2016 DESIGN AND VERIFICATION™ DVCON CONFERENCE AND EXHIBITION EUROPE

Addressing the Complex Challenges in Low-Power Design and Verification

Madhur Bhargava, Mentor Graphics, India (madhur_bhargava@mentor.com)

Durgesh Prasad, Mentor Graphics, India (durgesh_prasad@mentor.com)

Jitesh Bansal, Mentor Graphics, India (jitesh_bansal@mentor.com)

***Abstract* - One of the often-stated "fact" in chip development process is that debug consumes about 50% of development time and effort. Debugging is considered to be one of the toughest challenges faced in the semiconductor chip design and verification industry. Low-Power design and verification is now de-facto in the semiconductor industry. When it comes to low-power design and verification, the debug challenges are further complicated as a result of the sophisticated power management architectures and techniques that are used. Moreover, the traditional debug technology and methods focus on issues found in a design working in always-on mode and fails to address the new and complex power-related issues thereby consuming more engineering time. This paper focuses on providing a comprehensive analysis of various complex debug problems faced in low-power design and verification. By using relevant examples we will demonstrate how these issues can be either avoided or easily solved. We will also highlight some of the common pitfalls that low-power designers can avoid which otherwise can lead to complex low-power issues that are difficult to debug at later stages of the design cycle.**

Keywords - Power Management, Power Aware Verification, Debug, Debugging Challenges in Low-Power Designs.

I. INTRODUCTION

Power is one of the most important dimensions of complex modern chips. Designers use complex power aware techniques such as power gating, voltage scaling, and body biasing to save power and minimize heat dissipation. As a result the power-aware verification of chips is a fairly complex process. It is well known fact among design and verification engineers that the time spent trying to debug power management related simulation failures of a large SoC are fairly high and is a very difficult task to complete. The evolution of the UPF standard, from the original Accellera 2007 UPF 1.0 LRM, followed by the initial IEEE 1801-2009 UPF 2.0 LRM, updated by the release of IEEE 1801-2013 UPF 2.1 LRM and finally IEEE 1801-2015 aka UPF 3.0 LRM, has provided many new capabilities that have eased the power intent specification process as well as enabled new power management verification flows aligned with the needs of IP based SoC design today. Unfortunately the evolution of the UPF standard alone has not reduced the complexity of the power management verification task whatsoever.

Debugging of low-power designs still requires a combination of the brightest and best technical resources and requires world class power management-centric debug environments to perform the task. The difficulty of power management debug is reflected in the large number of EDA vendors that provide not only power aware front-end simulation and emulation capabilities but also by those focused on power management verification tools ranging from static analysis and rule based power checks, to power aware logic equivalency checking and layout tools. While there are debug challenges in verifying the behavior of a power managed SoC from the front-end design phase all the way through the back-end implementation phase, this paper focuses on the simulation based debug challenges that are likely to be encountered.

Power management debug has permeated all aspects of traditional HDL based simulation and nearly all major EDA vendors provide the following power specific debug features: visualization of UPF created objects, HDL signal colorization and highlighting in RTL source and wave windows, dynamic informational/status power related messages, automatic power specific assertion generation and power report generation. Below are some of the low-power debug challenges divided into various groups. While none of the debug challenges discussed in this paper will be new to those involved in power management verification, the intent of this paper is to provide a practical guide to the uninitiated based on actual user experiences.

II. UNWANTED X ON SOME SIGNALS

One of the most common debug problem in low-power design simulation is an unwanted 'X'. Debugging of such problems need some advanced features from verification tools. The most useful features in this regard are as follows.

A. Debugging 'X' values on signal

A major debug problem in low power simulations is root-cause analysis of unknown (X's) values. There can be several reasons why 'X' values appear on signals in low power simulations. These reasons can be either incorrect UPF specification (missing isolation/level-shifting retention cells or improper power domain partitioning) or UPF 1.0 to UPF 2.0 simulation sematic differences because of usage of initial blocks (re-evaluation and races between signal initialization and domain power up). In order to distinguish a normal unknown value on a signal, most simulation tools have the ability to highlight unknown values in wave windows caused by power domain corruption. In normal simulations, unknown X signal values are typically displayed using either a single mid-high red line or a red outlined box around the unknown value region in a wave window. In low power simulations the entire low-high region is filled in using red colored coarse or fine-grain cross-hatching to indicate the X value is a result of direct power domain corruption as shown in the figure below.

