

Chapter 2

Background

It is now time for the chip design and verification community to take a closer look at the semiconductor physics fundamentals, since Moore's law still holds true after a half century. Gordon Moore forecast in 1965 that transistor integration in chips will double every 2 years. One of the prime reasons the law still holds true, defying chip area congestions, is because of the drastic advancement of the integrated circuits fabrication and process technology. TSMC foundry is in fabrication process of 10 and 7 nm technology from early 2017.

A higher process node is definitely attractive as more functionality integration is possible in a smaller die area at a lower cost. Typically, technology scaling contributes to reducing gate delay, which in turn further contributes to increase frequency, increase transistor density, and reduce energy per transition.

However, in reality, this comes at the cost of exponentially increasing leakage power. Because the minimum gate-to-source voltage differential that is needed in CMOS devices to create a conducting path between the source and the drain terminals, which is known as threshold voltage, has been pushed to its limit. The reason is to keep pace with smaller transistor and interconnect sizes and increasing operating frequency. Though there are promising process integration technologies that are rolling out because of advancements in silicon-on-insulator (SOI) deployment on multigates like FinFET, FlexFET, and Tri-Gate-Transistors. These provide better electrical control over the conduction channel of CMOS devices and help reducing the leakage current as well as overcoming other short-channel effects. However, there are integration challenges with multigate devices in regular semiconductor manufacturing processes, which are subject to extensive research.

The fundamental semiconductor physics for threshold voltage that causes a new phenomenal shift to higher leakage power dissipation from the predominating dynamic power issues, are mostly due to the advancement in process nodes. In other words, the leakage power is a function of the threshold voltage and at smaller device geometries its contribution to the total energy dissipation becomes significant. The total energy dissipation in a CMOS circuit can be expressed as the summation of leakage and dynamic power integration over time t as shown below.

$$E = \int_0^t (CV_{dd}^2 f_c + V_{dd} I_{leak}) dt.$$

Here the leakage power is directly contributed from device supply voltage and leakage current;

$$Leakage\ Power = V_{dd} I_{leak}.$$

While dynamic power is contributed from the switching activity of the capacitive load on supply voltage and its switching frequency;

$$Dynamic\ Power = \alpha CV_{dd}^2 f_c, \text{ where } \alpha \text{ is activity factor.}$$

It is evident from these equations that minimizing dynamic power requires minimizing switching capacitance and switching activity; whereas minimizing leakage power requires increasing threshold voltage and improving process technology.

But again it may cause higher leakage power dissipations unless new inventions, like multigate devices, are available in advanced process nodes.

However, reducing supply voltage V_{dd} provides controllability for both leakage and dynamic power dissipation. This became a major power management key in the last few years as voltage is a design parameter where the chip design and verification community has more controllability compared to the integrated circuit fabrication process technology that relies more on foundries.

Several power dissipation reduction techniques have been adopted over the past few decades for addressing both the leakage and the dynamic power, namely clock-gating, multi-threshold, body-bias, and transistor-sizing. Despite of their wide deployment for reducing dynamic and leakage power, each has its own limitations and complexity. Since power or supply voltage is the key today, more advanced power management and reduction techniques have been adopted in the last few years based on reducing or controlling the supply power. The mainstream techniques adopted today, as shown below, are mostly based on design type and complexity demands for system-on-chip (SoC), ASIC, microcontroller units (MCU), or processor core design implementations.

Example 2.1 Sample List of Mainstream Power Dissipation Reduction Techniques

- Power Gating or Power-Shut Down
- Power Gating with State/Data Retention
- Low power Standby with State/Data Retention
- Multi-Voltage Design with Performance on Demand
- Dynamic Voltage (and Frequency) Scaling
- Adaptive Voltage (and Frequency) Scaling

Although the names of these techniques suffice to understand their operations and objectives in any of the above-mentioned design implementations, however,

adoptions and verification of these techniques were impossible until the UPF or similar power formats, like CPF,¹ became available from 2007 and onward. Although CPF is not an IEEE standard itself, many of the CPF semantics were contributed to the IEEE-1801 committee to help standardize the UPF.

2.1 The Power Intent

The mainstream power management and reduction techniques listed in the previous section are solely based on direct manipulation of voltage, in terms of supply power connectivity and voltage area or power network distributions on the chip. Nevertheless, none of these were sufficient to understand and reflect the power aware verification plan or power intent to start power or voltage aware verification at the register transfer level (RTL) or even after gate-level synthesis.

Usually, RTL in Verilog, VHDL, or SystemVerilog (HDL in general), are the golden reference design, and it is totally unconventional to add supply networks and their corresponding connectivity to these references. Eventually, it is also impossible to distribute certain instances of a design hierarchy to specific voltage areas unless the design reaches the floor-planning stage of design implementation. Evidently the industry was facing the absence of a methodology that would help to define supply power connectivity, voltage areas, and power network distributions in a design, without intervening with the HDL reference design and possibly starts from RTL of the design abstraction levels. Thus adoption of any of the new or mainstream power controllable design implementation techniques through a power specification was just impossible at any design abstraction level, from RTL to place-and-route (P&R). As well, verification on power network connectivity, voltage area distribution, etc. was also impossible unless post P&R power-ground (PG) connected netlist are available at a much later part of the design implementation cycle.

