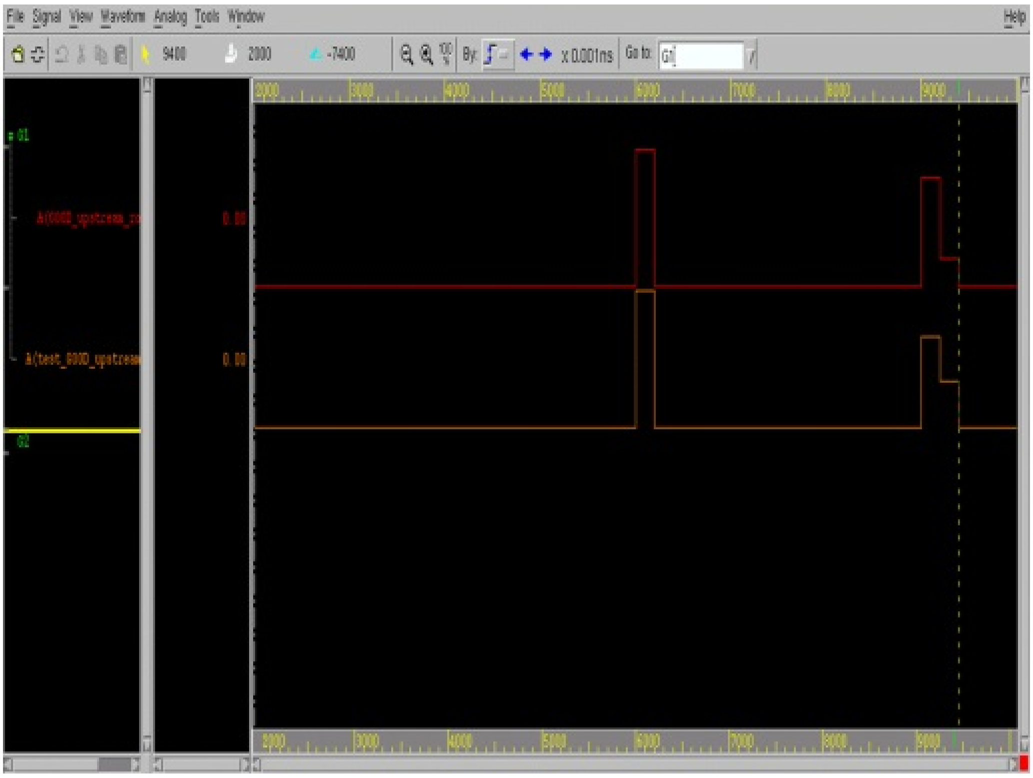

Figure 4.1.6: Activity waveform

Chapter 5

Results

This chapter concludes the results obtained by annotating switching activity information from VCD or SAIF file on the sequential cell nets of the design.There are various dependencies observed on the overall QoR of the design ,which are represented in the form of table and graph.The experiments are done on blocks having different sizes in terms of number of macros and standard cells ,including designs which are having multiply instantiated module.

5.0.1 Test Design I

Effect of setting different weights

This design has 20k standard cells ,2 macros.To estimate the power for this design accurate activity is annotated on the design and a comparison is carried out between estimated length and switching power. On Similar basis another experiment is carried out to check net weights assigned on the nets ,which performs low power placement.

net_weight	internal(in mW)	switching (in mW)	leakage (in mW)	Total (in mW)
1	103.1%	96.83%	99.01%	99.02%
2	102.1%	96.04%	99.01%	99.02%
3	102.1%	100.39%	98.8%	100%
default	100%	100%	100%	100%

Table 5.1: Net weight vs Switching power

internal(in mW)	switching (in mW)	leakage (in mW)	Total (in mW)	wire length (in um)
103.1%	96.83%	99.01%	99.02	94.464%
103.1%	96.83%	99.01%	99.02	77.054%
102.1%	100.39%	98.8%	100	79.075%
100%	100%	100%	100	100%