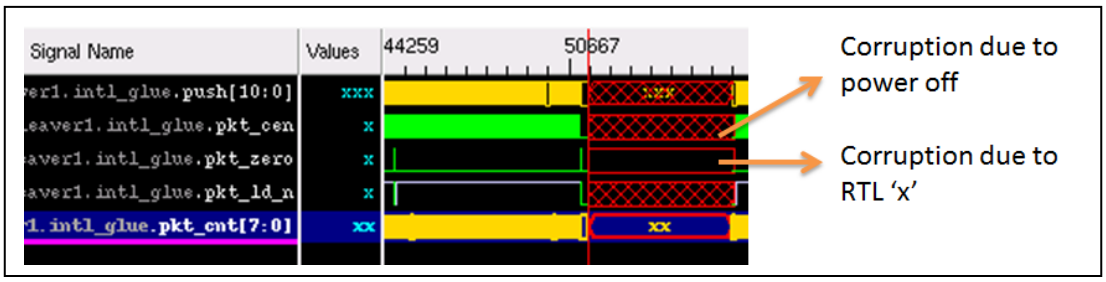

Typically the corruption highlighting is only visible on the direct output of the driving logic and is not displayed on the outputs of subsequent logic the corrupted net fans out to. Even when an unknown X value signal is not highlighted in the wave, the ability to trace the net connection backwards to where a highlighted signal is visible can greatly simplify the task of determining the origin of the X value. Most of the EDA vendors provide wave compare tool to identify root cause of these unwanted X values. Consider the code previously shown in the initial block revaluation section where 'X' value on the clk signal is because of power domain getting OFF. In this scenario, the clk can't toggle since the initial value is X instead of a 1'b0. If all the necessary power domains are functional, corruption highlighting won't be of much use in catching the X on the clk. However, wave compare between a normal and low power simulation could easily catch the X on the clk. Another area where wave compare is effective is right after power up a domain especially if retained register values need to be restored for full functionality. It is pretty easy to see differences in signals which remain X as a result of a missing retention element.

If wave compare is not available, another useful technique for catching unwanted X values is correlating their occurrence with changes in power domain sim-states or power states, power control signals including those for power switches, isolation enables, and retention save/restore signals. Most power aware simulation vendors provide the ability to print low power informational related messages. These messages typically are about the change of state in supply nets/ports, power switches, various power control signals, and power domains.

```
    Note: UPF_SWITCH_CTRL_INFO: Time: 794845 ns, Power Switch (glue_sw), Control Signal
(sw_en), switched to polarity (1), Power Switch state (FULL_ON)
```

In addition to informational messages many low power simulation vendors also provide the ability to create assertions to help in detecting sources of X values due to issues related to power control sequencing. For instance, turning off the retention supply to a power domain while it's OFF will corrupt the retained register values and cause an X values to be restored to the retention elements after power is restored.

```
    Error: UPF_PG_CHK:784845 ns, Power for Retention strategy: 'p_ret' of power domain:
PD_gluelogic' is switched OFF during retention.
```

Besides low power assertion checks to catch possible sources of X values during simulations, some vendors also provide a static analysis capability, when processing the UPF file, to determine if there are any missing, redundant, or invalid isolation/level shifters which can help catch them as well. An un-isolated power domain port will obviously propagate an X value to other power domains when it's powered down.

B. Trace the driver of the signal

Typically in a non-power aware simulation the driver of a signal is some RTL logic. However in the case of a power-aware simulation, it could be anything ranging from RTL logic to UPF inserted cells. For the cases of an unexpected value on a signal, it might be the effect of power-aware activity on that signal or it could be a propagated effect of power-aware activity on some of its driver signal/logic. In these cases the driver-tracing is a useful way to find out the driver signal and its value. Here is an example from Questa for finding out the driver signal/logic of '`/tb/out2`' which is an assign logic within the hierarchy '`/tb/TOP`'.

Tool Information – drivers /tb/out2

```
# drivers /tb/out2
# Drivers for /tb/out2:
#    StX  : Net /tb/out2
#      StX : Driver /tb/TOP/#ASSIGN#61
```

It is evident from the drivers information that 'x' on out2 is because of an assign statement going to 'x' (preferably due to power aware corruption). Another useful feature is dataflow/schematic debugging of a signal. It helps in tracing the value of a signal which is being driven from a distant logic. In power-aware debugging it is even more helpful as it also displays the UPF inserted cells in the dataflow path. Here is an example of dataflow from Questa showing an isolation cell and a level shifter encountered in the dataflow path of signal '`/tb/out1`':

