

LVS Verification across Multiple Power Domains for a Quad-Core Microprocessor

WEI LI, DANIEL BLAKELY, SCOTT VAN SOOY, KEVEN DUNN, DAVID KIDD,
and ROBERT ROGENMOSER

Broadcom Corp.

and

DIAN ZHOU

University of Texas at Dallas

A unique LVS (layout-versus-schematic) methodology has been developed for the verification of a four-core microprocessor with multiple power domains using a triple-well 90-nm CMOS technology. The chip is migrated from its previous generation that is for a twin-well process. Due to the design reuse, VDD and GND are designed as global nets but they are not globally connected across the entire chip. The standard LVS flow is unable to handle the additional design complexity and there seems to be no published literature tackling the problem. This paper presents a two-phase LVS methodology: a standard LVS phase where power and ground nets are defined as global nets and a multi-power-domain LVS phase where power and ground nets are treated as local nets. The first phase involves verifying LVS at the block level as well as the full-chip level. The second phase aims at verifying the integrity of the multi-power-domain power grid that is not covered in the first phase LVS. The proposed LVS methodology was successfully verified by real silicon.

Categories and Subject Descriptors: B.7.1 [**Integrated Circuits**]: Types and Design Styles—*Microprocessors and microcomputers*; *VLSI (Very large scale integration)*; B.7.2 [**Integrated Circuits**]: Design Aids—*Layout*; *verification*; J.6 [**Computer-Aided Engineering**]: *Computer-aided design (CAD)*

General Terms: Design, Verification

Additional Key Words and Phrases: LVS, physical verification, multi-core microprocessor

1. INTRODUCTION

As semiconductor technology advances into 90-nanometer and beyond, microprocessor design has evolved into today's very competitive battle to produce complex multicore architectures. Multiple processor cores on a single die can

Authors' addresses: W. Li, D. Blakely, S. Van Sooy, K. Dunn, D. Kidd, and R. Rogenmoser, Broadcom Corporation, 2451 Mission College Boulevard, Santa Clara, CA 95054; D. Zhou, Department of Electrical Engineering, The University of Texas at Dallas, Richardson, TX 75083; email: Weili@broadcom.com.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or direct commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 1515 Broadway, New York, NY 10036 USA, fax: +1 (212) 869-0481, or permissions@acm.org.

© 2006 ACM 1084-4309/06/0400-0490 \$5.00

ACM Transactions on Design Automation of Electronic Systems, Vol. 11, No. 2, April 2006, Pages 490–500.




Fig. 1. Chip detailed floorplan.

operate in various ways to yield scalable performance in a complexity- and power-efficient manner [Brooks et al. 2000]. The combination of continually rising complexity of chip design and tightening time-to-market schedule has forced designers to seek new methods to accelerate the development process. The use of core-based designs, the reuse of existing IP, and efficient verification methodologies have all helped to expedite the processor design cycle.

A system on a chip (SOC) BCM1280 has been designed with four 64-bit MIPSTM CPU cores, operating at Giga-Hertz frequency, as well as an on-die 1MB L2 cache, and dedicated networking hardware [Sano 2002; Moll 2004]. Figure 1 shows a detailed floorplan of the chip. BCM1280 is implemented with 164 million transistors on a $13.9 \times 14.1 \text{ mm}^2$ die. It is fabricated with a 90-nanometer triple-well bulk CMOS process with seven routing metal layers and a redistribution metal layer dedicated to power and ground nets from the package to the chip. The CPU core is migrated from a previous twin-well process design, in which the power grid (VDD/GND) was designed and verified as global nets across the entire chip. In this quad-core microprocessor, four CPU cores are identical and each sits in a separate triple-well region. Each single core, isolated from the others by a triple-well, forms a separate processor region. Therefore, under the full-chip context, each core is now built in a unique voltage region. In addition to the complexity of the multiple ground domains, the chip is divided into two logic voltage regions (1.1V). The previous generation chip, from which BCM1280 is migrated, was designed with one power supply only.

