

High-Level Energy Estimation for ARM-Based SOCs

Dan Crisu¹, Sorin Dan Cotofana¹, Stamatis Vassiliadis¹, and Petri Liuha²

¹ Computer Engineering Laboratory, Electrical Engineering, Mathematics and Computer Science Faculty, Delft University of Technology Mekelweg 4, 2600 GA Delft, The Netherlands
{dan,sorin,stamatis}@ce.et.tudelft.nl

² Nokia Research Center, Visiokatu-1, SF-33720 Tampere, Finland
petri.liuha@nokia.com

Abstract. In recent years, power consumption has become a critical concern for many VLSI systems. Whereas several case studies demonstrate that technology-, layout-, and gate-level techniques offer power savings of a factor of two or less, architecture and system-level optimization can often result in orders of magnitude lower power consumption. Therefore, the energy-efficient design of portable, battery-powered systems demands an early assessment, i.e., at the algorithmic and architectural levels, of the power consumption of the applications they target. Addressing this issue, we developed an energy-aware architectural design exploration and analysis tool for ARM based system-on-chip designs. The tool integrates the behavior and energy models of several user-defined, custom processing units as an extension to the cycle-accurate instruction-level simulator for the ARM low-power processor family, called the ARMulator. The models we implemented take into account the particular class, e.g., datapath, memory, control, or interconnect, as well as the architectural complexity of the hardware unit involved and the signal activity triggered by the specific algorithm executed on the ARM processor. Our tool can estimate at the architectural level of detail the overall energy consumption or can report the energy breakdown among different units. Preliminary experiments indicated that the estimation accuracy is within 25% of what can be accomplished after a circuit-level simulation on the laid-out chip.

1 Introduction

With the advent of mobile platforms for computing and communications, system designers and integrators were confronted with a massive shortage of tools that enable early energy consumption estimation for such systems. CAD tool support for embedded system design is still limited and it addresses mainly functional verification and performance estimation.

The intricacy involved by these new electronic appliances imposed a new design paradigm to cope with the specific requirements, e.g., low cost with fast time to market, and restrictions they have. Also, energy consumption is a critical factor in system-level design of embedded portable appliances. A hardware-software co-design framework must be employed to proceed with the design from the software applications intended to run on these appliances to the final specifications of the hardware that implements the desired functionality given the above-mentioned constraints. Studies have demonstrated that circuit- and gate-level techniques have less than a $2\times$ impact on power,

while architecture- and algorithm-level strategies offer savings of 10 – 100× or more [1]. Hence, the greatest benefits are derived by trying to assess early in the design process the merits of the potential implementation. Architecture optimization corresponds to searching for the best design that optimize all objectives. Since the optimization problem involves multiple criteria (power consumption, throughput, and cost) to reach the global optimum a set of Pareto points [2] in the design space have to be found. Ideally, when designing an embedded system, a designer would like to explore a number of architectural alternatives and test functionality, energy consumption, and performance without the need to build a prototype first.

Usually, typical portable systems are built of commodity components and have a microprocessor-based architecture. Full system evaluation is often done on prototype boards resulting in long design times. Power consumption estimation can be done only late in the design process, after the prototype board was built, resulting in slow power tuning turnarounds that doesn't meet the requirement of fast time to market. On the other hand, using field programmable gate array (FPGA) hardware emulators for functional debugging, with a fast prototyping time, can neither give accurate estimates of energy consumption nor of the performance.

Among the tools preferred for early performance assessment at the algorithmic and architectural level, in the last decade, were the cycle-accurate instruction-set simulators. Unfortunately, for power consumption estimation this approach was seldom easy to follow. There were only a few academic tools for power estimation (all based on or integrated in the SimpleScalar instruction set simulator toolset framework [3], [4], [5]) and almost no commercial products.