Table 5.2: Switching power vs net length estimation

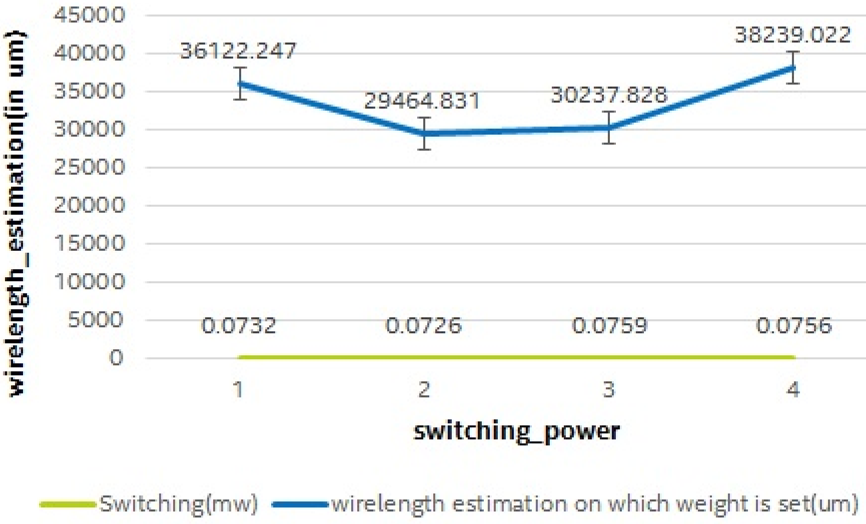

Figure 5.0.1: Wire length variation vs switching power

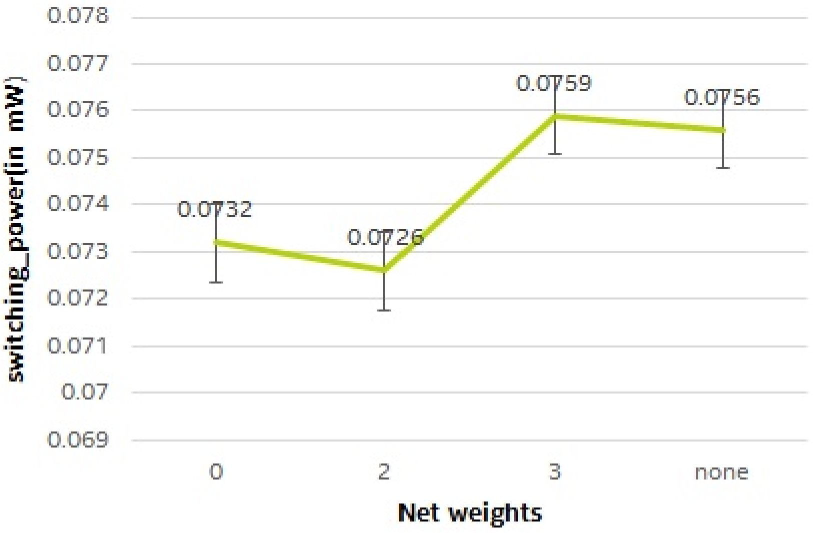

Figure 5.0.2: switching power variation with varying net weights

5.0.2 Test Design II

This design consists of around 150k standard cells and has 7 macros in it.Since the no of sequential cells in the design is more ,the chances of saving power here due to switching activity is estimated to be more.Also this exercise is performed with variation of weights in and examining the change in switching power and wire length estimation.

net_length	internal(in mW)	switching (in mW)	leakage (in mW)	Total (in mW)
1	100%	99.30%	99.56	96.49%
3	99.08%	99.44%	99.56	98.68%
4	99.08%	99.58%	99.56	98.85%
5	100%	99.72%	99.56	99.12%
default	100%	100%	100%	100%

Table 5.3: Switching power vs net weights

internal(in mW)	switching (in mW)	leakage (in mW)	Total (in mW)	wire length (in um)
100%	99.30%	99.56	96.49%	101.31%
99.08%	99.44%	99.56	98.68%	81.56%
99.08%	99.58%	99.56	98.85%	73.82%
100%	99.72%	99.56	99.12%	80.9%
100%	100%	100%	100%	100%