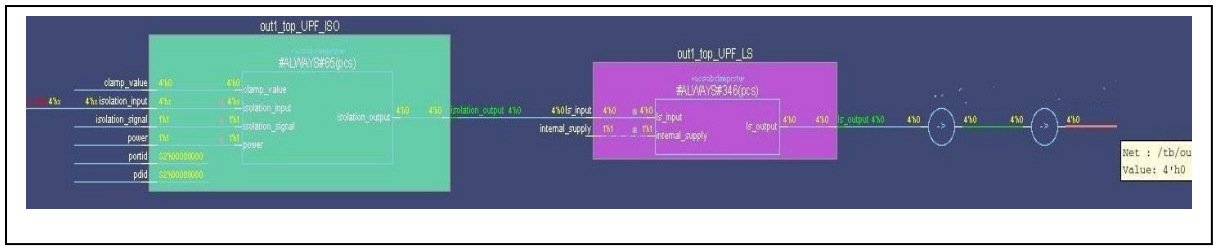

Figure 5: Dataflow of signal /tb/out1

III. SOME SIGNALS IN DESIGN ARE NOT CORRUPTED

One very common low-power debug issue is that certain part of design fails to switch off and the logic inside that part is never corrupted, although the user expects it to be switched off and show corrupted values. There could be multiple reasons for such behavior, some of which are as follows:

1) Incorrect power domain specification

The design element under concern has been put in a power domain that is not switchable. The first step for debugging such an issue is to identify the power domain to which this region belongs. This can be done either by looking at the power aware textual reports generated by the tool or using the tool's GUI capabilities. The next debugging step is to determine if the power of that domain is being switched off or not. It can be verified using the dynamic messages reported by the tool for whenever a power domain changes its state:

```
   # ** Note: (vsim-8902) MSPA_PD_STATUS_INFO: Time: 64 ns, Power domain 'pd' is powered
down.
```

2) Exclusion of behavioral model

EDA vendors provide mechanisms to skip the corruption on certain portion of the design. Some of the commonly used methods are by specifying the elements as DONT_TOUCH elements via UPF itself or use a separate file to exclude them from power-aware behavior. In such cases it is advised to look for the tool generated reports to find out if the design element has been excluded by the tool or not.

3) Simulation semantic disabled

Disabling simulation semantics of a design element means the tool will not impart any power shut-off/corruption on that element. The simulation semantics of a region can get disabled because of following reasons:

a) The design element is already a power-aware (All-pins model) [2] and it has its own power pins which are connected via UPF.

b) The design element is an 'always_on' cell as per its liberty specification.

c) The design element is a power-aware cell (isolation/level shifter) as per its liberty/HDL definition however it could not be mapped to any UPF strategy.

d) In all these cases the tool disables the power aware simulation semantic and hence there would be no tool injected corruption. All such regions can be easily identified by looking at the message thrown by the verification tool:

** Note: (vopt-9693) Power Aware simulation semantics disabled for chip_top/u_hm_top_0/u_ip_1

IV. SYSTEM TRANSITIONING TO ILLEGAL POWER STATE AND TRANSITIONS

Today's low-power designs have large number of operational modes. One of the major debugging tasks for low-power design is verification of the design's operational power states. This requires verifying that each defined power state of every power domain has been covered and functioning properly. It also requires verification of all power state combinations across all domains that comprise each operational power state. The complexities of the verification process increases multifold times as designs continue to increase in both the number of power domains and operational power states. While the Accellera UPF 1.0 standard supported the creation of power states using the three Power State Table (PST) commands, `add_port_state`, `create_pst`, and `add_pst_state`, there were several limitations that were not addressed in the standard which are:

a) No support for hierarchical power state definition/creation

b) No ability to update power state information

c) No ability to distinguish between legal and illegal power states or detect transitions to and from legal and illegal power states.

d) No support for bias states. Only supports the basic OFF (corrupt) and ON (normal) power states

e) No ability to merge multiple PST's together

Due to the many limitations inherent with UPF 1.0 power states, most low power simulation vendors extended their Finite State Machine (FSM) coverage capabilities to include both power states and power states transitions to do the job. The IEEE UPF 2.0 standard addressed all these limitations by providing the `add_power_state` and `describe_state_transition` commands. Not only does `add_power_state` support bias states, hierarchical power state creation, and an incremental update capability, it also allows any named power state to be declared as legal or illegal. Using these two commands require low power simulation vendors to issue run-time error messages whenever an illegal power state is reached or any illegal power state transition occurs.

However UPF 2.0 command "`add_power_state`" proved to be too much flexible with no way to specify the kind of objects it is being applied to, no restrictions on usage of objects in supply and logic expressions. It also lacked a way by which design engineers can provide information about the relationship of states of two objects. As a result of this flexibility, the debugging process of power states became even more complex. To address the problems caused by these flexibilities, the add_power_state definitions and concepts got updated and further revised in UPF 2.1 and UPF 3.0 respectively.