For several target general purpose processors a number of techniques emerged in the last few years. The processor energy consumption for an instruction trace was generally estimated by instruction-level power analysis [6], [7]. This technique estimates the energy consumed by a program by summing the energy consumed by the execution of each instruction. Instruction-by-instruction energy costs, together with non-ideal effects, are precharacterized once for each target processor. A few research prototype tools that estimate the energy consumption of processor core, caches, and main memory have been proposed [8], [9]. Memory energy consumption is estimated using cost-per-access models. Processor execution traces are used to drive memory models, thereby neglecting the non-negligible impact of a non ideal memory system on program execution. The main limitation of these approaches is that the interaction between memory system (or I/O peripherals) and processor is not modeled. Cycle-accurate register-transfer level energy estimation was proposed in [4]. The tool integrates RT level processor simulator with DineroIII cache simulator and memory model. It was shown to be within 15% of HSPICE simulations.

The drawback of all the above methods to estimate the power consumption is that they are based on certain architectural templates, i.e., general purpose processors and can be hardly adapted to model system-on-chip designs.

A new approach towards high-level power estimation is presented in this paper in the context of ARMulator [10], a cycle-accurate instruction-level simulator for the ARM low-power processor family. More in particular, we developed an energy-aware architectural design exploration and analysis tool for ARM based system-on-chip de-

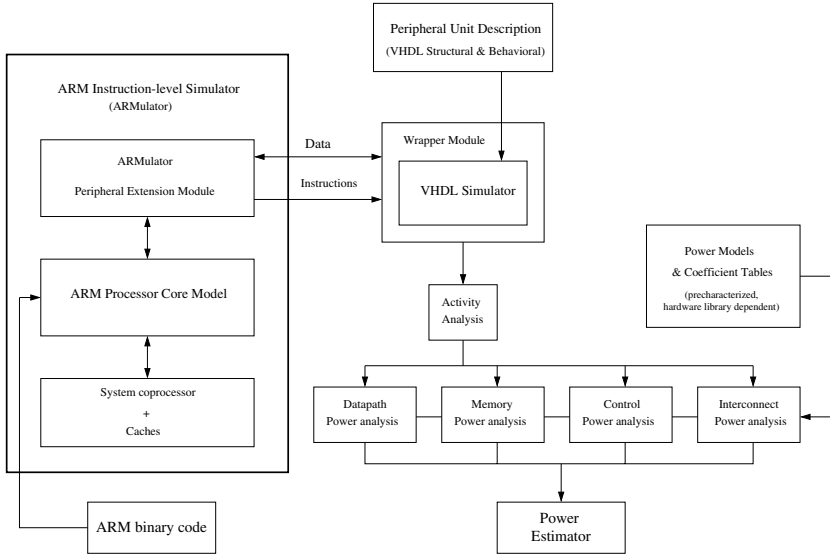


Fig. 1. System-on-chip simulator architecture.

signs. The tool integrates the behavior and energy models of several user-defined, custom processing units as an extension to the ARMulator. These models take into account the impact of design complexity and signal activity on datapath, memory, control, and interconnect power consumption. So far we have implemented only the tool framework and the power calculators for the datapath part. Experiments carried on employing a sample coprocessor design indicated that the accuracy of the results obtained by behavioral simulation is within 25% of that obtained using circuit simulators.

The rest of the paper is organized as follows. We present the system model and the methodology for cycle-accurate simulation of energy dissipation in Section 2. In Section 3, to validate the employed methodology we design down to the physical layout a sample coprocessor for an ARM1020T CPU core in order to run a number of realistic experiments and prove the effectiveness of the proposed high-level methodology. Finally, Section 4 presents the conclusions and describes future work in the area.

2 Proposed Design Exploration Framework

In this section we present our approach to estimate at the architectural level the power consumption of a coprocessor or peripheral unit coupled with an ARM CPU core on a system-on-chip.

2.1 System Model

Figure 1 gives an overview of the power analysis strategy that we propose.

