28th International Conference on VLSI Design

VLSI DESIGN 2015

14th International Conference on Embedded Systems

FirmLeak: A framework for efficient and accurate runtime estimation of leakage power by firmware

Arun Joseph, Anand Haridass, Charles Lefurgy[+], Spandana Rachamalla, Sreekanth Pai, Diyanesh Chinnakkonda, Vidushi Goyal[*]

IBM Systems & Technology Group, IBM Research[+], IIT Kharagpur[*]

Abstract

- Separating the dynamic and leakage power can enable new optimizations for cloud computing.
- We introduce FirmLeak, a new framework that enables accurate, real-time estimation of microprocessor leakage power by system software.
- FirmLeak accounts for power-gating regions, per-core voltage domains, and manufacturing variations.

Introduction

- Leakage power estimation in modern high-performance microprocessors, must account for significant manufacturing variation [2] as well as workload-induced temperature variation.
- Since firmware power management may operate on a longer time intervals than power-gating of cores or function units, the ability to derate the leakage estimate to account for the actual power-on time becomes important.
- Accurate runtime estimation of leakage power from [1, 3] requires intensive post-silicon power characterization and data collection.

FirmLeak

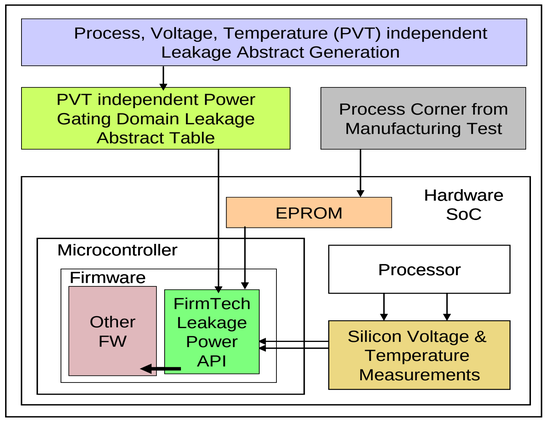

Fig. 1. FirmLeak Overview

- Introduces the use of Process, Voltage and Temperature (PVT) independent pre-silicon Power Gating domain Leakage Abstracts (PGLA) in firmware to accurately estimate per-device type contributions and total runtime leakage power.
- Enables significant reduction in post-silicon power characterization.
- Enables power-gating aware estimation by adding accompanying circuitry to track the time spent in power-gating states.

Motivating Scenarios

- Cloud-based billing

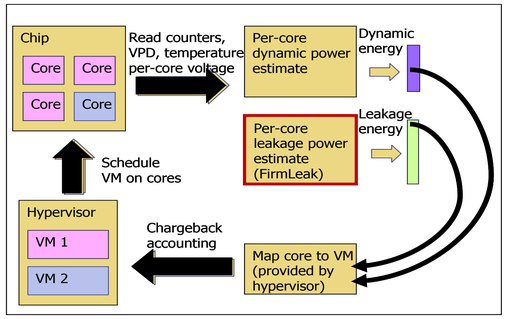

Fig. 2. Per core energy for VM energy chargeback

- Fan-based power optimization

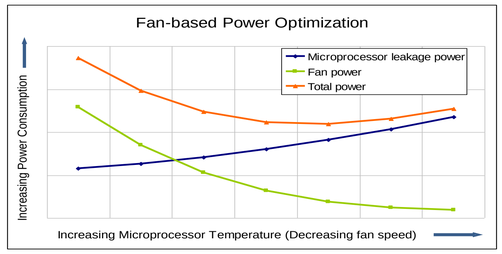

Fig. 3. Optimizing the sum of fan power and leakage power

PVT Independent Leakage abstracts for Power Gating Domains

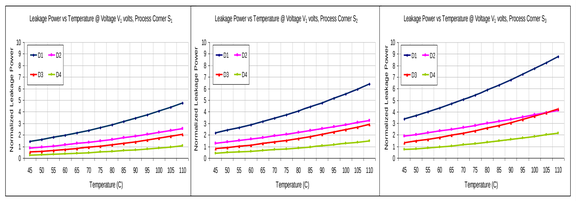

Fig. 4. Motivation for using per device type power gating abstracts in the runtime estimation of leakage power in firmware.

Power-Gating Domain	Rail	Device type					
		svt n	svt p	hvt n	hvt p	mvt n	mvt p
1	Vdd	W_1,l_1,C_1	W_2,l_2,C_2	W_3,l_3,C_3	W_4,l_4,C_4	W_5,l_5,C_5	W_6,l_6,C_6
1	Vcs	W_7,l_7,C_7	W_8,l_8,C_8	W_9,l_9,C_9	W_{10},l_{10},C_{10}	W_{11},l_{11},C_{11}	W_{12},l_{12},C_{12}
2	Vdd	W_{13},l_{13},C_{13}	W_{14},l_{14},C_{14}	W_{15},l_{15},C_{15}	W_{16},l_{16},C_{16}	W_{17},l_{17},C_{17}	W_{18},l_{18},C_{18}
2	Vcs	W_{19},l_{19},C_{19}	W_{20},l_{20},C_{20}	W_{21},l_{21},C_{21}	W_{22},l_{22},C_{22}	W_{23},l_{23},C_{23}	W_{24},l_{24},C_{24}

Fig. 5. Power gating domain leakage abstract (PGLA) table

- For each design unit that makes up the hardware block (example: microprocessor), we generate a leakage abstracts, using techniques described in [4].

