

UNIVERSITY OF THESSALY

DIPLOMA THESIS

Localized IR Drop Hotspot Identification

Author:
Vasileios SAMARAS

Supervisors:
Christos P. SOTIRIOU
Nestoras EVMORFOPOULOS
Georgios STAMOULIS

*A thesis submitted in fulfillment of the requirements
for the degree of Diploma*

in the

Circuits and Systems Laboratory (CASlab)
Department of Electrical and Computer Engineering

July 6, 2021

DISCLAIMER ON ACADEMIC ETHICS AND INTELLECTUAL PROPERTY RIGHTS

«Being fully aware of the implications of copyright laws, I expressly state that this diploma thesis, as well as the electronic files and source codes developed or modified in the course of this thesis, are solely the product of my personal work and do not infringe any rights of intellectual property, personality and personal data of third parties, do not contain work / contributions of third parties for which the permission of the authors / beneficiaries is required and are not a product of partial or complete plagiarism, while the sources used are limited to the bibliographic references only and meet the rules of scientific citing. The points where I have used ideas, text, files and / or sources of other authors are clearly mentioned in the text with the appropriate citation and the relevant complete reference is included in the bibliographic references section. I fully, individually and personally undertake all legal and administrative consequences that may arise in the event that it is proven, in the course of time, that this thesis or part of it does not belong to me because it is a product of plagiarism».

The declarant

Vasileios SAMARAS
July 6, 2021

DISCLAIMER

«All work presented in this thesis has been conducted in collaboration with Qualcomm Inc. Because of this, and by respecting the confidentiality agreement and the contract that the author has signed, the supervisors and the company (Qualcomm Inc.) have agreed to omit any information that would put these at risk. ».

The declarant

Vasileios SAMARAS
July 6, 2021

“Contrary to popular belief, I know exactly what I’m doing.”

Tony Stark (Robert Downey Jr.)

UNIVERSITY OF THESSALY

Abstract

Department of Electrical and Computer Engineering

Diploma Thesis

Localized IR Drop Hotspot Identification

by Vasileios SAMARAS

The end of Dennard scaling, which argued that one could continue to decrease the transistor feature size and voltage while keeping the power density constant, along with the end of Moore's law has raised a big challenge for large transistor count IC designs. As we push through lower technology nodes in the IC and chip design, the wire width goes thinner along with transistor size. Designers have been focusing on reducing power consumption and increasing the number of gates on the device. At the same time downscaling also benefits the performance of digital systems as smaller transistors can operate on higher frequencies improving system speed. This makes the wire resistance more dominant on 16nm and below technology nodes. This increasing resistance and the decreasing width of metal wires introduce many IR drop issues. These issues play major roles in reducing the lifespan of an electronic device and are the causes of functionality failure in any electronic devices with lower technology nodes. IR Drop is an electrical phenomenon that produces fluctuations in the supply voltage caused by the combination of parasitic resistive elements in the power distribution network and the current needed to operate the device. Moreover, there can be other factors that can cause IR drop, and identifying the real problem and devising a good solution is not an easy task. Every EDA company has its own tool which performs the IR analysis and based on the analysis, techniques for IR fixes are applied. In this thesis, we explore and exploit the locality of the IR Drop phenomenon in an attempt to capture the maximum number of IR Drop hotspots early in the design cycle through the generation of scenarios that provide us with realistic simulations of how our circuit would operate in the real world. Simultaneously, we address the need to present accurate simulation scenarios for every possible design, since the availability of real word scenarios from the designers is limited to a few critical parts of the integrated circuit that should definitely be examined due to the criticality and importance of it and argue that our methodology makes fast and predictable switching scenarios available for early IR hotspot detection with insignificant impact on runtime and tremendously accurate results that are close to a real-world measurement.

Acknowledgements

First and foremost, I would like to express my gratitude to my longtime collaborator and supervisor Prof. Christos P. Sotiriou for his continuous support and guidance throughout the years of my undergraduate studies. He provided me with the necessary stimuli to develop my skills and grow as an engineer, a researcher and a person. I am honored and consider myself extremely privileged to call him my mentor.

Moreover, I would like to thank Prof. Nestoras Evmorfopoulos and Prof. Georgios Stamoulis for their key suggestions, valuable comments and feedback. The assistance they provided during my undergraduate studies have been immensely important and I will forever be grateful to them.

Additionally, I would like to render my warmest thanks to my supervisor during my internship in Qualcomm's R&D center in Cork, Ireland, Dr. Cristiano Lopes Do Santos who trusted me, supported me and taught me all about Power Delivery Networks and IR Drop Analysis. He kept pushing me to improve and accomplish my goals, helped me develop my technical skills, my soft skills and showed me the way to conduct research in an industrial environment.

Furthermore, I have to acknowledge and thank my friends, old and new, for all the moments we have lived together all around the globe. They made these years unforgettable and joyful, assisted me to get out of tough situations and kept me going in this Marathon.

Finally, I would never manage to be here in the absence of the constant support, understanding and unconditional love of my family.

Contents

Abstract	vii
Acknowledgements	ix
1 Background	1
1.1 Power Distribution Network	2
1.2 IR-Drop	5
1.3 Static & Dynamic IR Drop	8
1.3.1 Static IR Drop Analysis	8
1.3.2 Dynamic IR Drop Analysis	9
2 Current IR Drop Analysis Methodologies	11
2.1 Static Analysis	12
2.2 Dynamic Vectorless Analysis	12
2.3 Dynamic Vector-based with critical cycle selection	13
3 Enhanced Vectorless Signoff	15
3.1 Enabling Multi-Cycle Analysis	16
3.2 Taking in account the locality of the phenomenon	16
3.3 Triggering a different portion of the design	18
3.4 Handling Clock Instances and Memories	19
4 Experimental Methodology & Results	21
4.1 Experimental Methodology	22
4.2 Metrics	22
4.3 Results	24
5 Conclusions	31
5.1 Conclusions	31
5.2 Future Work	31
Bibliography	33

List of Figures

1.1	PDN topology example	2
1.2	Simple model of a power grid	5
1.3	Current waveform of different transistors	9
2.1	Power waveform of an entire VCD	14
4.1	The ASIC Flow	22
4.2	Current when a device is turned on	23
4.3	Switching Coverage and Runtime	24
4.4	Switching Coverage breakdown	25
4.5	IR Outliers and Peak Current	26
4.6	Peak Cycle and Average Power	27
4.7	Omitted due to confidentiality	28

List of Abbreviations

IC	Integrated Circuit
EDA	Electronic Design Automation
VCD	Value Change Dump
ECO	Engineering Change Order
ECO	Digital Signal Processing
PG	Power Grid

Dedicated to my family

Chapter 1

Background

In this chapter I will be focusing on presenting some necessary terms and ideas so anybody can read this thesis.

1.1 Power Distribution Network

The PDN has been presented as one of the most important elements to model accurately because currents flowing through the PDN and the resulting supply voltage fluctuations determine the gate delays and the potential delay faults. As it is well known, PDN is a complex system of wires delivering power to the whole integrated circuit through different layers. A PDN is classically organized as a set of parallel large wires located in the upper metal layers covering the whole circuit surface [Rius, 2013a].

