Power Estimation and Power Optimization Policies for Processor-Based Systems

author: José L. Ayala, advisor: Marisa López-Vallejo

Departamento de Ingeniería Electrónica
Universidad Politécnica de Madrid (Spain)
{jayala,marisa}@die.upm.es

Continuing advances in semiconductor technology have allowed dramatic performance gains for general-purpose microprocessors and embedded systems. These improvements are due both to increasing clock rates as well as to advanced support for exploiting instruction-level parallelism and memory locality using the additional transistors available in each process generation. However, as a negative consequence, this causes a significant increase in power dissipation, due to the fact that the dynamic power is proportional to both clock frequency and to switching capacitance (which increases as more devices and on-chip components are included). Thus, despite continuous attempts to reduce voltages and to design lower power circuits, power dissipation levels have steadily increased with each new microprocessor generation. Moreover, a new problem arises because the power savings achievable with low level techniques are reaching their theoretical maximum.

The PhD Thesis described here addresses the problem of power optimization in processor-based systems from different views. First, the problem of estimating the energy dissipation has been studied providing a useful tool for power estimation in the cache hierarchy. Later, several power minimization techniques have been proposed, including architectural and compiler techniques for reducing energy consumption for in-order embedded processors, out-of-order high-performance processors, and multi-processor systems on a chip (MPSoC). In this way, the different abstraction and complexity levels in the processor-based paradigm are perfectly covered, from the SoC to the MPSoC design. While the approaches described for the SoC implementation point to local energy savings (in the register file), the MPSoC approach considers the problem of global energy minimization (in the whole system).

Power optimization can act on these different parameters and its reduction is essential for all portable applications (GSM, UMTS) or spatial telecommunications. It is well-known that algorithm transformations are more efficient than technological optimizations; furthermore, such modifications are less expensive and allow to restrict the time to market. So it is very important to evaluate the application consumption at the early stage of the system design.

VLSI designers need for advanced techniques and related tools for the early estimation of power dissipation during the design phases, in order to satisfy the power constraints without significantly reducing the global performance. The goal is to meet the design turn-around time deadlines, while exploring the space of possible design alternatives. Accuracy and efficiency of a high level analysis approach should be the "booster" to meet the power requirements, avoiding a costly redesign process. It has to be pointed out that relative accuracy in power estimation is more important than absolute accuracy since, usually, the main goal is to compare alternative design solutions. However, if battery design or cooling issues are involved, accuracy of the estimation is other main point.

When dealing with the processor of an embedded system, there are several parameters that have a strong relation with the performance-power tradeoff. Mainly, the target processor and the memory cache hierarchy. In this work we have accomplished the design and use of an accurate estimation tool for cache power consumption. The tool has been designed to be highly integrated in a retargetable design tool chain. With this work the most important source of power dissipation in the processor system is completely characterized from the very early design phases, and the subsequent design takes into account the power dissipation as a constraint. The proposed methodology is extensible to new architectures, providing a helpful framework for the design of power-constrained embedded systems.

Once the power estimation problem has been partially solved, the main goal of this PhD is the energy reduction in the processor, targeting the most power-hungry devices and modules. The register file consumes a sizable fraction of the total power in embedded processors and becomes a dominant source of energy dissipation when other power saving mechanisms have been applied. Also, while an extensive previous work for energy reduction in the cache hierarchy can be found, there is no such wide research for the register file. Register file power consumption in embedded systems depends very much on system configuration, mainly on the number of integrated registers, the cache size and the existence of a branch predictor table (i.e. depends on the relative size of other memory devices). In the Motorola's M.CORE architecture, the register file energy consumption could achieve 16% of the total processor power and 42% of the data path power. In out-of-order processors with a large number of physical registers which are implemented as part of the *Re-Order Buffer* (in Pentium III, for example), this structure dissipates as much as 27% of total energy according to some estimates.

The power reduction techniques for in-order and out-of-order processors are characterized by no execution penalty and low design time overhead. They are based on the observation that a register is only used when an instruction reads from it or writes to it, the register is "idle" at all other times. By keeping the idle registers in a low power (or "drowsy") state a significant amount of energy can be saved. Most registers are idle in any given cycle since at most three registers are accessed by an issued instruction.

Next, main chapters of this work will be described with some detail.

