

Instruction Level Energy Modeling for Pipelined Processors*

S. Nikolaidis¹, N. Kavvadias¹, T. Laopoulos¹, L. Bisdounis², and S. Blionas²

¹Department of Physics,
Aristotle University of Thessaloniki, 54124 Thessaloniki, Greece
snikolaid@physics.auth.gr
²INTRACOM S.A., Peania, Greece

Abstract. A new method for creating instruction level energy models for pipelined processors is introduced. This method is based on measuring the instantaneous current drawn by the processor during the execution of the instructions. An appropriate instrumentation set up was established for this purpose. According to the proposed method the energy costs (base and inter-instruction costs) are modeled in relation to a reference instruction (e.g. NOP). These costs incorporate inter-cycle energy components, which cancel each other when they are summed to produce the energy consumption of a program resulting in estimates with high accuracy. This is confirmed by the results. Also the dependencies of the energy consumption on the instruction parameters (e.g. operands, addresses) are studied and modeled in an efficient way.

1 Introduction

Embedded computer systems are characterized by the presence of dedicated processors which execute application specific software. A large number of embedded computing applications are power or energy critical, that is power constraints form an important part of the design specification [1]. Recently, significant research in low power design and power estimation and analysis has been developed. The determination of methods for the accurate calculation of the energy consumption in processors is required for system level energy consumption evaluation and optimization.

Power analysis techniques for embedded processors that employ physical measurements were firstly suggested in mid 90's. Significant effort on software optimization for minimizing power dissipation is found in [1]-[4] where a technique based on physical measurements is developed. Power characterization is done with the extraction of cost factors for the average current drawn by the processor as it repeatedly executes short instruction sequences. The base cost for an instruction is determined by constructing a loop with several instances of the same instruction. Inter-instruction effect induced when executing different adjacent instructions, is measured by replacing the one-instruction loops used in the measurement of base costs, with loops consisting of appropriate instruction pairs. The sum of the power base cost of the instructions

* This work was supported by EASY project, IST-2000-30093, funded by the European Union

executed in a program, refined by the power cost of the inter-instruction effects, are considered to provide the power cost of the program. This method has been validated for commercial targets based on embedded core processors.

The majority of work published on the field of measurement-based techniques, refers to the Tiwari method [1] as a base point. By Tiwari method only average power estimates can be utilized for modeling task, since the measurements are taken with a standard digital ammeter. Direct application of the Tiwari technique is found in [5] where an extensive study of the ARM7TDMI processor core is reported. In [6], physical measurements for the processor current are also obtained by a precise amperemeter. However, power modeling effort is more sophisticated, as architectural-level model parameters are introduced and integrated within the power model. These consist of the weight of instruction fields or data words, the Hamming-distance between adjacent ones, and basic costs for accessing the CPU, external memory and activating/deactivating functional units.

Instantaneous current is firstly measured in [7], where a digitizing oscilloscope is used for reading the voltage difference over a precision resistor that is inserted between the power supply and the core supply pin of the processor. Instantaneous power is then calculated directly from the voltage waveform from which average figures are extracted to guide instruction power modeling. Resistor-based methodologies suffer from supply voltage fluctuations over the processor which reduces the accuracy of the method.

All the above techniques acquire the current drawn by the processor on instruction execution. A complex circuit topology for cycle-accurate energy measurement is proposed in [8,9], which is based on instrumenting charge transfer using switched capacitors. The switches repeat on/off actions alternately. A switch pair is charged with the power supply voltage during a clock cycle and is discharged during the next cycle powering the processor. The change in the voltage level across the capacitors is used for the calculation of the consumed energy in a clock cycle. However, this method can not provide detail information for the shape of the current waveform, which may be significantly useful in some applications and also in case high quality power models including the architectural characteristics of the processor are required. In order to measure the energy variations, various (*ref*, *test*) instruction pairs are formed, where *ref* notes a reference instruction of choice and *test* the instruction to be characterized. This setup combined with the above modeling concept are then utilized to obtain an energy consumption model for the ARM7 processor.

2 Instruction-Level Energy Modeling for Pipelined Processors

The proposed method is based on the measurement of the instantaneous current drawn by the processor during the execution of the instructions. The instrumentation set-up for measuring the instantaneous current has been presented in [10]. The current sensing circuit is a high performance current mirror, which is inserted in the supply line and copies the drawn current on a resistor. In this way the supply voltage fluctuation at the processor is minimized increasing the accuracy of the measurements. A high-speed digitized oscilloscope is used for monitoring the voltage drop on the resistor.