Obviously, the electrical PDN model depends on the physical structure of the PDN. For the sake of simplicity, we assume the following classical structure:




FIGURE 1.1: PDN topology example

- In the top metal levels of the chip, high metal level # n and metal level # $n-1$ are exclusively composed of a set of parallel metal lines, these two sets having orthogonal directions. In the high metal levels through-vias that connect the two sets of orthogonal lines are regularly placed. The whole set of metal lines and through-vias creates a regular two-dimensional distribution network as represented in the Figure. Usually, in a given level, one line over two is dedicated to Vdd and every other line to Gnd [I. Polian and Becker, 2006] in this way, it is possible to analyze the Vdd distribution network and the Gnd distribution network as two independent three dimensional distribution networks. The figure illustrates the two orthogonal networks for Vdd PDN and Gnd PDN.
- In the bottom metal levels of the chip, metal #2 is commonly used for the Vdd and Gnd lines. The Vdd and Gnd lines in the metal #2 level have typically a small length corresponding to the mega-cell they feed [I. Polian and Becker, 2006].

In addition they have multiple parallel via connections to the upper regular three-dimensional network as illustrated in the Figure.

- Finally, intermediate metal level represents the interface between the upper regular three dimensional network and the lower irregular structure as illustrated in the Figure.

PDN topology is obviously optimized in the design phase to reduce the parasitic elements but they cannot be completely removed. And so, vertical crossings of conductive elements behave as capacitive elements and long wires and vias behave as resistive elements.

In addition, inductive elements may be considered although these are very small and often neglected. Among all these parasitic elements, the resistive ones are really predominant; reason why IR-Drop is one of the most analyzed electrical phenomena due to the current flows through the PDN.

Therefore, the distribution network model must contain an accurate representation and estimation of the PDN resistive elements.

From a modeling perspective, PDN topology can be divided into three areas:

- High metal level may be modeled as two independent two-dimensional resistive grids. The Vdd and Gnd lines in this level of the chip are very long, corresponding to the whole chip size. For this reason the parasitic resistances of the regular network are determinant in the current distribution through the PDN.
- Intermediate metal level may be considered as included in the two-dimensional grids
- Low metal level is made of short metal connection (in the order of the mega-cell dimension) in comparison with the high metal level wires (in the order of the chip dimension). Moreover, there are multiple parallel via connections to the metal #3 lines that enable to reduce the resistive behavior of metal #2 lines. For these reasons, the parasitic resistance of this level can be neglected in the model.

In conclusion, the electrical model of the PDN concerns the two independent supply networks, which are physically regular, intertwined and orthogonal from one level to the other, and correspond to the high metal levels. We can accurately model the PDN as a symmetrical structure composed of two grids with resistive elements.

However, this was a simple example.

To properly study a PDN including capacitive elements is mandatory. The first step is to analyze all the types of capacitive elements that could have an impact on the current distribution behavior. Capacitive elements can be parasitic capacitors (due to the physical superposition of electrical elements) or capacitors that are intentionally included during the design process:

- Parasitic capacitors of the physical PDN. Parallel conductive wires of the PDN behave as capacitive elements. The conventional way to model this parasitic capacitive element is to include small capacitors regularly in the PDN resistive model. The size of these capacitors is determined by the layout parameters of the PDN. These capacitors are connected to every node of the resistive grid.
- Intentional decoupling capacitors. In the design phase, on-chip decoupling capacitors are deliberately included in the PDN. Decoupling capacitors are an efficient way to reduce the power supply noise created by the transient elements in the IC. Moreover, placing some decoupling capacitors in the design reduces power supply fluctuations between different areas of the IC. Consequently, decoupling capacitors reduce the IR-Drop impact. Intentional decoupling capacitors are much larger than the parasitic capacitive elements of the PDN and their placement is usually determined during the PDN design phase using commercial tools. Therefore, value and placement of the decoupling capacitors are known.
- Intrinsic decoupling capacitors due to non-switching gates. The CMOS transistors of the logic gates have intrinsic decoupling capacitance elements due to their internal electrical parameters and the interconnection capacitance[M. Popovich, 2008]. The equivalent decoupling capacitor of a gate is a function of the internal transistor capacitance and the interconnection capacitance. It can be calculated using the layout information. In a single gate, the intrinsic decoupling capacitor is negligible, but it becomes very significant when a large number of non-switching gates are connected to the same node of the PDN.

1.2 IR-Drop

IR-Drop is defined as an electrical phenomenon associated with the switching of MOS transistors. A current draw appears in the power supply connection and/or the ground supply connection when transistors switch. The inherent parasitic elements of the PDN combined with this current draw produce fluctuations in the voltage level. IC scaling technology has increased the density of transistors and the functional frequency. Thus, there are more gates switching simultaneously and consequently, voltage fluctuations have also increased due to the increase of the amount of current flowing through the PDN.

As a consequence of these fluctuations, logic gates can be powered with a lower-than-normal Vdd or higher than-normal Gnd or both, reducing the gate swing and impacting logic gates by an increased delay.

Moreover the sensitivity of the gate delay to power supply noise increases with technology scaling. It has been reported that fluctuations of 10% in power/ground supply voltage increase gate delay by 8% in 180nm technology [R. Saleh and Overhauser, 2000], but fluctuations of 10% can cause up to a 30% increase in gate delay in a 130nm technology [S. Pant and Panda, 2003], and a 1% change in power supply voltage causes nearly 4% of additional gate delay in 90nm technology [C. Tirumurti and Change, 2004].

The impact of the IR Drop phenomenon has therefore become a critical concern. There are a variety of reasons that may cause IR Drop:

- Improper placement of Power/Ground pads.
- Wrong block placement.
- Bad global power routing.
- Insufficient Core Ring / Power Strap width.
- Lesser number of Power Straps.
- Missing Vias.
- Insufficient number of Power Pads.

Let us have a simple model of a power grid as shown below.




FIGURE 1.2: Model of a power grid

The resistors are the effective resistances in the VDD and VSS power grid distribution and G1-G4 are the gates switching with their own capacity. Ideally, the voltage

coming from the power supply to the gate G4 is the same. Actually, the voltage in that node is reduced due to current flowing through the resistances R11-R14 that cause a voltage drop. Vice versa, the voltage in the ground line will rise because the current for G4 to VSS will pass through R21-R24. This can be applied for every cell in the design.

This variation will be equal to:

$$\Delta V_{max} = I_{avg}(R_{Vdd} + R_{Vss}) \quad (1.1)$$

This is why it is called IR drop, since $V = I * R$. In this case, I_{avg} is the current consumed by the cell switching and R is the sum of the effective resistance of the power grid for the pair VDD and VSS[Cadence Design Systems, 2002].

Power Supply Noise (PSN) refers to the voltage fluctuations in the power and ground distribution networks (PDN). The voltage fluctuations due to power supply noise in the PDN are generally called voltage drop.

The power distribution network includes all the metal wires and vias that deliver power to every gate in the chip. This on-chip PDN is predominantly resistive but capacitive and inductive parasitic elements are also presented.

Power supply noise is induced by current flows through the PDN:

- IR-Drop is generated due to the resistive elements of the PDN.
- Ground bounce is generated due to the inductive elements of the PDN ($L*(dI/dt)$).

The design of a good, reliable on-chip PDN of a digital IC is a very complex task because designers cannot anticipate all the functional conditions Design researches about the power supply sensibility try obviously to reduce as much as possible the whole power supply noise effect at the chip level.