We use this strategy for the design of peripheral units that augment or complement ARM CPU core functionality. The instruction set architecture of the ARM family of processors offers room for extensions to be added by providing the so called coprocessor instructions. Referring to the Figure 1, the inputs from the user are a description of a candidate architecture for the desired peripheral unit given in behavioral or structural VHDL and the set of data and the application program for which a power analysis is desired. The provided program is then compiled using the ARM native compiler. Usually, the code will embed, beside ARM native instructions, specific peripheral unit instructions. These specific instructions, when executed on the ARMulator (the ARM instruction set simulator), will be recognized as non-native or coprocessor instructions and they will trigger callback functions, installed using the ARMulator API (application programming interface), so specific actions (e.g., new data or commands are fed to the hardware description simulated in VHDL) can be taken. Moreover, every clock cycle, the ARMulator will send signals to the VHDL simulator to advance the state of the simulated hardware description one more clock cycle. In this simple way the simulated hardware description will process its own data in lockstep with the ARM processor pipeline. Every clock cycle the activity on internal relevant signals is also collected and sent to the power analysis units. Rather than attempting to find a single power model for the entire chip, we take the approach of identifying four basic classes of components: datapath, memory, control, and interconnect. The total power consumption of the coprocessor or peripheral unit per program executed on the ARM processor is estimated.

The central elements of our architectural design exploration framework are:

- ARMulator, the cycle-accurate instruction set simulator for the ARM family of low-power processors;
- VHDL simulator, capable of saving the state of the simulated hardware description whenever it receives this command, also it is capable to reinitialize the hardware description with the previous saved state before it processes the new stimuli sent by the ARM CPU core simulated on the ARMulator;
- Wrapper Module, that handles the communication between ARMulator (using ARMulator API) and the VHDL simulator; it is also responsible of passing correct formatted data between ARMulator, the VHDL simulator, and the activity analysis module;
- Precharacterized Power models and Effective Capacitance Coefficient Tables Module, that contain for a library of hardware cells all the technology dependent information required by the power analysis modules to compute the power consumption; the tables are derived only once for a given library of hardware cells (more detailed explanations are given in Subsection 2.2);
- Activity Analysis Module, that feeds the Power Analysis modules (power calculators) with statistics about signal activity inside the simulated hardware description;
- Power Analysis Modules, that estimate the power consumption in the datapath, control, memory, and interconnect based on statistics received from the Activity Analysis Module and lookups in the effective capacitance coefficient tables;
- Power Estimator Module, that adds the estimates of power consumption of datapath, control, memory, and interconnect and offers the total figure of power consumption of the coprocessor or peripheral unit per program executed on the ARM processor;

The approach we have taken provide all the benefits of a co-design framework, moreover, it is also capable of power estimation:

- permits experimental partitioning schemes between features that must be provided by software and features that will be mapped in hardware;
- allows changing of the organization and order of execution of the algorithmic blocks to investigate and verify potential new architectures;
- provides methods of performance monitoring (in terms of throughput, and power consumption);
- accelerates the implementation of new algorithms and provides an environment in which to test them both individually and as part of an entire pipeline;
- allows tweaking the bit width precision and seeing potential impacts on result accuracy and performance factors.

2.2 Power Models

In this subsection, we will describe the methodology of modeling power consumption at the architecture level. We followed the methodology presented in [1]. The premise for the success of such methodology consists in the existence of a library of hardware cells consisting of various operators for the datapath part, gates for control logic, and bit-cells, decoders, sense amplifiers for memory cores. Depending on the estimation accuracy desired, the individual cells can be specified at gate-level if the gates employed are already characterized for power, or can be specified at the layout-level in order for the internal interconnect parasitics between individual constituent transistors to be accounted for. Once such a library exists, it can be precharacterized via gate-level, respectively circuit-level simulations, resulting in a table of effective capacitive coefficients for every element in the library. Then using only this tables and the activity statistics derived during the architectural-level simulation the power consumption can be estimated easily. This precharacterization has to be done only once and only the effective capacitive coefficients table are needed for power estimation. The precharacterization results are valid only for a specific library of hardware cells and a given IC technology.