The UPF 2.0 standard also stipulates that an unnamed or undefined power state is also illegal. The detection of undefined power states is especially useful if an unintended power state occurs when transitioning from one defined state to another. The occurrence of an undefined power state during a legal power state transition may result due to a race between changing a UPF supply via the UPF package `supply_on/supply_off` functions and switching of a power control logic signal at the same time. A race induced undefined power state likely indicates an area where voltage ramp up/down times versus logic switch times must be accounted for in order to ensure proper operation of the design.

UPF Code

```
add_power_state PD_ALU_SS -state ON4 { -logic_expr { !pwr_alu && !pwr_ram } -simstate
CORRUPT -illegal}
```

Simulation message

```
# ** Error: (vsim-8933) MSPA_UPF_ILLEGAL_STATE_REACHED: Time: 129 ns, Supply set
'PD_ALU_SS' reached an illegal power state 'ON4'.
# File: src/parser_test22/demo.upf, Line: 73, Power state:ON4
```

However at times, user skips to define all the fundamental power states of an IP and proceed for simulation, which resulted unexpected illegal state messages. User can make use of UPF 3.0 syntax which allows the specification of the set of power states for a given object to be marked as *complete*, which indicates that all fundamental states of the object have been defined as named power states. If the set of power states for an object is complete, then it shall be an error for the UNDEFINED power state to be the current power state of that object. It is also an error if a new fundamental power state is defined after the power states are marked complete. If the power states of given object are not marked as complete, it is assumed that all fundamental states have not been defined and will not be marked as erroneous behavior.

Likewise, the describe_state_transition (UPF 2.0/2.1) and add_state_transition (UPF 3.0) command allows any transition between two power states to be declared legal or illegal. UPF 3.0 also allows to define a group of related power states which can then be used in add_power_state command. The legal power states of a power state group define the legal combinations of power states of other objects in this scope or the descendant subtree, i.e., those combinations of states of those objects that can be active at the same time during operation of the design. This command can be used to define the illegal power state combinations.

UPF Code

```
create_power_state_group CPU_cluster
add_power_state -group CPU_cluster -state {RUN1 -logic_expr {CPU0==RUN && CPU1==SHD &&
CPU2==SHD && CPU3==SHD}} -illegal
```

V. Power Intent Specification Complexities

The specification of power intent for power management of low power designs has been addressed by the UPF (Unified Power Format); however the UPF standard is still evolving with new features, concepts and clarifications being added over the releases. It often poses problems related to backward compatibility, differences and migration issues which are then difficult to debug.

A. UPF 2.0 Migration Issues

In UPF 1.0, the UPF supplies defaulted to the ON state. Many verification and design engineers involved in UPF 1.0 based power aware simulations have unknowingly relied on this fact and as a result there has been a tendency to not use the UPF package defined `supply_on` function to explicitly turn them on. It is interesting to note that a previously passing UPF 1.0 based simulation might fail after switching to UPF 2.0 based power aware simulation semantics. A common reason for many of these failures is that in UPF 2.0, the UPF supplies default to the OFF state causing all power domains to be in a CORRUPT simstate. This common migration issue can easily be avoided by using the UPF package defined `supply_on` function to explicitly set both UPF state and voltage values for all the created UPF supplies. The `supply_on` and `supply_off` function, as well as other functions defined in the UPF package, are available for use by placing the import statement in the simulation test bench, as shown below.

HDL Code

```
module tb;
import UPF::*;
…
initial begin
   supply_on ("tb/dut_inst/VDD", 1.1);
   supply_on ("tb/dut_inst/GND", 0.0);
end
…
dut dut_inst (…);
…
endmodule
```

As shown above, the `supply_on/supply_off` functions are commonly placed in an initial block. The only issue involved with the use of the UPF defined `supply_on` function is passing the proper arguments to it. The `supply_on` function takes two arguments; the first is a string type that contains the full hierarchical path to the desired UPF supply_port or supply_net that needs to be turned on while the second argument is the real type voltage value. The example above, assumes that both the "`VDD`" and "`GND`" UPF supply_ports have been created in the UPF scope corresponding to the `dut_inst`. Another common mistake, even if the `supply_on` function is used to turn on the power supply_port, is not using a `supply_on` function to also turn on the ground supply port. Please note that that both power and ground supplies must have their UPF state set to FULL_ON in order for a power domains simstate to be in the Normal simstate.

One other potential UPF 2.0 migration issue is due to isolation of lower boundary power domain ports. In UPF 1.0, the `set_isolation -applies_to inputs/outputs` port filters only considered power domain ports that were aligned on a module boundary. In other words only the input and output ports of modules were isolated. In UPF 2.0, these isolation port filters have been extended to also include the lower boundary (child module instances in different power domains) input and output power domain ports as well. The concept of lower boundary power domain port isolation can potentially cause simulation failures as a result of unintended back-to-back isolation cells inferred by tools, especially if isolation strategies were included for any lower-level power domains.