In essence, the principle is to estimate the supply voltage drop due to the IR-Drop and try to adapt the PDN design to minimize this phenomenon. To this aim, some techniques are applied during design to decrease the power supply noise and to improve the noise immunity of the circuits as explained by Larsson in [Larsson, 1999]. The most widely used technique consists in adding decoupling capacitors between the power and ground supplies. Decoupling capacitors prevent the power noise from spreading through the PDN and their inclusion in the design allows to isolate different areas on the chip.

In this context most of the works try to develop algorithms to determine the optimal size and placement of the decoupling capacitors from the switching activities and the spatial correlations between different blocks. A power supply noise aware post-floorplanning methodology is proposed in [S. Zhao and Cheng, 2002] [H. M. Chen and Wong, 2003] [S. A. Moghaddam and Lucas, 2005]. Another research suggests to improve the traditional decoupling capacitor and to include active decoupling capacitors as a most effective technique to reduce the power supply noise [J. Gu and Kim, 2006]. Another classical way to reduce the power supply noise is to design a robust power distribution network. A lot of research works mark that traditional constraints in the PDN design are not enough to remove the IR-Drop timing faults.

In order to improve the PDN design, most of the works propose a mathematical model to address the most important issues in the PDN design: width and pitch of PDN wires [Guptaa and Kahng, 2006] [A. Mukherjee and Marek-Sadowska,

2002], size, number and location of pads [D.A. Andersson, 2008], [Bhooshan, 2007], [Bhooshan and Rao, 2007], [K. Shakeri and Meindl, 2003], [Shakeri and Meindl, 2005]. Wire sizing for power and ground networks considering the IR-Drop induced by both the clocking and computing components is suggested in [A. Mukherjee and Marek-Sadowska, 2002] and considering the IR-Drop and the area constrained in [Guptaa and Kahng, 2006]. To determine the size, number and location of pads, [Bhooshan, 2007] and [Bhooshan and Rao, 2007] propose a closed form model for the power distribution network in N-metal layer system for wire-bond and flip-chip packages in function of given design constraints (as power dissipation, power supply voltage or static IR-Drop).

Other works focus on the correlation between different parameters of the on-chip power distribution grid and their impact on noise [D.A. Andersson, 2008], [K. Shakeri and Meindl, 2003], [Shakeri and Meindl, 2005]. Results from these papers can be used as guidelines when designing a robust power distribution network.

Models proposed by [Rius, 2013b] and by [K. Shakeri and Meindl, 2003], [Shakeri and Meindl, 2005] exclusively focus on the IR-Drop phenomenon. In [Shakeri and Meindl, 2005] the author demonstrates that the PDN can be approximated as a continuous layer of conductive material and that IR-Drop can be calculated by solving a system of partial differential equations, i.e. Poisson equation, with the proper boundary conditions. The author proposes a compact physical IR-Drop model of the on-chip power distribution grid and an IR-Drop model is derived for the wire-bond and the flip-chip packages. In this paper, the tradeoff between the package and the on-chip power distribution network parameters is studied in details. The size and number of pad tradeoff is also analyzed. The optimal placement of these pads is derived to minimize the IR-Drop. In brief, Shakeri suggests the use of a large number of small pads for the power distribution network instead of a small number of large pads to reduce the IR-Drop.

Based on the conclusion of Shakeri, [Rius, 2013b] suggests another IR-Drop model to determine the average power consumption of a block. Initially, the IR-Drop is modeled in an infinite PDN. Then, the IR Drop model in a finite PDN is derived from the infinite model solution.

The suggested model provides an accuracy estimation of the average power consumption of a block for the wire-bond package. Models of Shakeri and Rius help designers in the early stage of the design to estimate accurately the on-chip and package resources that need to be dedicated to power distribution, reducing the cost of over-design.

In brief, most of these works are based on a vector-less approach and primarily target the spatial effect of supply voltage noise. Suggested statistical circuit models estimate the average current consumption at the chip level, allow to identify the critical areas and to adapt the PDN network design in order to avoid the undesirable voltage drop. Although the design approaches do not take into account the input vector dependence of the IR-Drop phenomenon, the statistical models provide an estimated supply voltage drop at the chip level.

It is a fact that a vector-based simulation at the chip level is non-viable due to prohibitive simulation costs.

1.3 Static & Dynamic IR Drop

Static IR-Drop is an average of the voltage drop of the circuit. Dynamic IR-Drop depends on the switching activity of the logic, hence is vector dependent. Dynamic IR-Drop depends more on the switching time and less on the clock period. Which means that IR-Drop should not become correlated with the average current which depends totally on the time period, where as dynamic IR-Drop depends on the instantaneous current which is higher while cells are switching.

In older technologies where natural decoupling capacitances from the PDN and non-switching logic were available, static IR-Drop constituted adequate sign-off analysis. Nowadays modern technologies include huge amounts of circuitry which is switching simultaneously. That switching results high current demands making dynamic IR-Drop a critical factor. These current demands can be brief and highly localized or the exact opposite, meaning long lasting and extending in a large area causing setup or hold time violations.

1.3.1 Static IR Drop Analysis

Static IR drop is the simplest analysis possible. It is used in the early stages of design development in order to give some initial information about the robustness of the power grid [S. K. Nithin, 2010], [Y. Ban, 2014].

In general, it consists of a simulation requiring little information, but it cannot guarantee a very realistic value of IR drop; however, a short time will be needed to simulate the circuit and it could provide an early estimation so as to highlight the most significant power grid problems. The verification with static voltage drop is not enough to ensure chip integrity because it considers only the parasitic resistance of the power grid and the average current, which is used to calculate the average voltage drop [K. Arabi, 2007].

The static analysis computes an estimation of the peak current in order to determine the maximum voltage drop on each node. A typical flow for a complete IR drop analysis involves simulating the chip with static analysis first, then verifying that the IR drop is below a user specified threshold and then, using the results as a baseline for the dynamic analysis.

In the static analysis, a typical approaches follows the following steps:

- Extract the parasitic resistance of the power grid.
- Create the power grid as a resistor matrix.
- Calculate the average current for each gate connected to the power grid.
- Distribute this average current around the resistance matrix, according to the physical location of considered gate.
- Apply for each power pad a source of VDD.
- Calculate the voltage on each cell using the current and the resistance matrix.

1.3.2 Dynamic IR Drop Analysis

Dynamic voltage drop analysis is the next step for the IR drop analysis. Dynamic analysis is used to detect integrity problems, evaluating the IR drop caused by peak current demand, due to simultaneous switching of gates, and considering the impact of decoupling capacitors and inductors [Y. Ban, 2014].

To understand the very basic difference between static or dynamic analysis we can see a simple example from [Zhu, 2004].

Let us consider two timing diagrams.




FIGURE 1.3: Current waveform of different transistors

These diagrams show the current pulse for different transistors. Each current pulse has the same magnitude but a different timing. In fact, in the first one all pulses occur at once, in the second one they are spread over a period of time. In both cases, the current has the same average value for all transistors.

For this reason, the static analysis will not see any difference. On the contrary, the dynamic one will see a worst IR drop in the first case, caused by the simultaneous pulse of the transistors.

