GREENER MOBILE SYSTEMS BY CROSS LAYER INTEGRATED ENERGY MANAGEMENT (GEMSCLAIM)

Goal of the project

The GEMSCLAIM project aims at introducing novel approaches for reducing the "greed for energy" of modern battery powered systems, thereby improving the user experience and enabling new opportunities for mobile computing

Mobile terminals and consumer devices are among the fastest growing markets in computing. In the long term, further growth is endangered by the "power/ energy wall". The purpose of GEMSCLAIM is to explore new techniques in energy optimization via an interdisciplinary vertical approach: a novel combined optimization across the major HW/SW system layers (compiler/OS/HW platform).

Short description of the project

- The ever-growing need for energy efficient computation requires adequate support for energy-aware thread scheduling that offers insight into a systems behavior for improved application energy/ performance optimizations. Runtime accurate monitoring of energy consumed by every component of a multi-core embedded system is an important feature to be considered for future designs. Although, important steps have been made in this direction, the problem of distributing energy consumption among threads executed on different cores for shared components remains an ongoing struggle.
- We aim at designing a generic low-cost and energy efficient hardware infrastructure which supports thread level energy consumption monitoring of hardware components in a multi-core system.
- The proposed infrastructure provides upper layers (operating system and application threads) with per thread and per component energy accounting API (Application Programming Interface), similar with performance profiling functions. Implementation results indicate that the proposed LEM (Load and Energy Monitor) adds around 10% resource overhead to the monitored system. Regarding the power estimates, the one derived by LEM achieve a correlation degree of more than 95% with the ones obtained from physical power measurements.

Project implemented by

Mobile Computing, Sensors Network and Embedded Systems Research Laboratory, Computer and Software Engineering Department, Faculty of Automation an Computers

Implementation period

Sep. 2012-Aug.2015

Main activities

In a world of de-facto standards as well as huge amounts of legacy HW and SW, it is very difficult to achieve real breakthrough in system-wide energy savings beyond fragmented point solutions, e.g. at the HW or OS level.

GEMSCLAIM´s mission is to overcome this hurdle by a novel cross layer energy optimization approach that combines the following major research activities:

- Development of an energy-aware optimizing and parallelizing compiler;
- Component aware energy-efficient operating system and
- Customizable HW modelling with energy monitoring facilities.