Due to the manual nature of the layout construction in a full-custom design, physical verification is very critical to the success of the chip design. One of the most important key steps to perform physical verification on the layout design is to verify that the layout matches the schematic. This is so called LVS (layout-versus-schematic), which involves checking the consistency between the original schematic netlist and the one extracted from the layout. Analyzing

Table I. Process Characteristics

Process	Triple-Well Bulk CMOS P+ Substrate
Feature size	90 nm
Clock frequency	1 GHz
Metal Layers	7 copper layer 1 redistribution layer
Supply Voltages	1.1V, 2.5V and 3.3V
Typical Power	25W
Package	Flipchip

the multi-power-domain power grid in LVS was a great challenge to us during the design of BCM1280. Because of the design reuse and the implementation of the multiple power/ground domains, VDD and GND are designed as global nets, but they are not globally connected throughout the chip. The standard LVS flow, which verifies VDD and GND as global nets, is no longer functional in verifying power grid at the full-chip level for this four-core microprocessor. In an extensive literature search, we found no published work tackling the considered problem. Hence, in order to reuse the existing design and make the fewest changes to the reused design, we need to come up with a special LVS methodology. In this article, we present an LVS methodology for a four-core microprocessor designed with multiple power domains, given the fact that the microprocessor core is previously designed for a chip with a single ground domain.

The rest of this article is organized as follows: The process technology in manufacturing this chip is briefly introduced in Section 2. The structure of the power grid is described in Section 3. The implementation of multiple power and ground domains is discussed in Section 4. The LVS methodology is presented in detail in Section 5. We conclude the presented work in the last section with some final remarks.

2. PROCESS TECHNOLOGY

The BCM1280 is fabricated in a 90-nanometer triple-well CMOS technology on a P-substrate. The major process characteristics are shown in Table I.

Triple-well technology allows P-wells to be placed inside a deep N-well. This type of structure electrically isolates the wells from the substrate, providing a means to isolate the n-type devices that would normally exist in the P-substrate as in a twin-well design. When a triple-well CMOS technology is used, separate voltage lines can be used to connect to the NMOS and PMOS bulk regions. Then, these lines can be modulated to provide a back-bias voltage for regions that are not operating. Large blocks of a chip can be disabled when they are not needed in order to save energy by reducing leakage current. A cross-section picture of the triple-well structure is shown in Figure 2.

3. POWER DISTRIBUTION

Power distribution systems (power grid) are designed to provide required voltages and currents to the transistors that perform the logic functions of a chip. With every technology shrink, power distribution networks do not scale due to the fact that the current demand per unit area of a design is increased. As a result, more silicon area needs to be dedicated to distributing higher currents




Fig. 2. Triple well structure provides an isolating N layer between the local P-well and the P-substrate.

at lower voltages [Kam et al. 2000]. Hence, power consumption has become a more and more critical issue in microprocessor design. An efficient power distribution network on large microprocessor is one of the key factors in achieving overall high performance targets of the chip.

On the BCM1280, Metal-6 and metal-7 are extensively used in distributing power and ground across the chip. Ground (GND) and power (VDD_{chip}, VDD_{switch}) lines are routed alternately on Metal-6 in the vertical direction and Metal-7 in the horizontal direction. This forms the power grid structure at the chip top level. The power and ground lines in the local sections are grided on Metal-2 through Metal-5, on which the local power grids are formed at the block level. The local power grids are left floating until final assembly at the full-chip level. At the final stage of the layout design, an in-house automated procedure is applied, connecting local logic to the top-level power grid, to implement a complete full-chip power distribution network.

Due to the trend of ever increasing die size in microprocessor design, efficient power delivery from the supply source to the chip has become very important. Flipchip technology is chosen for packaging this chip. Power and ground are supplied to the package through two VDD inputs (VDD_{switch} and VDD_{chip}) and one GND input, respectively. Pads arrayed across the active face of the die are used to connect to I/O, power and ground. Solder bumps are attached to pads. A redistribution metal layer is designated for distributing power and ground from the bump to the chip level power grid (Figure 3). With flipchip technology, power and ground bumps are connected directly to the on-chip power and ground distribution system; and these connections can be made across the entire die so power/ground is delivered wherever needed. This significantly reduces overall chip IR drop by eliminating the IR drop associated with delivering power and ground from peripheral wirebond pads.