The method is developed for pipelined processors like the ARM7 (three-stage pipeline). However, its application for non-pipelined processors is straightforward.

2.1 Energy Consumed in Pipelined Processors

The current drawn by the processor in a clock cycle is the result of the concurrent process of many instructions due to the pipeline structure of the processor. So, it is rather impossible to isolate the current drawn from the pipeline stage, which processes the test instruction. Instead, we have to find a method for calculating the energy consumed by an instruction from the current drawn in the clock cycles needed for its execution.

In each clock cycle the instantaneous current, $i(t)$ is monitored and the energy consumed in that cycle, E_{cycle} is calculated as:

$$E_{cycle} = V_{DD} \int_0^T i(t) dt \quad (1)$$

where V_{DD} is the value of the supply voltage and T is the clock period. The energy which is consumed as an instruction passes through a n -stage pipeline, $E_M(Instr)$, is distributed in n clock cycles and it is found by summing the energy components of the n cycles:

$$E_M(Instr) = E_{cycle_1} + E_{cycle_2} + \dots + E_{cycle_n} . \quad (2)$$

In n clock cycles n instructions are executed. Suppose a set of instructions with only one instance of the *test* instruction, and with the others being the same, let us say a reference instruction (probably the NOP), running through the pipeline. Then in n clock cycles one *test* instruction and $n-1$ reference instructions are executed, and the energy corresponding to the test instruction is calculated as:

$$E(Instr) = E_M(Instr) - (n-1)E(ref) . \quad (3)$$

Since the energy budget of a program corresponds to the sum of the energy of each instruction, such a model seems to lead to an overestimation because instead of modeling the circuit state changes from one instruction to another, we model every instruction as a circuit state change (change from test to NOP instruction), which results in counting the number of circuit state changes twice. To overcome this problem the inter-instruction effect is also modeled in a similar way, and in this case the value of this effect is expected to be rather (at least in most cases) negative compensating this overhead.

2.2 Base Instruction Cost

In the proposed method as base instruction cost we mean the amount of energy which is consumed for the execution of an instruction after the execution of a reference instruction (NOP). In order to calculate the base instruction costs, loops including an instance of the test instruction and a number of instances of reference instruction are executed on the processor. Figure 1 shows the pipeline states during the clock cycles

required for the execution of the instruction on a three-stage pipelined processor. The instruction is executed in three cycles. In each clock cycle the current is monitored and the energy is calculated.

Pipeline Stages	3-stage pipeline operation				
IF	NOP	NOP	Instr	NOP	NOP
ID	NOP	NOP	NOP	Instr	NOP
EX	NOP	NOP	NOP	NOP	Instr
Clock Cycles	n-1	n	n+1	n+2	n+3

Fig. 1. Pipeline states during the execution of instructions

The energy consumed in these three stages is calculated as:

$$E_M(Instr) = E_{cycle_n+1} + E_{cycle_n+2} + E_{cycle_n+3} \tag{4}$$

In these three cycles one test instruction and two reference instructions are executed. This amount of energy contains inter-cycle effects due to the change of the pipeline states and it can be expressed as

$$E_M(Instr) = E_{Instr} + 2E_{NOP} + E_{NOP,Instr} + E_{Instr,NOP} + E_{NOP,NOP} \tag{5}$$

where E_{inst} corresponds to the real energy of the instruction, E_{NOP} to the energy of the NOP instruction and the others to the inter-cycle effects between the cycles $n, n+1$ ($E_{NOP,Instr1}$), $n+1, n+2$ ($E_{Instr1,NOP}$) and $n+2, n+3$ ($E_{NOP,NOP}$).

The energy of the instruction is calculated as:

$$E(Instr) = E_M(Instr) - 2E_{NOP} \tag{6}$$

and corresponds to

$$E(Instr) = E_{Instr} + E_{NOP,Instr} + E_{Instr,NOP} + E_{NOP,NOP} \tag{7}$$

The energy of the NOP instruction is:

$$E_{NOP} = E_{cycle_n} \tag{8}$$

2.3 Inter-instruction Effect Cost

The instruction sequence for measuring the inter-instruction effect is shown in Figure 2. Appropriate loops will be executed on the processor. In this case, in four clock cycles one *Instr1*, one *Instr2* and two reference instructions are executed.