Dynamic analysis gives a more realistic value of IR drop and helps identifying which input activates a certain problem in the design.

Depending on the input data, we can distinguish two different dynamic analyses: Vector-based and Vectorless.

The vector-driven approach requires a preliminary functional simulation of the device. The back-annotated gate-level or post-layout simulation (i.e., considering cell and propagation delays) precisely identifies the occurrence of signal switching in time.

With this data, IR-drop analysis can accurately evaluate voltage values on the power grid in time and space. From an operative point of view, signal and cell activity data is produced by the simulator in a standard file format known as Value change dump (VCD). Knowing which cells are switching and when, the tool can create a current waveform based on the internal power consumption of each cell.

Generally, the VCD file generation can be a very time-consuming process, and the result is a huge database. For this reason, the application of the analysis process using the VCD associated with a long simulation is impracticable, due to the limited resources; one method is selecting a time interval from the available dataset, based on the highest activity and the worst-case power [Cadence Design Systems, 2018].

On the other hand, the vectorless approach can be very useful for an early-stage and full-chip verification [W. Zhao, 2014]. In the vectorless approach, through a probabilistic calculation, the analysis tool generates a worst-case input vector, without using actual simulation. To do that, it needs timing and activity information in order to know when and how often an instance is switching. The timing information is extracted from the data generated from static timing analysis, done during all the phases of the implementation in order to check the timing correctness.

The switching activity is set in a statistical way.

Once the power analysis is done, the tool will generate time-varying current waveforms for all instances based on the internal power defined in the cell libraries. The calculated waveforms are then fed into the dynamic analysis tool for the calculation of the IR drop. In general, the vectorless approach is less accurate than the VCD based one, because the switching activity comes from a statistical distribution. However, constructing a worst-case dynamic voltage profile in an early stage can be very useful and offers a great improvement in terms of run time over the vector-based approach, which is not feasible for a full chip verification [X. Xiong, 2010].

Chapter 2

Current IR Drop Analysis Methodologies

In this chapter I will be focusing on presenting the current methodologies that are being used in the Industry and in the Academia for IR Drop Analysis.

2.1 Static Analysis

The static analysis is the first step for a complete IR drop analysis.

The static IR drop analysis is widely used early in the flow to detect and correct problems in the power grid, such as missing vias, shorts, insufficient power routing width, etc. It is also used to simulate the average impact of the design operating for a long period; in fact, it calculates the average voltage drop associated to a user-defined circuit activity.

Even if static analysis is not optimal in precision, depending on how many input data we give to the simulation, we can have different levels of accuracy. In the analysis without functional simulation data (VCD file), the exact value of the switching current of cells is not known, because we do not know when they will switch and their exact activity. In this case, the activity is calculated with probabilistic computations, so the power dissipated will be an approximate value.

Moreover, the gate characterizations give better accuracy, since the tool will know the exact current behavior of each cell, and so on. It must be noted, however, that the objective of static analysis, generally, is not the highest accuracy, because the detection of power grid problems can be accomplished even without the computation of exact voltage drop values.

Since the leakage current will contribute to the value of power dissipation, it is easy to understand how the temperature can influence the voltage drop. The leakage current increases exponentially with temperature, thus changing the voltage drop value as well.

Therefore, the static analysis is performed using low, nominal and high temperature.

In order to do that we need to change the corners for each simulation. These corners are a triad of values, representing the performance of process (worst, nominal and best), voltage, and temperature.

They are usually identified with the PVT acronym (from process, voltage and temperature).

By changing corner, it is possible to simulate the circuit in different conditions.

Usually the worst corner for the IR drop analysis is best process, high voltage and high temperature [Y. Ban, 2014].

2.2 Dynamic Vectorless Analysis

The Vectorless analysis is a good surrogate for a dynamic analysis if simulation data are not available, which is extremely common.

In this type of analysis, the tool does not compute the average voltage drop over a clock cycle; instead, it performs the voltage drop calculation considering the sum of peak currents. These currents are calculated every time instant, called steps, and for each step, the voltage drop is computed. Since we are not working with average voltage drop, we will expect higher values with respect to the static analysis. Like the static analysis, different levels of accuracy can be obtained depending on the data

input given to the tool.

As an example, in the vector-free analysis current gate characterization has an important weight for the same reason seen before, that is because we are performing instantaneous peak current calculations. In fact, the triangular or trapezoidal model can lead to inaccurate results, while more detailed models can better approximate the exact current values. In this analysis it is usually possible to obtain current waveforms for the entire chip, analyzing the time interval with highest simultaneous activity.

We will see that this graph can be compared to the graph generated from the analysis with a VCD.

2.3 Dynamic Vector-based with critical cycle selection

The dynamic vector-based analysis is the optimal analysis in terms of accuracy, but it requires more resources and time.

The simulation vectors (VCD) are used to provide the most accurate switching information in order to construct a realistic switching scenario. In fact, every net or instance activity is determined entirely from the VCD.

Similarly as for the vectorless analysis, the tool analyzes peak currents, but now it knows the exact activity of each net or instance, so the analysis can be more accurate and realistic.

However, the simulation becomes much longer than in the previous case. In fact, since the large number of instances that need to be simulated, generally it is unfeasible to use the complete set of data deriving from a functional simulation of the circuit.

So, the analysis is applied in the worst case only.

One possible method involves the selection of the time interval during which it is present the higher activity, and perform the analysis in that interval.

We can see Figure 2.1 as an example: this represents the power waveform acquired performing the power calculation of an entire VCD. This is called FSDB (Fast Signal Database) and represent the power demand of the entire design with respect to the entire simulation time of the VCD. From that waveform, we clearly see that there are two time intervals (inside the red rectangles) with high power demand. Therefore, one way to go is to apply the analysis in those intervals only.

If there is the possibility to produce the FSDB, one can select the intervals manually. It must be noted that this waveform may be difficult or time consuming to obtain, or maybe the waveform coming from the simulation can be quite regular, so selecting the right interval can be very challenging.

Another option is the critical cycle selection method, which is a feature offered by some analysis tools that select automatically the worst power interval. During power calculation, the tool calculates the power demand for each time span of a fixed length. Then it will apply the voltage drop analysis in the selected interval, i.e., the one with most average power consumption.

Moreover, it will report a ranking in terms of average power demand and voltage change per cycle for future debug. The latter represents how much the voltage changes from one cycle to another and this can be a further cause of IR drop.




FIGURE 2.1: Power waveform of an entire VCD

The manual method is expected to be more accurate than the automatic way. The reason is that the power calculation done in the cycle selection is built to be fast, so it is less precise than the one used for the production of the FSDB, produced by a simulation of the VCD.

Finally, it is important to clarify that maximum power does not mean that the analysis will identify the maximum number of hotspots, since the activity may be located in a really small area that is very power hungry.

Chapter 3

Enhanced Vectorless Signoff

In this chapter I will be focusing on the methodology that was developed in order to address the challenges that arise when engineers try to efficiently stress a PG so they can identify structural PG weaknesses and localized IR drop hotspots that may cause malfunctions. As explained earlier, IR drop may cause timing violations, cause the chip to overheat or even destroy the IC.

Due to limitations that the confidentiality of this work introduces, explaining the methodology in-depth and sharing every step is not possible since this methodology is considered a competitive advantage and company confidential information.