B. UPF2.0 list/wildcard expansion issues

Another debug challenge in power intent specification arises because of the usage of list/wildcard expansion in various UPF commands. It often happens that an incorrect list of signals is created as a side effect of usage of the wrong pattern in wildcards. This problem can be avoided if the user relies on the UPF command "`find_objects`" to create a list of signals and uses the tcl command "`puts`" to print the contents of the element list. Another way to get the list of expanded signals is to use the "`save_upf`" command. Consider the following design example:

HDL Code

```
module dut;
…
Ip_module my_ip1();
Ip_module my_ip2();
Ip_module my_ip3();
endmodule
```

In this particular case `my_ip1` and `my_ip2` have same power requirements; however `my_ip3` has some different power requirements. The intended UPF usage is to add `my_ip1` and `my_ip2` to one power domain, while creating another domain for `my_ip3`. The following incorrect and Non-LRM usage of a wildcard can result in the addition of all three instances of `Ip_module` to same power domain.

UPF Code

```
create_power_domain pd_dut -elements {my_ip*}
```

Though correction of such an issue requires a change in the UPF itself, however the first and foremost problem is to know what all elements have been added to `pd_dut`. In certain cases, some verification tool may explicitly list out the expanded elements; however that is tool dependent and may not be supported by all the tools. One of the ways to know which elements are being added to power domain `pd_dut` is to use the "`save_upf`" command in the UPF file. With this UPF command, the verification tool will dump out interpreted commands in a new UPF file. This UPF file will contain lists of expanded elements. Another way could be to avoid using Non-LRM wildcard usage and rely on UPF command `find_objects`. As shown in the below UPF code, the user can get the list of expanded elements in a tcl variable and print out that variable to know the expanded list.

UPF Code

```
set my_element_list [find_objects . -pattern {my_ip*}]
puts "Elements added to pd_dut : $my_element_list"
create_power_domain pd_dut -elements $my_element_list
```

Stdio output as a result of processing above commands:

```
Elements added to pd_dut: {my_ip1 my_ip2 my_ip3}
```

C. Incorrect vct specified

In case of a macro model in the design, the power supplies along with the power aware functionality are present inside the model itself. When integrating this macro model in a SOC, the integrator has to connect these HDL supplies to the UPF nets. Since the UPF supplies are of type supply_net_type having state and voltage values and the supplies defined in HDL are of wire type, the connection between UPF and HDL net requires a VCT (value conversion table). The VCT defines the mapping between the state of UPF net and value of HDL port/net. The user can either rely on verification tools to apply a default VCT or explicitly specify which VCT to be used using the UPF command `connect_supply_net -vct`. The problem arises when the same VCT gets used for power/ground/pwell/nwell supply nets. This causes power up failure because ground/nwell supply nets are active low and expect a ground specific vct. Consider the design example:

HDL Code: Power Model

```
module macro_model(in1, in2, out);
…
wire VDD = 1;
wire GND = 0;
if (VDD && !GND)
    out = in1 & in2;
else
    out = 'x';
endmodule
```

UPF Connections

```
connect_supply_net upf_VDD -ports {hm_inst/VDD} -vct UPF2SV_LOGIC
connect_supply_net upf_GND -ports {hm_inst/GND} -vct UPF2SV_LOGIC
```

Here "UPF2SV_LOGIC" is a predefined vct by UPF to convert UPF supply_net value to SV logic value as per following rules:

```
create_upf2hdl_vct UPF2SV_LOGIC
    -hdl_type sv
    -table {{UNDETERMINED X}
            {PARTIAL_ON X}
            {FULL_ON 1}
            {OFF 0}}
```

Since wrong VCT(UPF2SV_LOGIC) is getting applied on ground net (GND), the value driven on HDL net GND is "1" when the UPF supply (upf_GND) is FULL_ON. In order to correct this issue, the user needs to specify the ground specific VCT.

UPF Connection

```
connect_supply_net upf_GND -ports {hm_inst/GND} -vct UPF_GNDZERO2SV_LOGIC
```

If the user is relying on verification tools to apply the VCT, then tools needs to be guided that the particular supply is a "ground" supply. This information can come from the liberty file attribute "`pg_type`" or it can be specified using the UPF command "`set_port_attributes`" as follows:

pg_type example (Liberty Specification):

```
pg_pin (GND) {
    pg_type: primary_ground
}
```

set_port_attribute (UPF Specification):

```
set_port_attributes -pg_type primary_ground -ports {hm_inst/GND}
```

VI. SUPPLY NETWORK RELATED ISSUES

At the implementation stage, the supply network in a power-aware design is often huge and highly complex. At times, it gets buggy and difficult to debug the improper connections or any other issue in the supply network (either present in just UPF files or already implemented into the design). The complexity in UPF supply network is due to following reasons:

- Due to connections between different supply nets, supply ports, power switches and power aware cells. Also the connections can cross multiple hierarchies making it more cumbersome.