2.2 The Abstraction of UPF

In early 2007, the Accellera Systems Initiative introduced the Unified Power Format (UPF), also known as UPF 1.0 that allows users to define and manage power for designs without any direct interference with the golden HDL reference design. The concept of UPF inherits, from deploying the power intent of any design directly overlying the HDL references. The overlay needs to be done without direct intervention, through a methodological approach to abstractly model the exact power supply networks, voltage, or power area distributions and corresponding power states. The methodologies are based on power specification or intent at the RTL and allow

¹CPF: Common Power Format Specification is an initiative from Si2 (Silicon Integration Initiatives) organization.

designers to continuously refine the power network distributions of voltage areas on the power domain boundaries throughout the entire design implementation flow, specifically at the post synthesis and post P&R phases.

In general, UPF provides the notion of power management of the entire design from the design implementation and verification perspective. Hence power aware (PA) design verification and implementation automation tools are also built on UPF semantics and language references.

The early UPF 1.0 initiatives were mostly driven from the physical design perspective in terms of explicit power supply networks, supply ports, supply nets, and their power states. These physical entities are mostly absent at the RTL or higher levels of design abstraction, unless the design has already been rolled out for synthesis and P&R. This contradicts as well as delays the original power management and verification objectives that possibly require to begin at least from the RTL.

In 2009, the IEEE Standard Organization published the IEEE-1801-2009 or UPF 2.0, the first standard methodology for implementing power intent, truly for any design abstraction level. The UPF 2.0 first provides the notion of supply sets, an abstract collection of power supply nets combining power, ground, and bias functionality that represents a potential supply net connection to a corresponding portion of a design. Supply sets can be incrementally refined and extended for different functionality, depending on design abstraction level starting from RTL.

Apart from unleashing supply sets, the IEEE standard itself evolved overtime, providing additional abstraction flexibility for the other rudimentary parts of design elements, like supply power, supply networks, and supply states for design groups, models, and instances, collectively known as objects. Some of the amendments are reflected in IEEE Standard 1801-2013 or UPF 2.1 published in May 2013, while the latest update is available in IEEE Standard 1801-2015 or UPF 3.0 published in December 2015.

However, it is important to note that each amendment is not necessarily backward compatible and each updated version is usually a superset of semantic and syntactic expressions from its predecessors.

The abstraction of UPF 3.0 or UPF in general, methodically consists of power specifications that categorically itemize the design elements for a particular power block, known as a power domain (PD), along with the PD boundaries. Further, it also specifies the power supply and supply states of the power domains or supply sets, whether the supply is in On, Off, or other potential states. Depending on the power management and reduction techniques the specification adopts for the final chip, the list further extends to specify the requirements of boundary strategies for intra or inter-domain communications. These strategies may include isolations (ISO), level-shifters (LS), enable level-shifter (ELS), always-on buffers (AOB), feed through buffers or repeaters (RPT), diode clamps, retention flops (RFF), power switches (PSW), and their corresponding supply network and locations details. In short, the UPF is a precise map that models the power specification of the design and converts the same to a power aware design.

To note, the strategies in UPF actually refer to special power management multi-voltage (MV), or power aware (PA) cells, and they are usually physically inserted into the design in post-synthesis or post-layout.

Figure 2.1 shows the HDL reference based design block diagram. Among several design blocks or instances only part of the design is distributed in four different power domains just overlaid on the blocks without intervention in the instance itself.

Figure 2.2 shows a typical UPF layout for the specific design of Fig. 2.1, where the PD_top represents the default UPF top power domain which contains the cpu_top




Fig. 2.1 Design block diagram with only a portion of the design overlaid by power domains with specific design instances




Fig. 2.2 Fundamental building block of UPF power domains, domain boundary, power network, and relevant strategies

design blocks as element. The PD_sub1, PD_sub2, and PD_sub3 power domains and PD_sub3.1 (the sub domain of PD_sub3) represent specific hierarchical design instances as their elements like udecode_top, ufetch_top, umem_top, and umem_top/umem_sub respectively. Hence the instance ualu_top naturally goes to the default PD_top power domain. The power domains also illustrate their respective ON or ON-OFF status and UPF strategies like PSW, ISO, LS, ELS (combination of ISO and LS), and RFF etc. are also shown with symbolic representation accordingly.

Epilogue: The **Chap. 1 Introduction** and **Chap. 2 Background** are the foundations of this book and are self-explanatory prologues. While introduction is purely intended to provide a brief overview of low power design and power aware verification concepts, this chapter, elaborated few fundamental aspects for power aware verification for the readers to comprehend and carry all over the book. The first crucial aspect is the list of **Mainstream Power Dissipation Reduction Techniques**. The second aspect is – power intent and how the manipulation of voltage on a design is mapped and controlled through the power intent “UPF”. The third is the abstraction of UPF – very basic concepts of power domains on a design and its surroundings.

Prologue: The next chapter, **Modeling UPF** is the entrance to the power aware verification world. This chapter is dedicated to provide readers the understanding of fundamental construct of UPF. The explanation starts with the detail elaboration of power domain and domain boundary as the concepts are directly inherited from HDL (Verilog, SystemVerilog and VHDL) perspective. The fundamental constructs further explains the power supply networks, power states and power strategies in light of both the UPF Language Reference Manual (LRM) and practical usage perspective. Consequently the methodical approach for UPF construct level incremental refinements and flow level successive refinements and hierarchical UPF are also explained with realistic examples.

Low-Power Design and Power-Aware Verification

Khondkar, P.

2018, XV, 155 p. 19 illus., 12 illus. in color., Hardcover

ISBN: 978-3-319-66618-1