Pipeline Stages	3-stage pipeline operation				
IF	NOP	Instr1	Instr2	NOP	NOP
ID	NOP	NOP	Instr1	Instr2	NOP
EX	NOP	NOP	NOP	Instr1	Instr2
Clock Cycles	n	n+1	n+2	n+3	n+4

Fig. 2. Pipeline states during the execution of instructions for measuring inter-instruction costs

Consequently, it is:

$$E_M(Instr1, Instr2) = E_{cycle_n+1} + E_{cycle_n+2} + E_{cycle_n+3} + E_{cycle_n+4} \quad (9)$$

which corresponds to:

$$E_M(Instr1, Instr2) = E_{Instr1} + E_{Instr2} + 2E_{NOP} + E_{NOP, Instr1} + E_{Instr1, Instr2} + E_{Instr2, NOP} + E_{NOP, NOP} \quad (10)$$

while the inter-instruction effect is figured as:

$$E(Instr1, Instr2) = E_M(Instr1, Instr2) - E(Instr1) - E(Instr2) - 2E_{NOP} \quad (11)$$

and corresponds to:

$$E(Instr1, Instr2) = E_{Instr1, Instr2} - E_{Instr1, NOP} - E_{NOP, Instr2} - E_{NOP, NOP} \quad (12)$$

2.4 Modeling the Energy of a Program

The total energy consumed when a program is executed can be estimated by summing the base costs of the instructions and the inter-instruction costs. Doing that, it is observed that the inter-cycle costs included in the expressions (7) and (12) cancel each other and finally a sum including only the actual energy components of each instruction and the corresponding inter-instruction effects appears. Indeed, only three additional components appear, one inter-cycle component of the first instruction of the program and one of the last instruction and an inter-instruction effect between NOP instructions. These three components will determine the maximum theoretic error of the method. An example is given in Figure 3 where the application of the method for a program consisting of three instructions ADD, MOV and CMP is analyzed. As it is shown, all the inter-cycle costs included in equations (7) and (12) are eliminated except the three ones mentioned before. In this way the actual energy components of the instructions are isolated from secondary effects and high accurate estimates for the processor can be created.

$$\begin{aligned}
E(ADD) &= E_{ADD} + E_{NOP,ADD} + E_{ADD,NOP} + E_{NOP,NOP} \\
E(MOV) &= E_{MOV} + E_{NOP,MOV} + E_{MOV,NOP} + E_{NOP,NOP} \\
E(CMP) &= E_{CMP} + E_{NOP,CMP} + E_{CMP,NOP} + E_{NOP,NOP} \\
E(ADD,MOV) &= E_{ADD,MOV} - E_{ADD,NOP} - E_{NOP,MOV} - E_{NOP,NOP} \\
+ E(MOV,CMP) &= E_{MOV,CMP} - E_{MOV,NOP} - E_{NOP,CMP} - E_{NOP,NOP}
\end{aligned}$$

$$\begin{aligned}
E(\text{Program}) &= E_{ADD} + E_{MOV} + E_{CMP} + E_{ADD,MOV} + E_{MOV,CMP} \\
&\quad + E_{NOP,ADD} + E_{CMP,NOP} + E_{NOP,NOP}
\end{aligned}$$

Fig. 3. Estimation of the energy of a program consisting by three instructions

It has to be mentioned that a hypothesis was used so that the inter-cycle costs can cancel each other. The inter-cycle effect as shown in the expressions of the base cost and the inter-instruction cost is not the same since in these two cases the pipeline stages do not necessary correspond to the same state. However, according to our results this approximation does not influence significantly the accuracy of the method.

3 Accurate Instruction-Level Energy Modeling

As it was mentioned above, the energy consumed during the execution of instructions can be distinguished in two amounts. The base cost, energy amount needed for the execution of the operations which are imposed by the instructions, and the inter-instruction cost which corresponds to an energy overhead due to the changes in the state of the processor provoked by the successive execution of different instructions. Measurements for determining these two energy amounts for each instruction of the ARM7TDMI processor were taken and presented in [11]. However the base costs in [11] were for specific operand and address values (zero operand and immediate values and specific address values to minimize the effect of 1s). This base cost is called *pure* base cost.