- Multiple supply ports can drive a single supply net and based on the resolution function of the net the final value is computed. UPF2.0/2.1 provides some predefined resolution rules (one_hot/parallel) for a supply net but UPF3.0 has relaxed it further to allow any custom resolution function.
  e.g
```
create_supply_net  N1 -resolve my_resolve_func
connect_supply_net N1 -ports { bot1/P1}
connect_supply_net N1 -ports { bot1/P2}
```

- There could be connection from UPF supply net to HDL single bit nets and hence they would require a vct either automatic one or explicit ones. These vcts induces an extra layer in supply flow and makes the debugging cumbersome.
- Although, UPF provides commands to connect supply nets to ports but it doesn't tell directly who is driver and who is receiver. This information is computed based on the direction and location of the nets and ports.

  For example, "`connect_supply_net N1 -ports { bot1/P1}`"
  Just by looking at this command one can not infer whether 'N1' is driving port 'P1' or 'P1' is driving net 'N1'.
- The ports/nets are neither in UPF nor in HDL but they are defined in liberty and UPF has commands to connect them.

Following section describes how user can debug the supply connections.

A. Static debugging of the connection

First of all we should ensure that all connections,vcts and resolution functions are proper and there are well identified root supplies.

To ensure these connections first of all we need to find out driver receiver relationship.This can be done by finding the the direction specified on supply ports. From there one can start populating the connection cones of driver and receiver.

```
e.g
create_supply_port P1 -direction out
create_supply_port P2 -direction in
...
connect_supply_net N1 -ports { bot1/P1}
connect_supply_net N1 -ports { mid1/P2}
```

So one can deduce that P1 being an output port is driving N1 and P2 being an input port is being driven by N1.

That is, P2 <= N1 <= P1.

One should also make a note of the root supply drivers which is P1 in this case.
Most of the EDA tools provide some kind of connection report to understand the complete supply network. Here is a snippet from Questa:

```
/chip_top/U11/vddx [LIB1] primary_power (MSPA_CONNECT_UPF_2_HDL_PWR) <= /chip_top/vdd_cx_int
[UPFSN1]. ./src/chip_top.main.upf(212)

This verbose report states that supply net " /chip_top/vdd_cx_int " is a driver net driving
the liberty pin "/chip_top/U11/vddx" which is a primary power of the cell and a default
automatic vct " MSPA_CONNECT_UPF_2_HDL_PWR" is applied.


Another kind of reports which is root supply reports are also generated by the EDA tools :
e.g
Root supply: /chip_top/u_hm_top_0/vdd_int [UPFSN10]
    Connected Supplies:
        /chip_top/u_hm_top_0/U3/vddx [LIB16]
        /chip_top/u_hm_top_0/u_hm_core/u_hm_sub/\doutbus2_A2_reg[0] /vddx [LIB17]
```

```
/chip_top/u_hm_top_0/u_hm_core/u_hm_sub/\doutbus2_A2_reg[1] /vddx [LIB18]
/chip_top/u_hm_top_0/u_hm_core/u_hm_sub/\doutbus2_A2_reg[2] /vddx [LIB19]
/chip_top/u_hm_top_0/u_hm_core/u_hm_sub/\doutbus2_A2_reg[3] /vddx [LIB20]
/chip_top/U11/vddx [LIB1]
/chip_top/vdd_cx_int [UPFSN1]
```

B. Dynamic debugging of supply connections:

This is the place where a good UI from EDA tools can help a lot. Here is a snippet from Questa which shows only supply network and connections with the values flowing across them.
[TBD]

VII. DESIGN POWER UP FAILURES

A common debug scenario that commonly occurs is when a low-power design fails to power up after a power down period. The problem can be because of any of the following reasons.