The power estimation methodology presented in [1] analyzes separately the four main classes of chip components: datapath, memory, control, and interconnect. For the first two classes, a model called the Dual Bit Type (or DBT) model was developed which demonstrated good accuracy results, with power estimates typically within 10 – 15% of results from switch-level simulations. The DBT model achieves its high accuracy by carefully modeling both physical capacitance and circuit activity. The key concept behind the technique is to model the activity of the most significant (sign) bits and least significant bits separately due to the fact they exhibit different statistical behavior as presented in Figure 2. The least significant bits are modeled as uniform white noise (UWN). The DBT model applies only to parts of the chip that manipulate data. A separate model is introduced to handle power estimation for control logic and signals. This model is called the Activity-Based Control (ABC) model. The method relies on the observation that although the implementation style of the controller (e.g., ROM, PLA, random logic, etc.) can heavily impact the power consumption, it is still possible to identify a number of fundamental parameters that influence the power consumption regardless of the implementation method. In a chip, datapath, memory, and control blocks

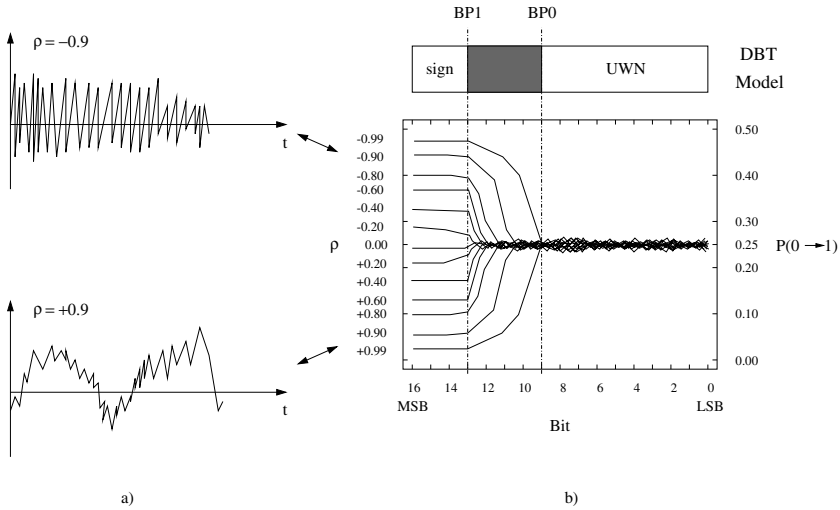


Fig. 2. Bit transition activity for 2 two’s complement data streams modeled as Gaussian processes with different temporal correlation ρ : a) Activity for positively and negatively correlated waveforms. b) Bit transition activity for data streams with varying temporal correlation.

are joined together by an interconnect network. The wires comprising the network have capacitance associated with them and, therefore, driving data and control signals across this network consumes power. The precise amount of power consumed depends on the activity of the signals being transferred, as well as the physical capacitance of the wires. The DBT and ABC models provide the activity information for control and data buses, but the physical capacitance depends on the average length of the wires in each part of the design hierarchy. The average length of the wires is estimated based on the Rent’s rule.

Having the library of hardware cells specified, for example, at the layout-level the library is precharacterized first with a circuit-level simulator. During the precharacterization stage of the library of hardware cells a black-box model of the capacitance switched in each module for various types of inputs is produced. If desired, these capacitance estimates can be converted to an equivalent energy, $E = CV^2$, or power, $P = CV^2 f$. The black-box capacitance models can be parameterized, i.e., taking into account the size or complexity of the module. The model accurately accounts for activity as well as physical capacitance. As a result, the effect of the input statistics on module power consumption is reflected in the estimates.

For illustrative purposes, we will exemplify only briefly the modeling of the capacitance of a ripple-carry subtractor. Intuitively, the total power consumed by a module should be a function of its complexity, i.e., size. This reflects the fact that larger modules contain more circuitry and, therefore, more physical capacitance. The amount of circuitry and physical capacitance an instance of this subtractor will contain is determined by its word length, N . In particular, an N -bit subtractor can be realized by N one-bit full-subtractor cells. The total module effective capacitance function should, therefore, receive an argument proportional to the word length as shown here:

$$C_T = f(\text{activity_statistics}, C_{eff}N) \quad (1)$$

where C_{eff} is the average capacitive coefficients per bit table, f represents the total module effective capacitance function, and C_T represents the total module effective capacitance. The average capacitive coefficients per bit table is obtained after a process of data fitting (using least-squares approximation method) employing the effective capacitive coefficients tables generated for several sample width of the datapath (e.g., 4 bits, 8 bits, 16 bits, 32 bits). The average capacitive coefficients per bit table will be generated for the subtractor during the precharacterization stage and stored for the time when the estimate of the power consumption is needed. Many modules besides the subtractor also follow a simple linear model for the argument of the total module effective capacitance function. For example, ripple-carry adders, comparators, buffers, multiplexers, and Boolean logic elements all obey Equation (1). The DBT method is not restricted to linear capacitance models and can model non-linear modules like array multipliers and logarithmic shifters. The total module effective capacitance function f is actually the power model of the module under consideration and receives the activity statistics seen on the module terminals, the complexity parameters of that module (e.g., N), and a pointer to the average capacitive coefficients per bit table for that module. The reader is referred to [1] for more details. The total module effective capacitance C_T represents the effective capacitance switched by that module every clock cycle during the execution of an application program on the ARM processor.

3 Experimental Results

To verify the power consumption prediction accuracy of the architectural space exploration tool we designed, we need to provide the following experimental setup: a precharacterized library of hardware cells, the description in VHDL of the peripheral or coprocessor to be simulated, and the binary code of the program to be simulated on the ARMulator containing calls to the coprocessor.

We precharacterized parts of a datapath library of cells (including a ripple-carry subtractor) designed in UMC 0.18 μ m Logic 1.8V/3.3V 1P6M GENERICII CMOS technology. We extracted from the layout the circuit of the subtractor in three variants of the datapath width: for 4 bits, 8 bits, and 16 bits. After using the method presented in Subsection 2.2 we obtained the average capacitive coefficients per bit values presented in Table 1.

We modeled in VHDL a sample coprocessor for an ARM1020T CPU core. It was designed starting from the datasheet of AMD's Am2901 four-bit bipolar microprocessor slice. The coprocessor has a datapath width of 8 bits. The coprocessor consists of a 16-word by 8-bit two-port register file, an ALU and the associated shifting, decoding and multiplexing circuitry. The 9-bit microinstruction word is organized in three groups of three bits each and selects the ALU source operands, the ALU function, and the ALU destination register. The ALU provides various status flag outputs. The ALU is capable of performing three binary arithmetic ($R + S$, $S - R$, $R - S$) and five logic functions ($R \text{ OR } S$, $R \text{ AND } S$, $\overline{R \text{ AND } S}$, $R \text{ XOR } S$, $R \text{ XNOR } S$).

To generate the application programs we analyzed real trace data for environmental control realized with well known microcontrollers (Intel 8051 and compatible). We ex-

Table 1. Average capacitive coefficients per bit for the ripple-carry subtractor.

Transition Templates	Capacitive Coefficients (fF/bit)			
UU/UU	35.12			
UU/SS	34.05	51.26	49.73	17.26
SS/UU	0.00	28.32	41.02	29.56
SS/SS/SS	0.00	28.33	41.17	3.15
	41.92	0.00	18.61	0.00
	52.13	25.71	0.00	0.00
	0.51	0.00	0.00	0.00
	0.00	35.42	0.00	52.12
	16.00	38.02	36.21	15.69
	0.00	55.79	0.00	0.00
	49.10	27.42	0.00	0.00
	0.00	0.00	69.21	60.01
	0.00	0.00	43.17	0.00
	7.65	32.73	48.82	10.13
	46.18	0.00	19.29	0.00
	0.00	0.00	0.00	2.41
	0.00	0.00	54.92	72.12
	0.00	0.15	0.00	46.89
0.75	25.02	22.83	2.81	

tracted the recurrent patterns of control and data in these instruction flows and generated three instruction flows A, B, and C, along with the data using biased noise generators. We used biased noise generators because in the case of our sample coprocessor there is no already developed software available. We executed these instruction flows on the ARM processor family ISA and, using the framework described in Subsection 2.1, we obtained power consumption estimates for the subtractor. They are presented in the second column of Table 2.

Table 2. Power consumption results for the ripple-carry subtractor.