I. Power Estimation

The cache is one of the most power-hungry devices in the processor-based systems, and will be the focus of our estimation tool. Our estimation work is built over the CGEN ("Cpu tool GENerator") core to generate embedded system design tools with explicit information of power dissipation in caches. As CGEN successfully does with cross-design tools, the cache power estimator can be automatically generated for the whole set of available target processors in order to free the designer from this annoying task.

In order to get an accurate estimation of cache power dissipation, the exact memory access pattern must be known to feed an analytical model of power consumption. This model provides mathematical expressions for the several energy dissipation sources in the cache architecture and takes as arguments parameters of different kinds: architectural (cache size, way configuration, word line size, etc.), technological (line and device input/output capacitors, etc.) and statistical (switching bus activity, cache accesses, etc.). While the technological and architectural parameters are known from the architecture design phase, the access pattern depends on the particular data in use. Furthermore, to calculate the exact value of the switching activity, perfect knowledge of the transferred data and addresses is needed. Such simulation detail can only be achieved at instruction level, when the memory and register contents are also available.

To have access information, the running code must generate a detailed memory trace (memory address accessed, data transferred, type of the access) per memory access. To get a machine independent implementation of this trace facility, we have supplied the target machine with four new virtual registers. The resulting tool has been extensively used to perform several design explorations in different scenarios and to evaluate the impact of the switching activity, the input data correlation, the compiler optimizations or the cache hierarchy in the power consumption.

Our work in this area can be found in [1] and [2].

II. Power Reduction for In-Order-Processors

The first approach we propose for power-aware register file reconfiguration for in-order embedded processors is based on compiler support and code profiling. With a minor change in the ISA (*Instruction Set Architecture*), the compiler can dynamically select the most suitable register file configuration for application requirements.

The proposed solution to deal with the "scalable" register file works as follows. Instead of working with the whole register file, only the currently used registers are addressable and maintained on. The rest of the devices are put into a drowsy state where the static power consumption is reduced to a minimum and the clock distribution network is gated as well.

It is assumed that the ISA is augmented with one instruction to turn a set of registers on and off, therefore, one clock cycle is lost whenever the register file is reconfigured. This register file management technique requires a code exploration and code generation phase that marks the code sections where the register file reconfiguration should be performed. This phase has to detect those frequently executed code sections with reduced number of required registers and mark the beginning and end of such portion of code. This power aware compilation has been implemented in the GNU compiler, *gcc*.

The architectural modifications we propose in the second approach allow the drowsy registers to be turned back to the "active" state as the instruction accesses them. The proposed technique keeps the registers in the register file in the low-power state (drowsy state) until they are needed. Due to the necessity of restoring the register contents previously to the access, the accessed registers must be known at least one cycle before the access happens. During the fetch stage the instruction is loaded from I-cache into the instruction register. Our approach takes advantage of this fact by forwarding the operand fields of the instruction to the register file address decoders. As can be noticed, the main advantages of this approach are the no-time penalty on the predecode stage, the simplicity of the extra logic (only a few modifications to the memory cell, the gating circuit and the bypass network) that means small energy and area overhead and, finally, the possibility of applying this approach to different microprocessor architectures present in many embedded systems.

The work developed for in-order processors can be found in the following references: [3], [4] and [5].

III. Power Reduction for Out-Of-Order Processors

For out-of-order processors, the indetermination of the register renaming task makes unfeasible the direct use of the previous approach: the allocation of a free register is a non deterministic task and it is not possible to know in advance the rename buffer assigned to a destination register. The main goal of this work is to reduce this indetermination so that the required rename buffer can be known in advance to be turned on. Our proposal is based on a deterministic variation of the renaming algorithm which allows to turn-off many registers of this structure if they are not used, with the corresponding energy savings. The required registers are turned-back on in advance and no performance penalty is accomplished. Moreover, the design of a power-aware compiler is also presented what increases the energy savings achieved by the hardware technique.

The M-positions FIFO of free registers was split into four N-position FIFOs (with $N < M$). The size of these FIFOs has to be large enough to provide rename buffers to the destination registers but avoiding to become empty. In this way, the problem of indetermination has been mostly solved. Instead of assigning free rename buffers from the same FIFO to the whole set of logical registers, the first set of logical registers is assigned to the first FIFO, the second set to the second FIFO, and so on. The worst execution case happens when every operand coded in the instruction word belongs to two different FIFOs (spread assignment) and the expected savings are the lowest achievable.