A. Missing/incorrect isolation/level shifter

Today's SoCs have large numbers of power domains where each power domain is interacting with other power domains. Two interacting power domains may also be operating with different voltage ranges. The Voltage representing logic value '1' in driving domain may represent it as '0' in receiving logic. Level-shifters are inserted at a domain boundary to resolve such issues; they translate the logic values to proper voltage ranges. The translation ensures the logic value sent by the driving logic in one domain is correctly received by the receiving logic in the other domain. Interacting domains may also cause floating problems. If the driving logic is powered down, the input to the receiving logic may float between 1 or 0. An un-driven input can also cause functional problems if it floats to an unintended logic value. To avoid this problem, isolation cells are inserted at the boundary of a power domain. In such scenarios, it becomes necessary to have isolation and level shifter cells at domain boundaries for the proper functioning of the design. Any missing cell can cause lot of functional issues, which would be difficult to debug. The following techniques are used to debug such scenarios.

1) Static Verification at compile time

Many tools statically determine the need for isolation and level shifter cells at domain boundaries from the PSTs and power states described in UPF. The PSTs describe the valid interacting states between two domains. These states clearly define the voltage ranges in which two domains are interacting and also whether one domain is relatively ON to other domain or not. If a signal is going from a domain of low voltage range to high voltage range then there is a need of Level-Shifter with the "`low_to_high`" rule. Similarly there is a need for the "`high_to_low`" rule for signals going from high-voltage to low-voltage. It is an error scenario if no level shifter strategy is specified or a level shifter strategy with a different rule is specified.

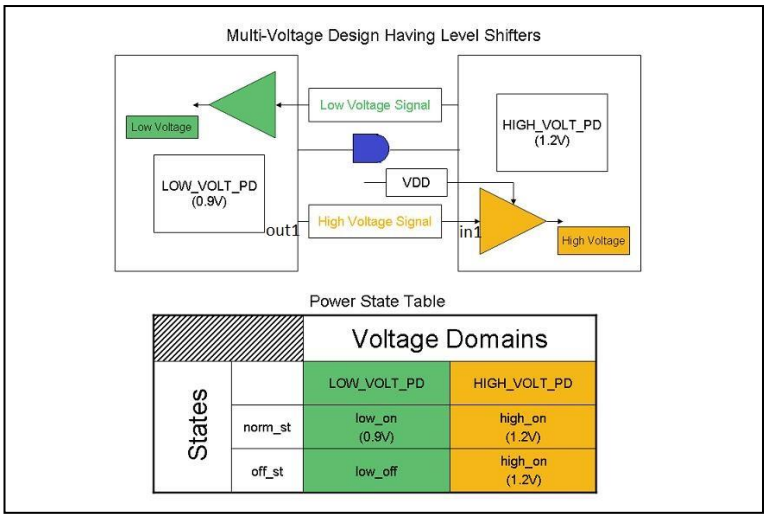

In the above example, two level shifter strategies are required, one with the rule "`low_to_high`" for the signal going from `LOW_VOLT_PD` to `HIGH_VOLT_PD` and other with the rule "`high_to_low`" for the signal going from

HIGH_VOLT_PD to LOW_VOLT_PD. Also HIGH_VOLT_PD is relatively ON to LOW_VOLT_PD, so there is a need of an isolation cell.

UPF Snippet :

```
add_pst_state norm_st -pst PST -state {low_on  high_on}
add_pst_state off_st -pst PST -state {low_off high_on}
```

The tool will issue the following type of message :

Missing Isolation Cell :

```
    Source power domain : LOW_VOLT_PD -> Sink power domain: HIGH_VOLT_PD
            1. Source port: out1 -> Sink port: in1
```

2) Tool generated assertions

System Verilog Assertions (SVAs) are a very powerful way to achieve the dynamic verification of low power designs. They can be used to validate power control logic sequences and also ensure that specific requirements are met before and after power mode transitions. Many of the EDA vendors provide automated, tool-generated assertions to check for missing or incorrect Isolation and Level shifter at run time. For each interface signal at the domain boundary an assertion is inserted that would check the need for Isolation and Level Shifter Cells.

The following types of errors are reported at run time:

```
   ** Error: UPF_MISSING_LS_CHK: Time: 90 ns, Missing level shifters for domain boundary,
LOW_VOLT_PD ( Operating Voltage: 5.000000 V ) => HIGH_VOLT_PD ( Operating Voltage: 5.200000 V
) for Source port : out1 -> Sink port: in1
    File: ./src/test.upf, Line: 32, Power Domain: LOW_VOLT_PD
```

3) UPF Bind Checker

UPF also provide a way to add custom assertions using the bind_checker command. The user can write their own assertions in a module and then bind that module to boundary instances to check whether a signal is isolated or not.

Checker module:

```
module checker_isolation(input op, src_supply, sink_supply) ;
    always@(src_supply)
        assert ( (src_supply == ON) || (sink_supply == OFF) || (op != 'X') )
else $error("Missing Isolation");
endmodule
```

The above checker module checks the value of a signal when Source Domain goes OFF and Sink Domain is ON. If a signal gets corrupted then it is case of Missing Isolation Cell.