We have observed in our measurements that there is a dependency of the energy consumption of the instructions on the values of their parameters (operand values, addresses). To create accurate models this dependency has to be determined. Consequently the dependency of the energy of the instructions on the register numbers, register values, immediate values, operand values, operand addresses and fetch addresses was studied. Additional measurements were taken to satisfy this necessity. It was observed by the measurements that there is a close to linear dependency of the energy on these parameters, versus the number of 1s in their word space. In this way the effect of any of the above *energy-sensitive factors* was efficiently modeled by a coefficient. Incorporating these effects in our models the proposed method keeps its promised accuracy while it becomes very attractive since it can be easily implemented in software as an estimation tool.

Making some appropriate experiments we observed that the effect of each energy-sensitive factor on the energy cost of the instruction is independent of the effect of the

other factors. The effects on the energy of these factors are uncorrelated as can be observed in Table 1. The distortion of our results from this conclusion is, most of the time, less than 2-3% and only in some marginal cases becomes more than 7%. According to this conclusion, the effect of the energy-sensitive factors can simply be added to give the total energy amount.

Table 1. Comparison of measured to calculated instruction energy costs (nJoules) due to additional dependencies

Instruction formation	Measured		Calculated		% diff
	$E_{base} + E_{opnd}$	$E_{base} + E_{opnd} + E_{regm}$	$E_{base} + E_{opnd}$	$E_{base} + E_{opnd} + E_{regm}$	
ADD Rd,Rn,Rs,ASR Rm	2.58	2.61	2.57	2.61	-0.04
ADDRd,Rn,Rs,ASR imm	1.57	1.60	1.55	1.60	-0.27
ADD Rd,Rn,Rs	1.51	1.59	1.51	1.56	1.71
ADD Rd,Rn,Rs,RRX	1.51	1.63	1.52	1.64	-1.19
LDR Rd,[Rn,Rs]	3.07	3.27	2.97	3.29	-0.63
STR Rd,[Rn,Rs]	2.28	2.48	2.23	2.43	2.01

Other sources of energy consumption are conditions of the processor which lead to an overhead in clock cycles because of the appearance of idle cycles. This is the case of the appearance of pipeline stalls, which was measured and modeled.

According to the above, the energy, E_i , consumed during the execution of the i instruction can be modeled as:

$$E_i = b_i + \sum_j a_{i,j} N_{i,j} \tag{13}$$

where b_i is the pure base cost of the i instruction, $a_{i,j}$ and $N_{i,j}$ is the coefficient and the number of 1s of the j energy-sensitive factor of the i instruction, respectively.

Each instruction may contain some of these energy-sensitive factors. The effect of all the factors included in an instruction have to be taken into account to create the energy budget of the instruction.

Having modeled the energy cost of the instructions, the energy consumed for running a program of n instructions can be estimated as:

$$PE = \sum_1^n E_i + \sum_1^{n-1} O_{i,i+1} + \sum \epsilon \tag{14}$$

where $O_{i,j}$ is the inter-instruction cost of the instructions i and j , and ϵ is the cost of a pipeline stall.

4 Pure Base Cost and Inter-instruction Cost Models – Results

The complete models for the instruction-level energy consumption of the ARM7TDMI created according to the proposed methodology can be found in [12]. Pure base costs of all the instructions and for all the addressing modes are given. Since the number of the possible instruction pairs (taking into account the addressing modes) is enormous, groups of instructions and groups of addressing modes according to the resources they utilize, have been formed and inter-instruction costs have been given only for representatives from these groups. In this way we keep the size of

the required model values reasonable without significant degradation of the accuracy (error less than 5% in the inter-instruction cost by using representative instructions).

Most of the values of the inter-instruction costs have negative sign as it was expected. The contribution of the inter-instruction costs, as they are calculated according to the proposed method, remains small. As it can be observed by our models most of the inter-instruction costs are less than 5% of the corresponding pure base costs while almost all the cases are covered by a 15% percentage. Also, there is no symmetry in the inter-instruction cost for a pair of instructions. For example, the execution of the LDR after the ADD present a cost of 0.064nJoule while the execution of ADD after LDR present an cost of -0.122 nJoule. This seems more reasonable from the case of symmetric inter-instruction costs as Tiwary method supposes.