4. MULTIPLE POWER/GROUND DOMAINS IMPLEMENTATION

BCM1280 is designed with two power supply inputs and one ground input at the package level. As a result of function and logic requirements, there are multiple voltage regions across the chip. This design style enables power-saving capability with a selectable configuration, allowing the chip to be an optimal fit into a broad variety of applications. For instance, for some applications where power consumption is a major concern, two of the four CPU cores and the correlated peripheral logic circuitry could be shut-off to achieve power efficiency.




Fig. 3. Power and ground are distributed from the bumps onto the chip-level power grid through a redistribution metal layer.




Fig. 4. An SB1 core inside triple-well with power switch.

4.1 Multiple Ground Domains Implementation

To take advantage of the triple-well technology, we have designed the chip with each microprocessor CPU core placed inside a triple well—a deep N-well, inside which the P-well is set apart from the P-substrate. Since the P-well connected to the core ground is isolated from the ground outside the core by a reverse biased deep N-well, the individual cores are electrically isolated from one another as well as the rest of the chip. As a result, BCM1280 has a total of four core grounds connected to the P-wells inside the four triple wells, respectively, and one chip ground connected to the P-substrate.

Figure 4 gives a top-down layout view of a processor core in the triple-well structure. This particular process offers a migration path for existing twin-well designs to triple-well structure by adding an NW-ring and a triple-well recognition layer (TWR) around the core. All N-wells inside the triple-well are shorted; therefore, they are all on the same potential.

A power switch structure is designed surrounding each core. As illustrated in Figure 4, the transistor's source is connected to the core ground inside the triple-well and the transistor's drain is tied to the chip ground inside the substrate. The power switch enables a single core to be turned on or off individually. This independently configurable feature is very desirable for a variety of applications

that may have different power consumption requirements, as power dissipation limits have emerged as a major constraint in the design and application of microprocessors [Brooks et al. 2000; Kam et al. 2000; Iyer et al. 2002; Karnik et al. 2002].

4.2 Multiple Power Supplies Implementation


BCM1280 has several power supplies that all share the same chip ground reference. All require adequate external decoupling capacitors with low-inductance paths. The ground reference GND is common for all the power supplies. As we described in the previous section, there are two power supply inputs VDD_chip and VDD_switch at the package level in this design. VDD_chip provides power for all the internal logic except the node controller, packet manager, switch and HT/SPI-4 ports. All four CPU cores operate at the same supply voltage, VDD_chip, but use a different internal ground net for each core. The different internal ground net for each core is separated from chip ground by an NMOS switch. VDD_switch provides power for all the internal logic of the node controller (NC), packet manager (PM), switch and HT/SPI-4 ports.

5. LVS METHODOLOGY


LVS is one of the most important physical verification steps in integrated circuit (IC) design. It serves to check the equivalency of a layout design against the corresponding schematic design. If the two designs are equivalent under the given LVS rule deck (LVS runset), then the LVS result is said to be “LVS clean”. LVS tools first generate a netlist based on extracting the connectivity of an input layout stream file. The extracted netlist is then compared to the schematic netlist to ensure the two views match.

Typically, power nets are globally connected throughout the design, and so are ground nets. VDD and GND are designed as global nets at the block level, which is done to simplify the schematic level design process; treating them as local nets would require creating specific hierarchical ports for VDD and GND within each schematic. At the full-chip level, however, under the multiple power/ground domains scenario we described in the previous sections, VDD and GND can no longer be treated as global nets. LVS methodology for the block level now fails at the full-chip level. For example, a short on GND between the different processor cores cannot be detected because GND is considered to be global. The design team needed an LVS methodology that enabled each block to be independently verified during the development. They also needed the LVS flow to be capable of verifying the integrity of the chip level power grid with minimum impact in altering the design.

In this section, we discuss a two-stage LVS flow used in the BCM1280 full-chip LVS verification. It consists of a standard LVS stage (standard-LVS) where VDD and GND are defined as global nets; and a multi-power-domain LVS stage (MPD-LVS) where VDD and GND are verified as local nets. Standard-LVS would suffice for verification at the block level. HERCULES [HERCULES 2003], one of the well-known commercial LVS tool, is used for the LVS verification.



(a) VDD is global net. LVS reports text shorts.



(b) VDD is local net. VDD needs to be texted at every level of the LVS equivalence point to get LVS clean.