Instruction Trace	Power Consumption (estimated)	Power Consumption (simulated)	Relative Error (%)
A	0.77mW	0.91mW	-15
B	1.02mW	0.84mW	21
C	0.63mW	0.61mW	3

In order to find the relative error of these estimations for the ripple-carry subtractor we have to compare the results obtained employing our design exploration tool with the power consumption estimated accurately with the HSPICE circuit simulator on exactly the same excitation patterns for the ripple-carry subtractor. For this purpose we designed down to the layout-level the sample coprocessor using the library of hardware cells. The layout of the sample coprocessor is presented in Figure 3.




Fig. 3. Sample coprocessor layout. From left to right and up to down: register file, control, and datapath.

The simulation results on the extracted netlist of the ripple-carry subtractor are presented in the third column of Table 2. The clock frequency for the sample coprocessor assumed throughout these experiments is 200MHz. We have to mention here that the circuit-level simulation of the subtractor took several hours for the three instruction traces executed on the ARM processor. This clearly indicate that a circuit-level simulation of the whole coprocessor, to obtain the power consumption directly, for an instruction trace executed on the ARM processor is computationally unfeasible. The relative error between the power estimated and the power consumption obtained by circuit-accurate simulation is presented in the last column of Table 2. The power prediction accuracy is good, well within 25% of a direct circuit simulation with HSPICE.

4 Conclusions

A new approach towards high-level power estimation is presented in this paper in the context of ARMulator, a cycle-accurate instruction-level simulator for the ARM low-power processor family. More in particular, we developed an energy-aware architectural design exploration and analysis tool for ARM based system-on-chip designs. The tool integrates the behavior and energy models of several user-defined, custom processing units as an extension to the ARMulator. The models we implemented take into account the particular class, e.g., datapath, memory, control, or interconnect, as well as the architectural complexity of the hardware unit involved and the signal activity triggered by the specific algorithm executed on the ARM processor. Our tool can estimate at the architectural level of detail the overall energy consumption or can report the energy breakdown among different units. Preliminary experiments indicated that the estimation accuracy is within 25% of what can be accomplished after a circuit-level simulation on the laid-out chip.

Our endeavor to accurately predict power consumption within the ARM-based system-on-chip designs is an ongoing work. We have to mention that, up to date, the precharacterization of the library is done in a manual way. In the future, we intend

to develop an automatic process of precharacterization. We believe that a command language with a dedicated grammar for the precharacterization process can be a possible solution. Within this approach every leaf cell in the library of hardware cells will be accompanied by a description file with specific commands for the precharacterization process. These issues will be addressed in the near future.

References

1. Landman, P.: High-Level Power Estimation. In: International Symposium on Low Power Electronics and Design, Monterey CA (1996) 29–35
2. Micheli, G.D. In: Synthesis and Optimization of Digital Circuits. McGraw-Hill (1994)
3. Burger, D., Austin, T.M.: The SimpleScalar Tool Set, Version 2.0. Technical Report Nr. 1342, University of Wisconsin-Madison Computer Sciences Department (1997)
4. Vijaykrishnan, N., Kandemir, M., Irwin, M.J., KIM, H.S., Ye, W.: Energy-Driven Integrated Hardware-Software Optimizations Using SimplePower. ISCA 2000 (2000)
5. Brooks, D., Tiwari, V., Martonosi, M.: Wattch: A Framework for Architectural-Level Power Analysis and Optimizations. In: Proceedings of the 27th International Symposium on Computer Architecture, Vancouver, BC (2000) 83–94
6. Tiwari, V., Malik, S., Wolfe, A., Lee, M.: Instruction Level Power Analysis and Optimization of Software. *Journal of VLSI Signal Processing Systems* **13** (1996) 223–238
7. Tiwari, V., Malik, S., Wolfe, A.: Power analysis of embedded software: A first step toward software power minimization. *IEEE Transactions on VLSI Systems* **2** (1994) 437–445
8. Li, Y., Henkel, J.: A Framework for Estimating and Minimizing Energy Dissipation of Embedded HW/SW Systems. In: Proceedings of Design Automation Conference. (1998) 188–193
9. Kapoor, B.: Low Power Memory Architectures for Video Applications. In: Proceedings of 8th Great Lakes Symposium on VLSI. (1998) 2–7
10. ARM Limited: ARM Developer Suite version 1.1. (1999)