To determine the accuracy of the method a number of programs with various instructions have been created. In these instructions the operand values were kept zero and thus the effect of energy sensitive factors wasn't taken into account. The energy consumed during the execution of each program was calculated directly from measurements of the instantaneous drawn current and also calculated by using the derived instruction-level energy models. The error was found to be up to 1.5%.

5 Models of the Energy Sensitive Factors Effect – Results

The dependency on the energy of the energy sensitive factors was also studied. Energy depends on the number of 1s in the word structures of these entities. The Hamming distance between the corresponding word fields of successive instructions is not considered here since according to the followed modeling methodology, where NOP is used as a reference instruction, Hamming distance equals to the number of 1s of the corresponding fields of the instructions.

The effect of each factor was studied separately from the others since, as it can be verified by the results, the correlation among the effect of these factors is insignificant. This energy dependency can be approximated with sufficient accuracy by linear functions. Coefficients were derived for each instruction for any energy sensitive factor. However, appropriate grouping of the instructions is used to keep reasonable the number of required coefficients to increase the applicability of the method without significant loss in the accuracy.

The grouping of the instructions for the derivation of the coefficients and the corresponding measurements is presented in [12]. Some results are given here. In Figure 4 the effect of the register number for data-processing instructions in register addressing mode is presented. Actual physical measurements were compared with the estimated energy values, as they are produced using the selected coefficients, for various instructions. The error in most cases was less than 3%. The energy consumption versus the number of 1s in operand values for the STR instruction is shown in Figure 5. The linear dependency is obvious in both figures.

To evaluate the absolute accuracy of our modeling approach, real kernels were used as benchmarks. The corresponding assembly list have been extracted from C programs by utilizing the facilities of the *armcc* tool, shipped with the ARM ADS software distribution. The energy consumption during the cycles of the program was calculated by direct measurement of the drawn current and by using the derived mod-

els and equation (14). According to our results the error of our approach in real life programs was found to be less than 5%.

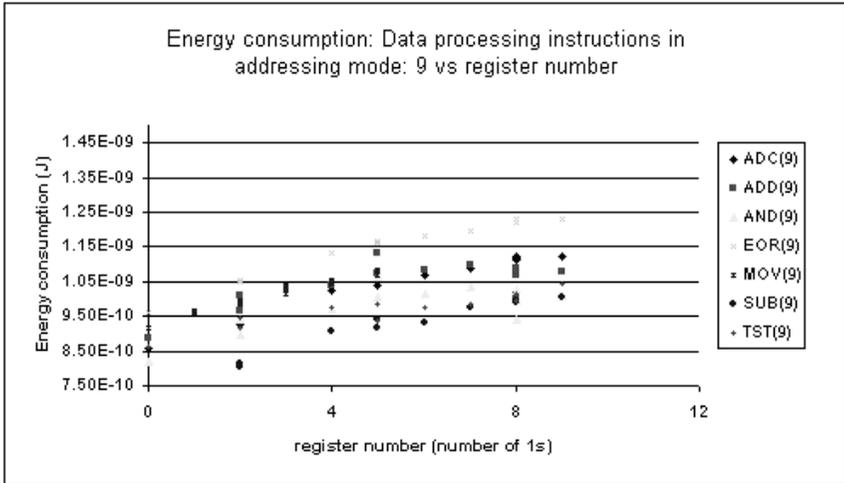


Fig. 4. The effect of register number for data-processing instructions in register addressing mode