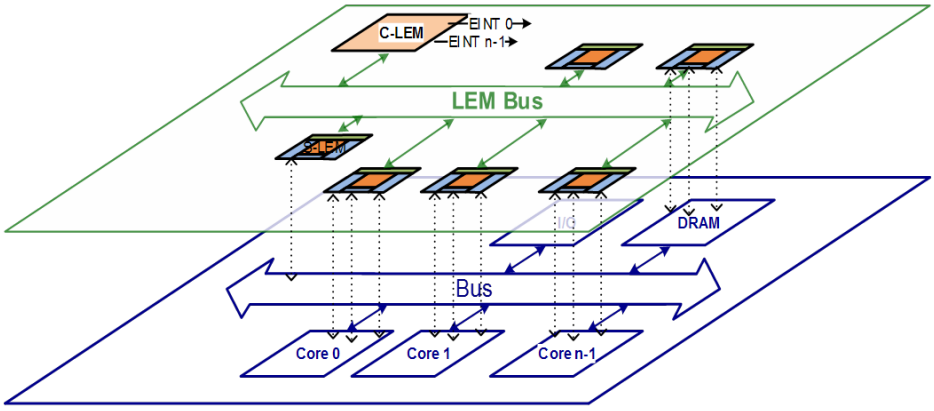

Overall solution architecture

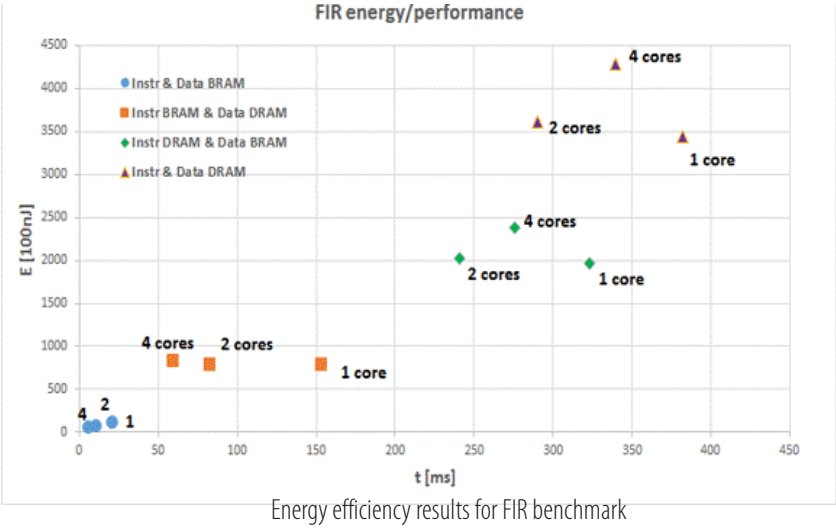

Energy efficiency results for FIR benchmark

Results

The contributions of this work are as follows: (1) hardware infrastructure for dynamic energy consumption monitoring in a heterogeneous multi-core system with per-thread energy accounting; (2) energy interrupt specification and design; (3) a use case on the software side (OS and drivers) for run-time per-thread energy accounting implementation on FPGA; and (4) validation of proposed infrastructure on a high-end FPGA board with physical energy measurements.

Per-thread energy accounting (PTEA) can be achieved by splitting the whole energy into processing energy (energy consumed by processing cores), data movement energy (energy consumed by interconnects to read and store data) and data storage energy (energy consumed by memories to store task data). The proposed infrastructure addresses all of these energy consumers: processing energy accounting, data movement energy accounting, and data storage accounting. Both processing and data movement accounting are performed per thread

In this project, we have introduced a cost effective LEM infrastructure for component level power and energy monitoring by providing adequate hardware and software support for PTEA and energy interrupt. The monitoring infrastructure implements two levels of energy accounting: processing energy and data movement energy. Per core energy accounting can be done using the LEM hardware infrastructure. The infrastructure can be further used in conjunction with OS drivers in order to, to implement thread-level energy accounting at OS level.

We have validated our infrastructure on a Zynq ZC702 evaluation board. We have developed systems consisting of 1 MB core, 2 MB cores, 4 MB cores and 8 MB cores. The results from the execution of WCET benchmarks has indicated a strong correlation between the LEM based energy estimates and the physical power board measurements of more than 95%.

The implementation results indicated that the overall overhead of the proposed infrastructure is around 10%, for 14 sensors attached to 4-cores reference design. The proposed LEM has lower cost with respect to the Xilinx based performance counters, while having increased flexibility and accuracy.

Applicability and transferability of the results

A number of hardware components described in Verilog have been developed and provided as IPs to FPGA community

Financed through/by

CHIST-ERA partnership projects, PNII-IDEI – 1/CHIST-ERA/01.10.2012

Research Centre

Research Center in Computer and Information Technology

Research team

Innsbruck University (LP), Queen's University Belfast,
RWTH Aachen University,

Politehnica University of Timisoara:

Assoc. Prof. Marius Marcu
Dr. Oana Boncalo
Dr. Sebastian Fuicu
Dr. Gabriel Garban
Dr. Alexandru Amaricai
Dr. Razvan Bogdan
Dr. Cosmin Cernazanu
Ing. Lucian Bara
Ing. Madalin Ghenea
Ing. Marian Ionascu

Contact information

Assoc. Prof. Marius MARCU, PhD
Department of Computer Science
Address: Str. Bd. Vasile Pârvan, No. 2, RO300223, Timisoara,
Phone: (+40) 256 403263
E-mail: marius.marcu@upt.ro
Web: www.cs.upt.ro/~mmarcu