Run-time Leakage Power Estimation

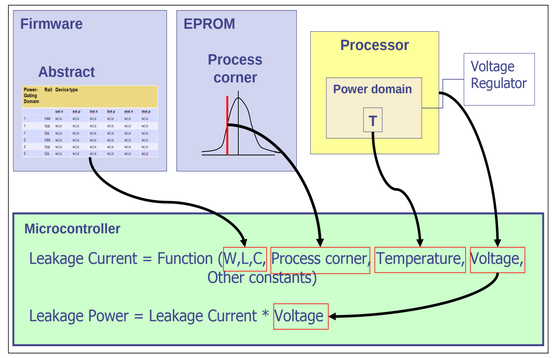

Fig. 6. Leakage power estimation via firmware

- The information for each device type in the domain, along with process corner information from test, temperature and voltage conditions from hardware sensors, technology specific constants configured in firmware, is passed into a device level leakage power computation model.

Derated Leakage Power Estimation

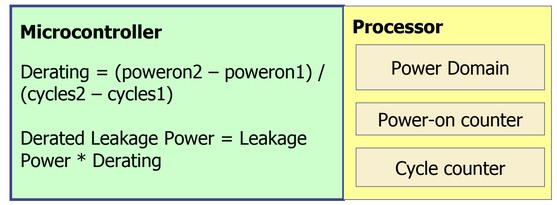

Fig. 7. Derated leakage power estimation

- The derating factor is the ratio of the incremented count in the power-on counter to the incremented count of the cycle counter during the interval, and will be a value in the range of 0% to 100%.

Experimental Setup

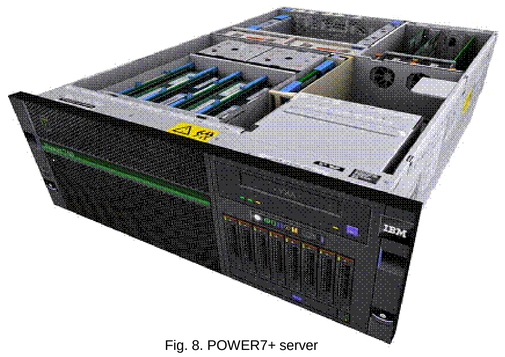

Fig. 8. POWER7+ server

- Power 730/740 class server [5] used for experimental evaluation.
- Uses 32nm POWER7+ processors.
- 2 socket entry-level SMP server with up to 16 processor cores.
- The system planar supports two POWER 7+ modules.

Experimental Results

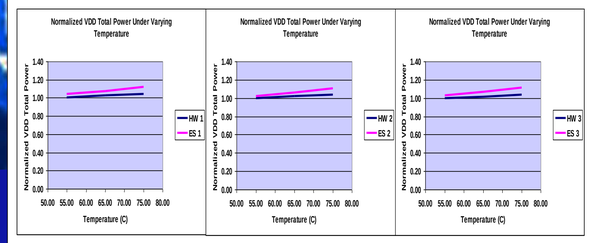

Fig. 9. VDD total power under varying conditions of temperature

- Average Error of ~5%.

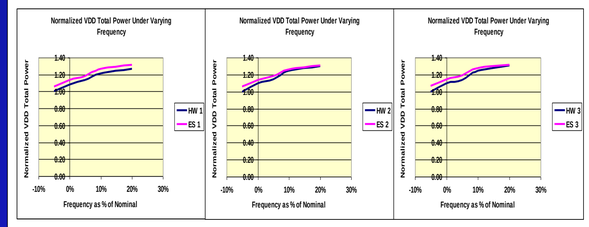

Fig. 10. VDD total power under varying conditions of frequency

- Average error of ~3.8 %.

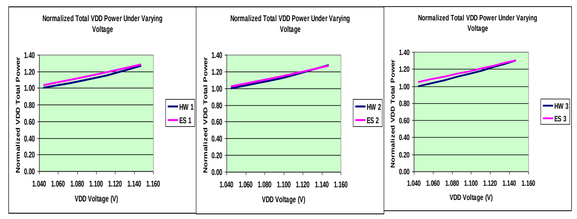

Fig. 11. VDD total power under varying conditions of voltage

- Average error of ~2 %.

References

1. Wei Huang et al. "TAPO: Thermal-aware power optimization techniques for servers and data centers," Green Computing Conference and Workshops (IGCC), 2011.
2. S. Borkar et al. "Parameter variations and impact on circuits and microarchitecture," in Proc. Design Autom. Conf., Jun. 2003.
3. US Patent 8,527,794: "Realtime power management of integrated circuits," Ibrahim, A. and Dwarakanath, A. and Shimizu, D.P .
4. Dhanwada, N. et al. "Efficient PVT independent abstraction of large IP blocks for hierarchical power analysis," Computer-Aided Design (ICCAD), 2013 IEEE/ACM International Conference on, pp.458,465, 18-21 Nov. 2013.
5. IBM Redbooks: "IBM Power 710 and 730 Technical Overview and Introduction," May 2013.