Fig. 5. Global texts vs. local texts in hierarchical LVS processing.

5.1 Global Nets versus Local Nets in LVS

In hierarchical LVS, layout text is viewed hierarchically. In a schematic netlist, nets listed as global are assumed to be connected throughout the design during netlist comparison. Similarly, nets defined as global in layout are treated as global during netlist extraction. Global net names are dispersed through connections up the hierarchy. Local or nonglobal nets only specify a net's value for a particular hierarchical level. If two or more texts are attached to one net, LVS reports this as a *short*. A text *open* is reported, if different physical nets have the same text.

Figure 5 gives a layout example to illustrate how the LVS tool handles global texts and local texts. In Figure 5, cell TOP is the top block of the design, A is instantiated in TOP and B is placed inside A. In case 5(a), since VDD is declared as a global net, it is propagated from cell B all the way up through the hierarchy to the top level. This results in text shorts between VDD and VDD_switch, and between VDD and VDD_chip. There are two strategies, which can be utilized to solve the text conflicts in this particular case. One possible solution is that VDD needs to be treated as a local net, which would require VDD to be labeled at every LVS compare equivalence level in both the layout and the schematic. This is not desirable since it involves significant and time-consuming effort in redesigning the layout and the schematic. Another method is to use two LVS runs. In the first LVS run, VDD is kept as a global net and VDD_switch/VDD_chip labels are ignored. This will cause opens at the top level, which can be ignored at this stage, but the design can be verified without layout or schematic modifications. In the second LVS run, VDD is treated as a local net




Fig. 6. Chip-level LVS verification flow chart.

and VDD_switch/VDD_chip are treated as global nets. This ensures the opens from the first LVS run are not real and the transistors are connected to the correct power domain.

5.2 Detailed LVS Flow

As Figure 6 shows, our LVS flow consists of two phases: a standard-LVS phase, to check the consistency between the layout and the schematic under a unitary power domain, and a multipower-domain LVS phase, to verify the top-level power grid is open/short free.

The original schematic that standard-LVS reads is a *SPICE* format netlist, where VDD and GND are specified as global names. They are assumed to be globally connected throughout the schematic netlist; therefore, they are not declared as ports in *subckt* definitions across all the netlist hierarchies. The layout stream file input to the LVS tool is in *GDSII* format. Global texts automatically propagate up and down the layout hierarchy. VDD and GND nets are designed




Fig. 7. An example shows that physical opens and shorts are not found by standard-LVS.

as global in the layout, so they are not labeled at every hierarchical level in the layout. The special power and ground names (such as VDD_chip, VDD_switch, GND_core0, etc.) used in the multi-power-domain design are labeled on the bump at the chip top level.

The LVS runset file is an ASCII file, which contains instructions for the LVS tool to get the input files and to perform the LVS check against a set of rules defined in the runset. In the LVS runset_1, power and ground names of the multi-power-domain are replaced with VDD and GND, respectively, and VDD and GND are declared as global names for both schematic and layout. During the standard-LVS run, since different power names are treated as VDD and different ground names are treated as GND, we expect to see text opens for VDD and also text opens for GND in the LVS error report. The standard-LVS is considered to be “clean” with VDD and GND text open errors only. However, we are not guaranteed that the text opens reported by LVS are the opens we created intentionally by replacing the texts. These could be real physical opens unintentionally added. It is also possible for the standard-LVS flow to hide real shorts, as illustrated in Figure 7. In the picture, A1, A2, A, B and C are placed inside TOP. All the top-level texts are replaced with GND. There are a real open for GND_A in Figure 7(a) and a real short for GND_A and GND_B in Figure 7(b). However, these real errors are masked by the text replacement.