B. Wrong Retention behavior

The incorrect retention also results in power-up failure. Sometimes the restore event does not happen and flop remains in corrupt state causing power-up failure. The easiest way to debug is to dump custom retention assertions using UPF bind_checker command and UPF generics. UPF generics provide the automated way to specify generic bind_checker statement for all the sequential elements of design. Users do not have to worry about clock/reset conditions for each sequential element. Tool will automatically do the identification of clock/reset conditions for each sequential element and apply the bind_checker statement accordingly.

Retention strategy:

```
set_retention RET
-domain PD
-save_signal {SAVE posedge}
-restore_signal {RESTORE  posedge}
-restore_condition {!UPF_GENERIC_CLOCK}
-elements {Q1 Q2}
```

Checker module:

```
module checker_retention #(parameter elem_name) (input ret_ff, restore_sig, clock) ;
    always@(posedge restore_sig)
        assert (clock != 1) else $error("Incorrect restore protocol for %s", elem_name);
endmodule
```

Bind_checker statement:

```
array set RET_STRTGY [query_retention RET -domain PD -detailed]
set ELEMENTS $RET_STRTGY(elements)
set RESTORE  $RET_STRTGY(restore_signal)

bind_checker ret_checker_inst
      -module ret_checker -elements $ELEMENTS \
      -parameters { {elem_name UPF_GENERIC_ELEM_NAME} } \
      -ports {{ret_ff   UPF_GENERIC_OUTPUT} \
              {restore_signal RESTORE} \
              {clock UPF_GENERIC_CLOCK}}
```

Above bind_checker will bind checker instance for each element by replacing UPF_GENERIC_CLOCK by corresponding clock of sequential element. Assertion will fail whenever retention protocol fails for any sequential element.

C. Macro Cell Corruption

The corruption of macro cells is governed by liberty attributes power_down_function and related_power_pin /related_ground_pin. These attributes use the primary supplies of macro cell to define the corruption semantics. So the primary supplies of macro cells should be properly connected for the correct functionality of liberty cells. One of the main reason for cell power-up failure is wrong supply connections and it is necessary to verify supply connections of each cell. Most of the EDA vendors dump the connection report for all supply connections. One must verify that supply ports are properly connected by checking the report.

Most common reasons for wrong supply connections are –

1) Unknown pg_type of primary supplies
   EDA tools normally do the automatic connections of macro cells by matching the pg_type properties from liberty file and if pg_type is not specified then no connection would happen.
2) Wrong inference of supply port direction
   In many cases, the direction of primary supply ports is defined as inout and tool needs to infer proper direction for automatic connections. But sometimes, tool is not able to infer the direction correctly and makes wrong connections.
3) No bias supplies defined in primary supply set
   related_bias_pin and power_down_function use bias supplies for corruption semantics. So, primary supplies of power domains should have nwell/pwell functions for proper bias connections.
4) No connections because tool is not able to infer Isolation/Level-shifter/Retention/Switch cells
   The cells corresponding to strategies are connected with strategy supplies and if tool is not able to infer the strategy and then supply ports would not be connected.
5) Incorrect backup supply connections in Isolation/Retention cells
   Isolation/Retention cells use backup_power/backup_ground for corruption semantics and they should be correctly connected according to strategy semantics.
6) No explicit connection to backup supplies in AON (always_on) cells
   AON cells use backup_power/backup_ground for corruption semantics and there are no UPF semantics to apply automatic connections on backup supplies. So, user must specify explicit connections for backup supplies for proper functionality of AON cells.

VIII. CONCLUSION

Debugging of low-power design still stands tall as one of the toughest challenges currently faced in the semiconductor industry. In this paper we have highlighted various challenges faced in debugging of a low-power design so that low-power issues can also be detected early thereby saving lot of design cycles. With the help of relevant examples, we have also demonstrated how such low-power issues can be either avoided or fixed thereby significantly increasing the productivity of the debug process.

IX. REFERENCES

[1] IEEE Std 1801™-2013 for Design and Verification of Low Power Integrated Circuits. IEEE Computer Society, 29 May 2013
[2] Power Aware Models: Overcoming barriers in Power Aware Simulation (Mohit Jain, Amit Singh, J.S.S.S.Bharath, Amit Srivastava, Bharti Jain), DVCon-Europe 2014.
[3] Static and Formal Verification of Power Aware Designs at the RTL Using UPF (Rudra Mukherjee, Amit Srivastava and Stephen Bailey), DVCon 2008.
[4] IEEE Std 1801™-2015, IEEE Standard for Design and Verification of Low-Power, Energy-Aware Electronic Systems.