A Design Platform for 90-nm Leakage Reduction Techniques

Philippe Royannez, Hugh Mair, Franck Dahan, Mike Wagner, Mark Streeter, Laurent Bouetel, Joel Blasquez, H. Clasen, G. Semino, Julie Dong, D. Scott, B. Pitts, Claudine Raibaut, Uming Ko
Texas Instruments
821 Avenue Jack
KILBY - BP 5 - 06271 Villeneuve-Loubet Cedex
+33 6 86 36 88 47
p-royannez@ti.com

ABSTRACT
Methodology, EDA Flow, scripts, and documentation plays a tremendous role in the deployment and standardization of advanced design techniques. In this paper we focus not only on leakage reduction techniques but also on their deployment as a worldwide infrastructure as the added-value resides not only in the techniques themselves but also in the way they are implemented to build an efficient, re-usable, robust, low cost and portable platform. Techniques have been silicon proven on the 90-nm TI CMOS technology and is commonly used to design SoC with complexities over 100 Million transistors

Categories and Subject Descriptors
B.7.2 [Design Aids]: Integrated Circuits, Design Aids

General Terms
Design,Verification

Keywords
SoC Design, Leakage Power Management, Wireless Application processor

1. LEAKAGE REDUCTION TECHNIQUES
The typical set of dynamic and leakage power reduction techniques we propose includes 1) idle (clock stopped), 2) standby (logic and memory retention with fast re-start using header and footer diodes), and 3) deep sleep (power off mode for ultra low power consumption using both voltage scaling and power gating) [1].

2. DESIGN PLATFORM
The design platform that enables the integration of these techniques is the alliance of library cells and memory compilers, of EDA scripts and integration guidelines. Tools and Flow plays a tremendous role especially for verification. Spyglass™ rule checker at both RTL and gate level is efficient at catching errors such as "forgotten isolation cells". Other tools have been developed to verify that the current surge that happens during the power restore sequence does not cause any issues.

The complexity of SoC level leakage management can be handled by a divide-and-conquer approach. The SoC is split into several power managed IPs. We defined an IP as PM (Power Managed) if it can be put in retention or in off mode using, voltage scaling, power gating and Diode footed SRAM. We qualify an IP as being part at the PM framework if it is compliant with a set of rules for regularity (homogenous metal grid, embedded power switches on a regular pitch …) and present a standardized set of signals acting as a power control interface (similar to the 5 pins JTAG standard interface). Once a complete library of PM IPs (including CPU, DSP core, memories …) is available, SoC integrators can build up their design with the usual IP-based approach. From the physical point of view, each component can be placed according to regular metal grids. From the logic point of view, each power control interface will be plugged to a global "power controller". This dedicated block centralizes the power management and its scalable architecture allows it to open as many slots as required to plug all the PM IPs. Each slot contains small sequencers to generate autonomously the standard power down and wake up control sequences in slave mode. All these sequencers are supervised and synchronized by the PWC (power controller). The PWC is always on and is controlled at software level where all the intelligent power management is actually happening.

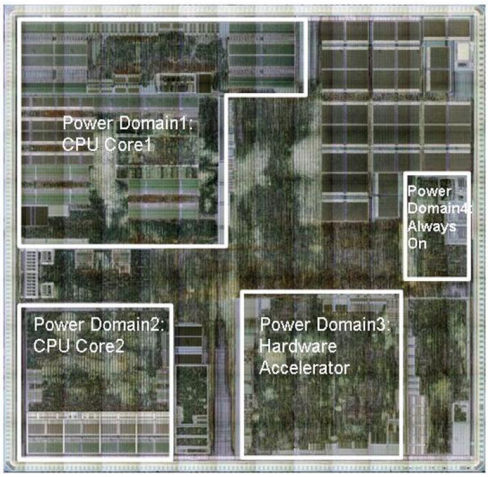

Figure 1: Die Microphotograph and Power Domains

DAC 2005, June 13–17, 2005, Anaheim, California, USA.
ACM 1-59593-058-2/05/0006.

3. APPLICATION EXAMPLE

As an application on a product, all of the techniques presented have been integrated on the most recent generation of OMAP™ wireless application processors. The chip has been fabricated on a 90nm TI CMOS technology. It is based on dual-core architecture ARM11+ TI DSP together with both 3D graphic and image and video accelerators. The logic portion of the chip is divided into five independently managed power domains as follows: one domain for each CPU core, one for the graphics hardware accelerator, one for a domain which is always active, and one for the rest of the chip which includes signal distribution and various peripherals (Figure 1).

It implements both a retention mode (fast wake up) and an off mode for ultra low leakage. Silicon measurements are showing an off-state leakage of 30uA standby current, which is a 40x leakage reduction versus active leakage at room temperature.

4. REFERENCES

[1] P. Royannez, H. Mair, F. Dahan, M. Wagner et. al.; "90nm Low Leakage SoC Design Techniques for Wireless Applications"; ISSCC'05, Feb 2005