3.1 Enabling Multi-Cycle Analysis

When performing Vectorless Dynamic IR Drop Analysis, the user has to decide if the analysis will be :

- Single-cycle.
- Multi-cycle.

Single-cycle applies the same switching scenario in every clock cycle, for the whole duration of the simulation. This is not realistic since the chance of a circuit having the same instances switching continuously for, let's say 20 clock cycles, is pretty low. In addition to that, by continuously triggering the same instances for the whole duration of the simulation only concentrates the switching activity in a small area, achieving low spatial coverage and evaluating only a small fraction of the PG.

Multi-cycle applies a different switching scenario in every clock cycle, for the whole duration of the simulation. This is closer to the operation of the circuit in the real world and due to the constant change of the switching portion of the design, high spatial coverage is achieved. In addition to that, by continuously triggering the same instances for the whole duration of the simulation only concentrates the switching activity in a small area, achieving low spatial coverage and evaluating only a small fraction of the PG.

3.2 Taking in account the locality of the phenomenon

Power has to feed every transistor within a chip. That power is distributed around the chip using the metal layers. As fabrication technologies have become smaller, the size of the wires has also been getting smaller, while physical chip dimensions have stayed roughly the same.

This means that wires have become thinner but have not gotten shorter.

That leads to an increase in resistance per unit length. There is an almost 10X increase in resistance between a 28nm chip and a 7nm chip, and that will follow an exponential increase for smaller geometries.

At the same time, the total power consumed by chip has also remained fairly flat.

As current flows through a resistor, the voltage drops – this is what is referred to as IR drop. When the voltage at a transistor drops, it becomes slower and this could impact the circuit timing. When this happens on a critical path through a design, it can also lead to a functional failure and thus needs to be avoided.

Another trend in semiconductor design has been a reduction in operating voltage, meaning that small changes in supply voltage may represent an increasing percentage of the digital swing and potentially lead to incorrect logic values being seen.

In extremely advanced technology nodes, even when an extremely small number of instances that are close to each other switch, significant IR Drop occurs that may lead to violations and/or malfunctions of the circuit. That is why we decided to pick small portions of the design in each cycle and concentrate the switching activity on these areas. This allows us to efficiently test the robustness of the PG and to identify any outliers early in the design cycle.

TEXT RETRACTED DUE TO CONFIDENTIALITY

3.3 Triggering a different portion of the design

In this section, a short and abstract analysis of the methodology is presented. Unfortunately, more info cannot be disclosed.

The user has the ability to divide a design into regions and produce different scenarios for each region. This provides the user with the necessary freedom to develop a scenario that will be very close to a real-world scenario. A portion of the design may require higher switching activity while a different portion may have lower. For example, a DSP application would demand plenty of Fourier transformations. Therefore, if a design is a DSP core, the user would expect to see a lot of switching activity in the area where the blocks that conduct the Fourier transformations are located and see lower activity in the separate floating-point unit.

These defined regions are then divided once again into smaller portions that cover a user-defined area. The user has the ability to define the shape and dimensions of these chunks. These chunks are then grouped through the utilization of a heuristic technique and only a certain number of these chunks is allowed to switch per clock cycle. Therefore, in each of these groups there are always some members that are toggling.

The user is also responsible to provide the methodology with a desired Toggle Ratio for the region that will follow the generated scenario and the desired switching activity that will be focused in a chunk.

Finally, the user must define the number of cycles so the methodology can generate a scenario for this particular number of clock cycles.

Our methodology divides the layout of the design into smaller areas

TEXT RETRACTED DUE TO CONFIDENTIALITY

3.4 Handling Clock Instances and Memories

It is mandatory to handle the clock instances and memories that exist in a design.

For the clock instances two different flavors were evaluated:

- Triggering every clock instance of our design in every cycle.
- Triggering every clock instance that lies on an "activation area" in each cycle.
- Allow the methodology work freely, taking advantage of the randomness.

Triggering every clock instance of the design in every cycle, would lead us to have 100% switching of clock instances in every cycle. This would increase the power consumption and would also cause our analysis to be unrealistic since in modern IC designs, clock instances do not toggle in every clock cycle through the utilization of smart design techniques, such as clock gating.

Triggering every clock instance that lies on an "activation area" in each cycle was also not applied. Through our experiments, we realized that the activation of that many clock instances in a small area affected negatively the ultimate goal, which was high spatial coverage and extremely high switching coverage. By switching that many clock instances in a small box, there was no room to trigger that many combinational and/or sequential instances that lied on the same area.

The final call was to allow the methodology to work freely, taking advantage of the randomness that is embedded in its process to decide which instances will switch.

It is **guaranteed** that each time a small area is chosen to toggle, a different set of instances is activated and each instance is **equally** likely to be picked.

For the memories, we produce scenarios that ensure that in each simulation cycle, the memories switch on a different custom state. Unfortunately, since the memories used in our design are fully custom and are considered a company competitive advantage, more info cannot be disclosed about their operation and the scenarios that are applied when these blocks are tested.

TEXT RETRACTED DUE TO CONFIDENTIALITY

Chapter 4

Experimental Methodology & Results

In this chapter I will be focusing on the experimental methodology that we used in order to benchmark our approach and compare it to three different approaches. Subsequently, results will be presented and analyzed.

Comparisons between the number of outliers that the analysis identified, switching coverage, runtime, peak current, peak cycle power and average cycle power shall be presented. Due to limitations that the confidentiality of this work introduces, all values will be presented compared to one another.

4.1 Experimental Methodology

In order to have a fair comparison and evaluation of the proposed methodology, every experiment was executed on the same machine, without any other processes running in the background so we can have a good estimation of the total runtime. The total runtime presented here is the mean of 20 execution times, without taking in account the minimum and maximum runtime out of those 20 runs.

Moreover, every design was implemented in the same technology node, following precisely each and every step of the ASIC flow without any variations or design specific ECOs.

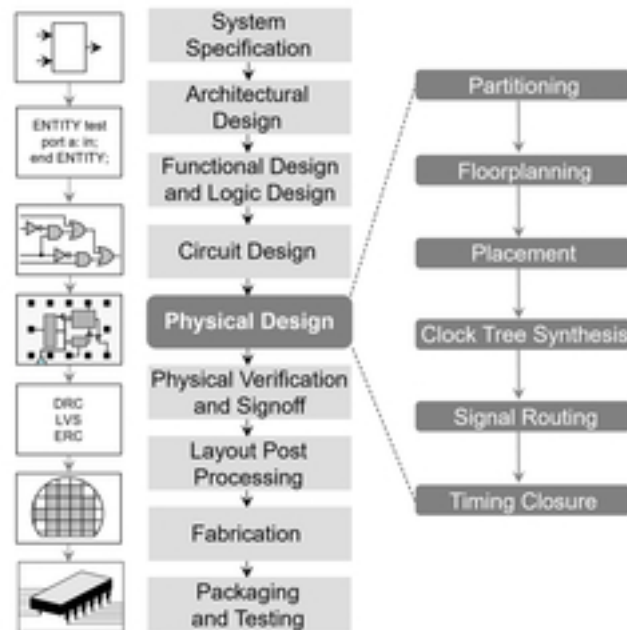


FIGURE 4.1: A flow diagram presenting the ASIC flow

4.2 Metrics

In this section, I would like to elaborate on the importance of each and every metric that was used in order to evaluate our methodology. To avoid misunderstandings, all metrics are considered equally important and there is no strong preference towards the "betterment" of a particular metric.