The power-aware compiler also designed for supporting this technique is able to increase the energy savings expected for the hardware approach by modifying the register assignment. In this work the register assignment performed by the GNU compiler (*gcc*) has been modified in order to minimize the concurrent use of more than one FIFO. Therefore, increased energy savings can be obtained for a solution closer to the theoretical maximum (only one FIFO active per instruction, what can provide a 75% of energy savings).

Our work in this area can be found in [6] and [7].

IV. Power Reduction for MPSoC

Computation-intensive parallel applications are pushing the development of complex multi-processor Systems-on-Chip (MPSoCs), which rely on the high integration densities made available by deep sub-micron technologies. In the context of heterogeneous MPSoCs, performance and energy efficiency of each architecture's processing sub-system cannot be evaluated in isolation, since its effectiveness can be substantially impacted by the interaction with the other sub-systems. Therefore, many decisions about the core's features have to be made concurrently and jointly assessed. Setting the operating frequency of MPSoCs' processing elements is an example of such critical design issues. Heterogeneous MPSoCs will integrate several cores running at different speeds.

The proposed approaches for MPSoC make use of the frequency and voltage scaling techniques to reduce the overall power consumption of the system. The analysis of second order effects with high impact on the system performance and energy consumption, as they are bus congestion or processor synchronization, count with no previous analysis in the scientific comunity.

This work focuses on the system level interaction of cores running at different frequencies, assessing the impact on overall system performance, energy dissipation and energy-delay product. When one processor is relatively slower or faster than the others in presence of benchmarks requiring inter-processor synchronization and communication, the negative impact on bus utilization has been demonstrated. The power-performance trade-off has been explored when multiple cores are allowed to work at different speeds at the same time: degradation of execution time is counterbalanced by a lower energy consumption as the number of frequency-scaled processors increases. We have provided frequency selection guidelines to minimize energy dissipation when an application workload can be partitioned among a variable number of processors. Also, the relative impact of OS support has been highlighted (execution time almost doubled for communication-intensive benchmarks), and the benefits of increasing clock frequency for processors with heavy asymmetric workloads in terms of both execution time and energy dissipation have been thoroughly analyzed. Finally, the knowledge acquired during the design and analysis phases has been exploited to devise several power management policies which minimize energy consumption under performance constraints.

Our work in this area can be found in [8].

V. Supporting Paper and Other Publications

The supporting paper [5] belongs to a journal publication in the *International Journal of Parallel Programming* (Kluwer Academic Publishers). This publication presents both the compiler and architectural approaches designed for in-order embedded architectures. Other publications refered to the dissertation topic are the followings:

References

[1] J. L. Ayala and M. López-Vallejo, "A unified framework for power-aware design of embedded systems," in *IEEE International Workshop on Power and Timing Modeling, Optimization and Simulation*, September 2003, also published as Lecture Notes on Computer Science, vol. 2799 (Springer Verlag).

[2] J. L. Ayala, M. López-Vallejo, and C. L. Barrio, "A case study on power dissipation in the memory hierarchy of embedded systems," in *Design of Circuits and Integrated Systems Conference*, November 2003.

[3] J. L. Ayala, M. López-Vallejo, A. Veidenbaum, and C. A. López, "Energy aware register file implementation through instruction predecode," in *IEEE International Conference on Application-Specific Systems, Architectures and Processors*, June 2003.

[4] J. L. Ayala, M. López-Vallejo, and A. Veidenbaum, "Reducing register file energy consumption using compiler support," in *IEEE Workshop on Application Specific Processors*, November 2002.

[5] ——, "Power-aware compilation for register file energy reduction," *International Journal of Parallel Programming (Kluwer Academic Publishers)*, vol. 31, no. 6, pp. 449–465, December 2003.

[6] ——, "Energy-efficient register renaming in high-performance processors," in *IEEE Workshop on Application Specific Processors*, December 2003.

[7] ——, "Low-power register assignment for out-of-order architectures," in *IEEE Workshop on Innovative Architectures*, May 2004.

[8] J. L. Ayala, , D. Bertozzi, L. Benini, and M. López-Vallejo, "Omitted for blind review," in *Submitted to ISLPED*, August 2004.

VI. Personal Data

José L. Ayala (Departamento de Ingeniería Electrónica, Universidad Politécnica de Madrid (Spain), jayala@die.upm.es)
advisor: Marisa López-Vallejo (marisa@die.upm.es)

estimated graduation date: Beginning 2005