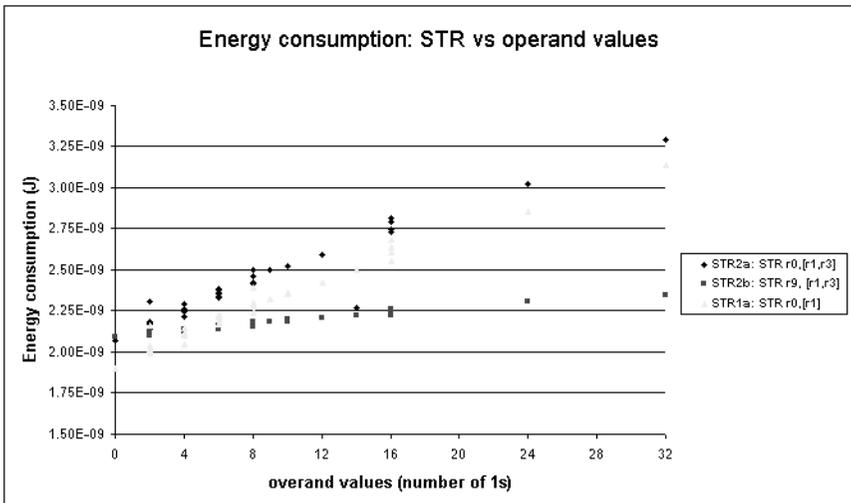


Fig. 5. Energy consumption versus the number of 1s in operand values for the STR instruction

6 Conclusions

High quality instruction level energy models can be derived for pipelined processors by monitoring the instantaneous current drawn by the processor at each clock cycle.

Knowing the current waveform the corresponding energy consumed at a clock cycle can be calculated. However, energy components of an instruction exist in different clock cycles. An efficient method is proposed to extract the actual energy components, which are used for the calculation of the energy consumption of complete software programs. The instruction base costs and inter-instruction costs are modeled in relation to a reference instruction. Also, the dependency of the instruction energy on the instruction parameters (e.g. register numbers, register values, operand values) were studied. The proposed method has been applied to the ARM7TDMI processor providing an error in estimating of energy consumption of real software kernels less than 5%.

References

1. Tiwari, V., Malik, S., Wolfe, A., "Power Analysis of Embedded Software: A First Step Towards Software Power Minimization," IEEE Transaction on Very Large Scale Integration (VLSI) Systems, Vol. 2, No. 4, (1994) 437–445
2. Tiwari, V., Malik, S., Wolfe, A., Lee M.T-C. "Instruction Level Power Analysis and Optimization of Software," Journal of VLSI Signal Processing, Vol. 13, No. 2-3, (1996) 223–238
3. Lee M.T-C, Tiwari, V., Malik, S., Fujita, M., "Power Analysis and Minimization Tehniques for Embedded DSP Software," IEEE Transactions on Very Large Scale Integration (VLSI) Systems, (1997) 123–135
4. Tiwari, V., Lee T.C., "Power Analysis of a 32-bit Embedded Microcontroller," VLSI Design Journal, Vol. 7, No. 3, (1998)
5. SOFLOPO, "Low Power Development for Embedded Applications", ESPRIT project: Deliverable 2.2: Physical measurements. By Thanos Stouraitis, University of Patras, Dec 1998
6. Steinke, S., Knauer, M., Wehmeyer, L., Marwedel, P., "An Accurate and Fine Grain Instruction-Level Energy Model supporting Software Optimizations," Int. Workshop on Power and Timing Modeling, Optimization and Simulation, Yverdon-les-bains, Switzerland (2001)
7. Russell, J.T., Jacome, M.F., "Software Power Estimation and Optimization for High Performance, 32-bit Embedded Processors," Int. Conf. On Computer Design (1998)
8. Chang, N., Kim, K., Lee, H-G. "Cycle-Accurate Energy Consumption Measurement and Analysis: Case Study of ARM7TDMI," IEEE Transactions on VLSI Systems, Vol. 10, No. 2, (2002) 146–154
9. Lee, S., Ermedahl, A., Min S-L, Chang N., "An Accurate Instruction-Level Energy Consumption Model for Embedded RISC Processors," ACM SIGPLAN Workshop on Languages, Compilers and Tools for Embedded Systems (2001)
10. Nikolaidis, S., Kavvadias, N., Neofotistos, P., Kosmatopoulos, K., Laopoulos, T., Bisdounis, L., "Instrumentation set-up for Instruction Level Power Modeling," Int. Workshop on Power and Timing Modeling, Optimization and Simulation, Seville, Spain, Sept. 2002
11. S. Nikolaidis, N. Kavvadias, P. Neofotistos, "Instruction level power measurements and analysis", IST-2000-30093/EASY Project, Deliverable 15, Sept 2002.
<http://easy.intranet.gr>
12. S. Nikolaidis, N. Kavvadias, P. Neofotistos, "Instruction level power models for embedded processors", IST-2000-30093/EASY Project, Deliv. 21, Dec 2002.
<http://easy.intranet.gr>