First and foremost, **switching coverage** demonstrates "how much" of the design is switching throughout the entire simulation. In order to stress the PG and efficiently assess its attributes and identify its weaknesses, we would like to ensure the highest possible spatial coverage to identify the maximum number of hotspots that may occur.

Secondly, **runtime** is an important metric that has to be taken in account so we can have the results of a simulation at a reasonable time. There is no point in executing something that will need an infinite amount of time to present its results, even if

those results would be the most accurate and realistic results possible.

Thirdly, the **total number of IR outliers** represents the number of instances that have a drop greater than a specific percentage, which flags them as violating instances. This percentage comes from measurements and data from the past years, which led to the "calibration" of the internal signoff methodologies.

Fourthly, the **peak current** is the maximum amount of current which output is capable of sourcing for brief periods of time. When a power supply or an electrical device is first turned on, high initial current flows into the load, starting at zero and rising until it reaches a peak value, known as the peak current. This is usually much higher compared to the steady-state current. The current will then decrease gradually from the peak current value to the steady state where it stabilizes.




FIGURE 4.2: The steady state, inrush and peak current when a device is turned on

The difference between the steady-state value and the peak current is what is referred to as the inrush current. This is a property that designers must pay attention to when selecting the components. The peak current should be within safe limits and only last for a short duration to avoid stressing the components or causing excessive heating or damage.

Fifthly, the **peak cycle power** represents the maximum power consumption in one clock cycle that occurred during our simulation. This power consumption has to be a reasonable value, otherwise someone may claim that the analysis is not realistic. For example, if in one clock cycle the design consumes, let us say 30Watts, the methodology is probably flawed and reaches an absurd power consumption level in one clock cycle.

Last but not least, the **average power** demonstrates the average power consumption of our design during the analysis. If we wish to have a simulation which is close to the real-world behavior of our circuit, then, this value should respect the power budget that the design specifications define.

4.3 Results

In this section, I will be presenting the results of our work.

In a similar order as the one that was used previously, where I elaborated on the significance of each metric, I will be showcasing the great improvement that was achieved when using our methodology on an industrial benchmark.

The following figure is a bar chart that shows the switching coverage that each approach managed to achieve and the total runtime of each approach as a separate line, shown on the same figure.

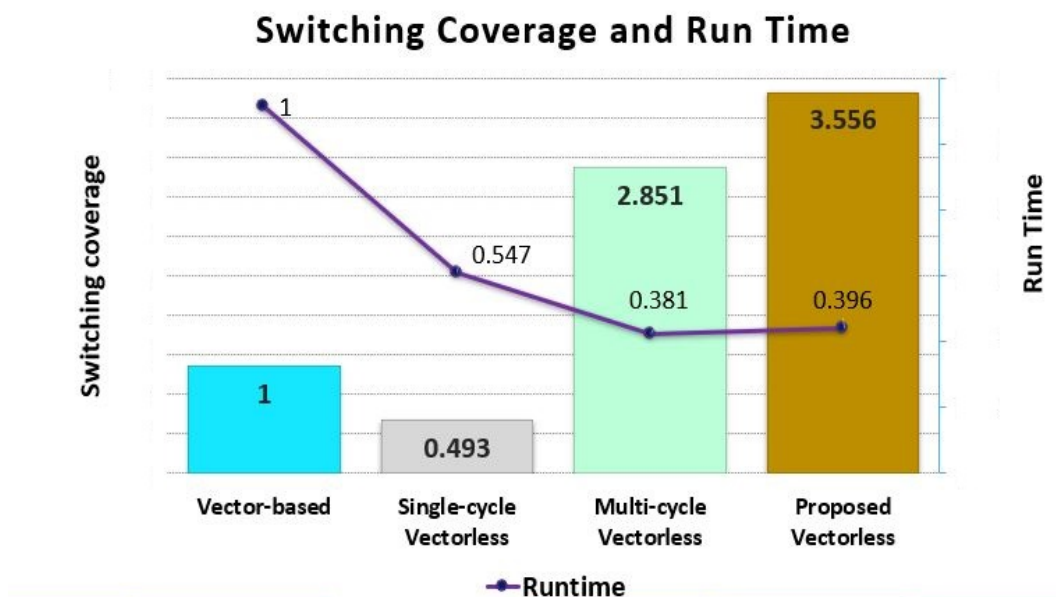


FIGURE 4.3: Switching coverage and runtime achieved by each methodology

This figure clearly shows the tremendous improvements on switching coverage that the proposed vectorless methodology achieves compared to all the other approaches. A single-cycle Vectorless approach manages to reach 0.493x the coverage of a VCD, the multi-cycle Vectorless achieves 2.851x and our proposed approach manages to cover 3.556x more than the Vector-based approach. Moreover, the impact on runtime is insignificant since the proposed methodology has a similar runtime as the native Multi-cycle vectorless, which is slightly faster in a design that is tremendously large. These results reassure us on the **scalability** that this methodology has and proves its ability to handle humongous IC designs.

The following figure presents a breakdown of the switching coverage of each methodology for the same benchmark. Disclosure of the total number of instances is not allowed.

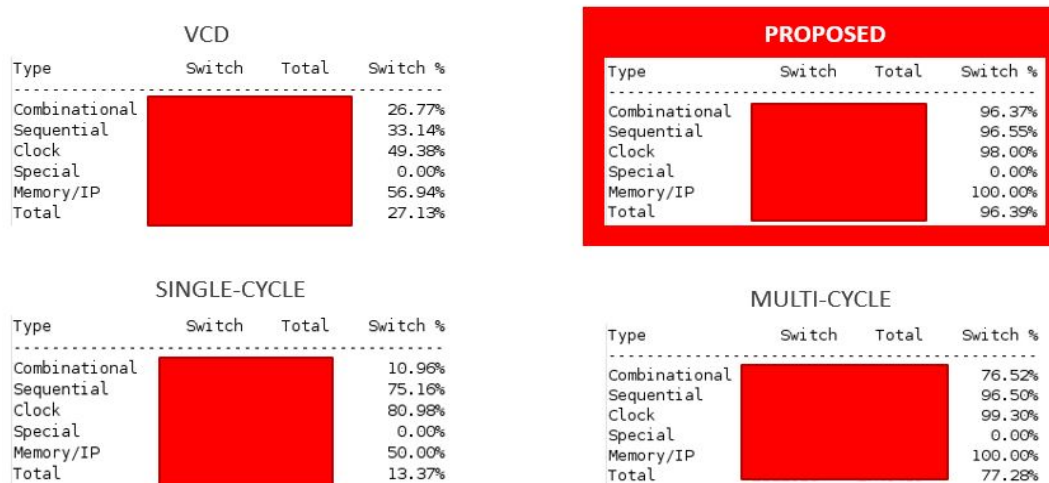


FIGURE 4.4: Switching coverage breakdown for each methodology

The above figure confirms that the achieved switching coverage that the proposed vectorless methodology causes, compared to all the other approaches, is better at benchmarking the PG. The single-cycle Vectorless approach manages to cover only 13.37% of the design, while only exciting 10.96% of the combinational logic of the circuit. The multi-cycle Vectorless approach improves the coverage and reaches 77.20% total coverage of the design, while exciting 76.52% of the combinational logic of the circuit. That is a decent result, however, why would someone settle for the better approach, when the best approach is available?