Table 5.4: Switching power vs net length estimation

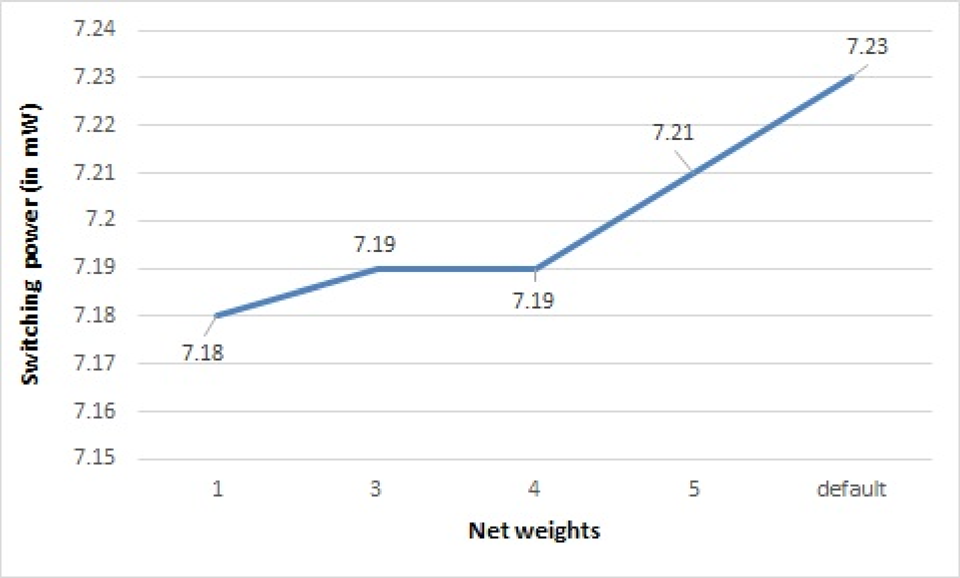

Figure 5.0.3: net weight vs switching power

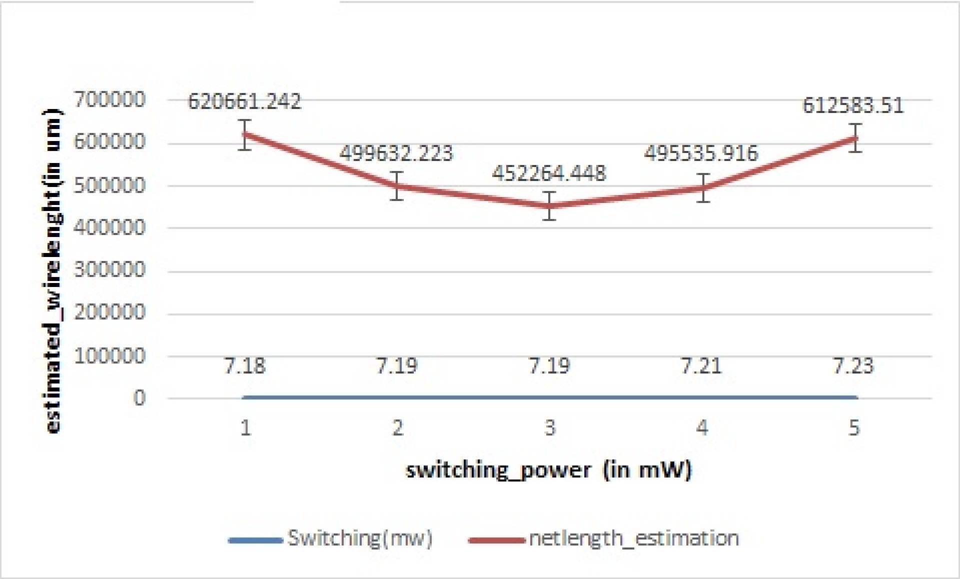

Figure 5.0.4: estimated wire length vs switching power

Bibliography

[1] Solvnet Synopsys.

[2] Process Variability at the 65nm node and Beyond – Sani.R.Nassif IEEE 2008 CICC.

[3] The International technology Roadmap for Semiconductors, 2007.l.

[4] Static Timing Analysis for Nanometer Designs - A Practical Approach by J. Bhasker, Rakesh Chadha.

[5] Rakesh Chadha, J. Bhasker auth. An ASIC Low Power Primer Analysis, Techniques and Specification.

[6] S. Bobba and I. N. Hajj, “Estimation of maximum current envelope for power bus analysis and design,” Proc. of ISPD, pp 141-146, Apr 1998

[7] Himanshu Bhatnagar, “Advanced ASIC Chip Synthesis: Using Synopsys Design Compiler Physical Compiler and Primetime”, 2nd Edition, Kluwer Academic Publishers, ISBN: 0792376447

[8] Synopsys, Inc, “Using the Synopsys® Design Constraints Format”, Application Note, Sept 2005.

[9] Najm, F.N, “Transition Density, a stochastic measure of Activity in Digital Circuits”, DAC, pp. 644-649, June 1991

[10] F. Najm, “A survey of power estimation techniques in VLSI circuits, ”IEEE Trans. VLSI System., vol. 2, pp. 446–455, Dec. 1994.

[11] Rabe, D; Jochens, G.; Kruse, L.; Nebel, W, „“Power-simulation of cell based ASICs: accuracy- and performance trade-offs”, Proceedings of Design automation and test in Europe, Feb 1998

[12] J.-Y. Lin et al., “A cell-based power estimation in CMOS combinational circuits,” in Proc. IEEE Int. Conf. Computer-Aided Design, 1994, pp. 304–309.

[13] Dresig, F. Lanches, P. Rettig, O., et al, “Simulation and reduction of CMOS power dissipation at logic level”, Design Automation, 1993, with the European Event in ASIC Design. Proceedings, pp. 341-246, Feb 1993

[14] An-Chang Deng Yan-Chyuan Shiau Loh, K.-H, “Time domain current waveform simulation of CMOS circuits”, IEEE international conference on Computer aided design 1988, pp. 208-211, Nov 1988.

[15] F.N. Najm, R.Burch, P. Yang, and I.N. Hajj. “Probabilistic Simulation for Reliability Analysis of CMOS VLSI Circuits”. IEEE Transactions on CAD, 9(4):439-450, April 1990.

[16] Min Zhao; Panda, R.V.; Sapatnekar, S.S.; Edwards, T.; Chaudhry, R.; Blaauw, D, “Hierarchical analysis of power distribution networks”, DAC, pp. 150-155, June 2000

Given a set of numbers, there are elementary methods to compute its Greatest Common Divisor, which is abbreviated GCD. This process is similar to that used for the Least Common Multiple (LCM).