Multi-power-domain LVS is the second phase of the LVS flow, which targets verifying the integrity of the chip-level power grid. Unlike the standard-LVS,

Table II. A Comparison of the Two LVS Runsets

Standard LVS Runset (runset 1)	Multi-power-domain LVS Runset (runset 2)
• VDD and GND are global nets	• VDD and GND are local nets
• Multi-power-domain power and ground names labeled at the bumps are replaced to VDD or GND	• Original multi-power-domain power and ground names labeled at the bumps are declared as global

in which VDD and GND are treated as global names, multi-power-domain LVS analyzes power and ground nets as local nets. A netlist parser is implemented to create a schematic netlist with localized VDD and GND nets, and also the multi-power-domain power and ground names are added as ports at the appropriate level of netlist hierarchy based on the information provided in a configuration file. As listed in Table II, where the major differences in the two LVS runsets are highlighted, the global net definition of VDD and GND are removed in LVS runset.2; VDD and GND nets are now local just like other signal nets. The actual multi-power-domain power and ground names labeled at the bumps are used in the second phase LVS. Hence, any opens or shorts on the power-grid will be reported upon completing the multi-power-domain LVS verification.

6. SUMMARY

Due to the importance of the stringent time-to-market demand, design reuse has become a common practice in today's IC industry. Part of the BCM1280 chip, including the microprocessor cores, is migrated from the previous design. Under this circumstance, our LVS strategy was designed to use two phases. A standard LVS phase is for a complete LVS verification at the block-level and a partial LVS verification at the chip top-level, while a multiple-power-domain LVS phase serves to verify the integrity of the power grid that is not covered by the standard-LVS phase.

The proposed two-stage LVS methodology allows us to fully verify the entire chip with minimum changes to reused design data. The price we pay is to have two LVS runs before the tapeout and to manipulate the full-chip netlist to localize power and ground nets for the second LVS run. Theoretically, it is sufficient at the full-chip level to run only the multi-power-domain LVS to completely verify the chip, however, we find it useful in practice to run both LVS runs. When VDD and GND nets are treated as local nets in the multi-power-domain LVS verification, while they are physically designed as global nets, it would be very difficult for LVS debug. The standard-LVS run at the chip level is useful to debug LVS problems before the chip is fully LVS clean. In the future projects, it is likely that multiple-power-domain application will become our standard design flow. If this trend is to become a reality, we will need to design VDD and GND as local nets by texting them in the layout at every LVS equivalency hierarchical level, and design the schematics accordingly using localized VDD and GND nets.

The LVS methodology we present in this article was successfully employed on implementing BCM1280, a high-end embedded four-core multiprocessor designed with multiple power domains. The proposed LVS flow is efficient in delivering the verification results before the tapeout and it is verified by real

silicon. In laboratory tests, we are able to successfully shut off power to any of the four CPU cores. This enables the BCM1280 to be used in many different applications.

ACKNOWLEDGMENTS

The authors acknowledge the contributions of the whole BCM1280 design team at Broadcom Broadband Microprocessor group. We feel privileged to present this work. Special thanks to Susan Kwan for her efforts in implementing the layout and debugging the LVS.

REFERENCES

- BROOKS, D. M., BOSE, P., SCHUSTER S. E., JACOBSON, H., KUDVA, P. N., BUYUKTOSUNOGLU, A., WELLMAN, J. D., ZYUBAN, V., GUPTA, M., AND COOK, P. W. 2000. Power-aware microarchitecture: Design and modeling challenges for next-generation microprocessors. *IEEE Micro* 20, 6, 22–44.
- HERCULES User's Reference Manual* 2003. Synopsys Inc.
- IYER, A. AND MARCULESCU, D. 2002. Microarchitecture-level power management. *IEEE Trans. VLSI Syst.* 10, 3, 230–239.
- KAM, T., RAWAT, S., KIRKPATRICK, D., ROY, R., SPIRAKIS, G. S., SHERWANI, N., AND PETERSON, C. 2000. EDA challenges facing future microprocessor design. *IEEE Trans. Comput. Aid. Des. Integ. Circ. Syst.* 19, 12, 1498–1506.
- KARNIK, T., BORKAR, S., AND DE, V. 2002. Sub-90nm technologies—Challenges and opportunities for CAD. In *Proceedings of IEEE International Conference on Computer Aided Design* (San Jose, CA, Nov.). IEEE Computer Society Press, Los Alamitos, CA, 203–206.
- MOLL, L. 2004. The next-generation line of Broadcom's SiByte™ multiprocessor SoCs. In *Proceedings of the Fall Microprocessor Forum* (San Jose, CA, Oct.).
- SANO, B. 2002. Chip combines four 1GHZ cores. In *Microprocessor Report*.

Received September 2004; revised May 2005; accepted December 2005