The proposed methodology covers 96.39% of the design, while switching 96.37% of the combinational instances, 96.55% of the sequential instances and every memory/IP that is instantiated in the design. Finally, the VCD, reaches 27.13% total coverage while exciting 26.77%, 33.14% and 56.94% of combinational, sequential and memory/IP instances respectively.

These results back us up on the argument that this proposed methodology excites almost every instance of the design and has the ability to identify **small localized hotspots** and efficiently **evaluate** the PG for weaknesses through the completeness of the analysis.

This next figure presents a bar chart that shows the total number of IR Outliers that each approach managed to identify achieve and the peak current of each approach as a separate line, shown on the same figure.

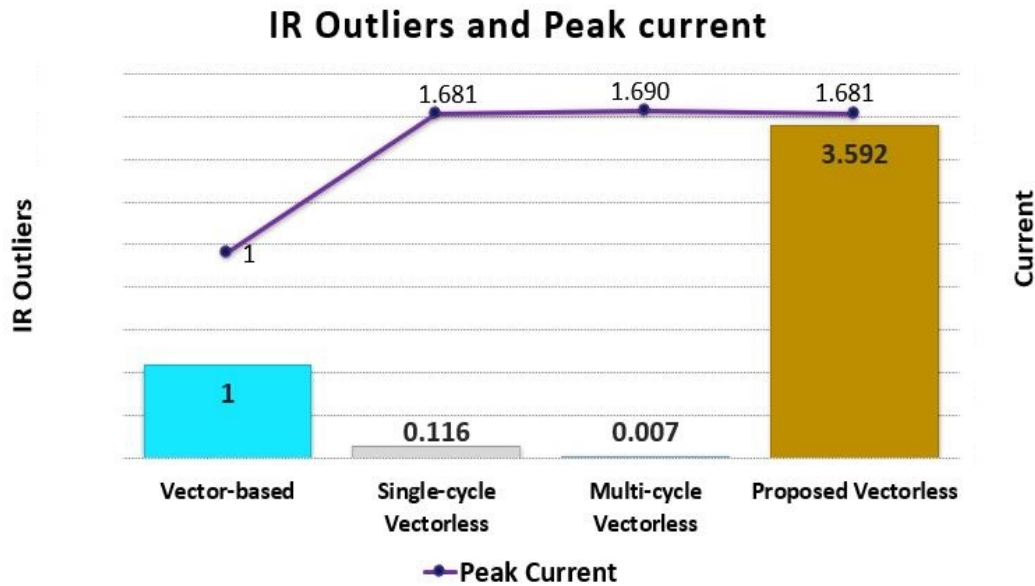


FIGURE 4.5: IR Outliers and Peak Current for each methodology

By examining the above figure, we can see that the total number of IR Drop Outliers identified by the proposed vectorless methodology is 3.592x the number of outliers identified by the vector-based approach, which is a simulation of a real-word application. In comparison, the number of outliers identified by the single-cycle and multi-cycle vectorless methodologies are 0.116x and 0.007x the number of outliers that the vector-based approach identifies respectively.

Does that mean that all these hotspots that were identified will cause malfunctions or that these drops will happen while the IC is operating and cause any timing issues or any other issue?

No, definitely no.

But, and this is a huge but, **no one can claim that these hotspots are not there**, or that they are fictional or false positives. And no one can guarantee that there are no paths that may be excited and cause issues. Therefore, a smart designer would prefer to extensively benchmark the PG of the circuit and ensure that the IC will not have any dangerous regions that may cause trouble.

In addition to that, we see that the peak current that the proposed vectorless methodology has insignificant differences with the two other vectorless approaches. Therefore, we see that the native flows of an EDA vendor and our methodology reach the same values and we can claim that our approach is safe to use.

This next figure presents a bar chart that shows the peak cycle power that each approach managed to identify achieve and the average power of each approach as a separate line, shown on the same figure.

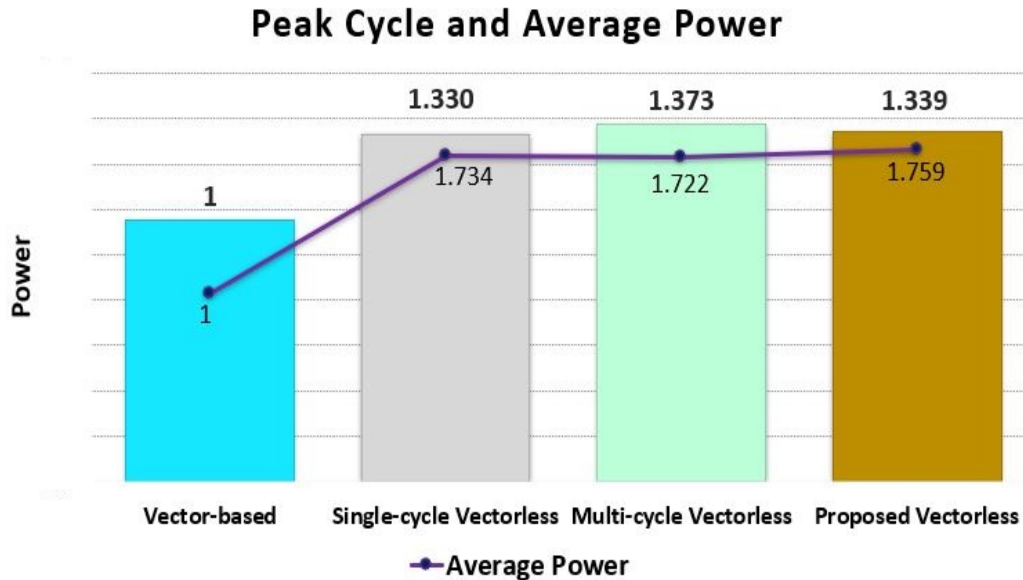


FIGURE 4.6: Peak Cycle Power and Average Power for each methodology

After evaluating the above figure and examining the design specifications, that cannot be disclosed, we can see that our simulations all respect the Power Budget. Therefore, no one can accuse us that in order to get more outliers we are increasing the power budget up to unrealistic values. This technique is commonly used from lazy engineers that do not want to examine thoroughly their design and are taking the easy route or do not have any proper methodology to address these challenges. Furthermore, the same applies for the peak cycle power that the simulations have. This guarantees that during the simulation, there was no crazy behavior like having a few extremely power hungry cycles and then the rest of the cycles in the simulation having minimal power consumption. That kind of behavior would compromise the validity of our analysis.

Unfortunately, due to the confidentiality agreement that we must respect, there is no way that IR Drop maps can be presented. Therefore these maps have been removed from this section, along with their analysis.



FIGURE 4.7: Omitted due to confidentiality

TEXT RETRACTED DUE TO CONFIDENTIALITY

TEXT RETRACTED DUE TO CONFIDENTIALITY

TEXT RETRACTED DUE TO CONFIDENTIALITY

Chapter 5

Conclusions

5.1 Conclusions

In this thesis we focused in exploring methodologies to identify localized IR drop hotspots that are encountered in cutting edge technology nodes the semiconductor industry uses.

Utilizing the locality of the phenomenon, that comes from its nature, we managed to provide scenarios that efficiently reproduce the behavior of the circuit as it operates in the real world and identify all the weaknesses that the PDN of the Integrated Circuit has. This allows us to provide feedback, face the challenges that arise and come up with solutions to make sure that our IC will be fully functional. Therefore, we ensure great yield, save money and have the opportunity to reassure the customer that the product will meet all the expected guidelines and requirements.

Moreover, since the proposed methodology does not depend on input vectors, we save time and eliminate the dependency of engineering teams to others. We are now able to ensure the robustness of the power grid and achieve extremely high coverage of the circuit in a simulation, without introducing any significant overhead or violating any restrictions such as the power budget.

To summarize, the proposed analysis methodology can improve the sign-off quality through the utilization of the locality of the IR Drop effect, identify structural PG weakness and provide realistic results. The impact on runtime is small and power budget violations were not seen on any experiment. Extensive switching coverage is ensured, achieving numbers greater than 90% and the ability to identify structural PG weaknesses is provided.

5.2 Future Work

In the future, there is significant value to extend the research presented in this thesis towards the following directions:

- Guarantee that the analysis respects the power budget.
Power budget violations were not encountered in our experiments, but that does not mean that there should not be a well formed and deterministic approach towards this.
- Hotspot driven VCD window selection.
There is tremendous value to the selection of a "good" portion of a vector in order to analyze a circuit. Limiting the Vmin and identifying which part of the input vector is the one that presents this result and selecting the portion of the vector that will manage to capture IR Drop hotspots instead of picking the portion that will cause our analysis to achieve maximum power.

Bibliography

- A. Mukherjee K. Wang, L.H. Chen and M. Marek-Sadowska (2002). "Sizing Power/Ground Meshes for Clocking and Computing Circuit Components". In: *Design Automation and Test in Europe*, pp. 176–183.
- Bhooshan, Rishi (2007). "Novel and Efficient IR-Drop Models for Designing Power Distribution Network for Sub-100nm Integrated Circuits". In: *International Symposium on Quality Electronic Design ISQED*, pp. 287–292.
- Bhooshan, Rishi and Bindu P. Rao (2007). "Optimum IR drop models for estimation of metal resource requirements for power distribution network". In: *Very Large Scale Integration VLSI-SoC*, pp. 292–295.
- C. Tirumurti S. Kundu, S. K. Susmita and Y. S. Change (2004). "A Modeling Approach for Addressing Power Supply Switching Noise Related Failures of Integrated Circuits". In: *Proceedings of Design, Automation and Test in Europe Conference*, pp. 1078–1083.
- Cadence Design Systems, Inc. (2002). "VoltageStorm Transistor-Level PGS User Guide". In: *N/A*.
- (2018). "Voltus IC Power Integrity Solution User Guide". In: *N/A*.
- D.A. Andersson L.J. Svensson, P. Larsson-Edefors (2008). "Noise-Aware On-Chip Power Grid Considerations Using a Statistical Approach". In: *International Symposium Quality Electronic Design ISQED*, pp. 663–669.
- Guptaa, P. and A.B. Kahng (2006). "Efficient Design and Analysis of Robust Power Distribution Meshes". In: *19th International Conference on VLSI Design*, pp. 337–342.
- H. M. Chen L. D. Huang, I. M. Liu M. Lai and D. F. Wong (2003). "Floorplanning with Power Supply Noise Avoidance". In: *Design Autom. Conf. Asia and South Pacific (ASP-DAC)*, pp. 427–430.
- I. Polian A.Czutro, S. Kundu and B. Becker (2006). "Power Droop Testing". In: *ICCD*, pp. 243–250.
- J. Gu, R. Harjani and C. Kim (2006). "Distributed Active Decoupling Capacitors for On-Chip Supply Noise Cancellation in Digital VLSI Circuits". In: *Proc. Symp. Very Large Scale Integr. (VLSI) Circuits*, pp. 216–217.
- K. Arabi R. A. Saleh, X. Meng (2007). "Supply Noise in SoCs: Metrics, Management and Measurement". In: *IEEE Design & Test of Computers*.
- K. Shakeri, R. Savari and J. D. Meindl (2003). "Compact Physical IR-drop Models for GSI Power Distribution Networks". In: *IEEE International Interconnect Technology Conference*, pp. 54–56.
- Larsson, P. (1999). "Power Supply Noise in Future IC's: A Cristal Ball Reading". In: *Proceedings of the IEEE Custom Integr. Circuits Conf.* Pp. 467–474.
- M. Popovich A. V. Mezhiba, S. Koseeby (2008). "Power distribution networks with on-chip decoupling capacitors". In: *FRIEDMAN*, p. 110.
- R. Saleh S. Z. Hussain, S. Rochel and D. Overhauser (2000). "Clock skew verification in the presence of IRDrop in the power distribution network". In: *Computer-Aided Design of Integrated Circuits and Systems 19.6*, pp. 635–644.

- Rius, J. (2013a). "IR-Drop in On-Chip Power Distribution Networks of ICs with Nonuniform Power Consumption". In: *IEEE Trans. VLSI Syst.* 21.3, pp. 512–522.
- (2013b). "IR-Drop in On-Chip Power Distribution Networks of ICs with Nonuniform Power Consumption". In: *IEEE Trans. VLSI Syst* 21.3, pp. 512–522.
- S. A. Moghaddam, N. Masoumi and C. Lucas (2005). "A Stochastic Power-supply Noise Reduction Technique Using Max-flow Algorithm and Decoupling Capacitance". In: *Int. Workshop System-on-Chip Real-Time Applications*, pp. 265–269.
- S. K. Nithin, G. Shanmugan (2010). "Dynamic Voltage Drop Analysis and Design Closure: Issues and Challenges". In: *11th International Symposium on Quality Electronic Design (ISQED)*.
- S. Pant D. Blaauw, V. Zolotov S. Sundareswaran and R. Panda (2003). "Vectorless Analysis of Supply Noise Induced Delay Variation". In: *Proceedings of IEEE International Conference on Computer-Aided Design*, pp. 184–191.
- S. Zhao, K. Roy and K. K. Cheng (2002). "Decoupling capacitance allocation and its application to powersupply noise-aware floorplanning". In: *IEEE Trans. Comput.-Aided Design Integr. Circuits Syst* 21.1, pp. 81–92.
- Shakeri, K. and J. D. Meindl (2005). "Compact Physical IRDrop Models for Chip/Package Co-Design of Gigascale Integration (GSI)". In: *IEEE Transactions on Electron Devices* 52.6, pp. 1087–1096.
- W. Zhao Y. Cai, J. Yang (2014). "Fast Vectorless Power Grid Verification using Maximum Voltage Drop Location Estimation". In: *IEEE Design & Test of Computers*.
- X. Xiong, Jia Wang (2010). "A hierarchical matrix inversion algorithm for vectorless power grid verification". In: *IEEE Design & Test of Computers*.
- Y. Ban, C. Choi (2014). "Analysis of Dynamic Voltage Drop with PVT Variation in FinFET Designs". In: *International SoC Design Conference (ISOCC)*, pp. 512–522.
- Zhu, Q. (2004). "Power Distribution Network Design for VLSI". In: *Wiley*.