

An Instruction-Level Methodology for Power Estimation and Optimization of Embedded VLIW Cores

A. Bona[§] M. Sami[‡] D. Sciuto[‡] C. Silvano[‡] V. Zaccaria[‡] R. Zafalon[†]

[§]ALaRI, Lugano, Switzerland, Italy

[‡]Politecnico di Milano, Milano, Italy

[‡]Università degli Studi di Milano, Milano, Italy

[†]STMicroelectronics, Agrate B. (MI), Italy

The overall goal of this work is to define an instruction-level power macro-modeling and characterization methodology for VLIW embedded processor cores. The approach presented in this paper is a major extension of the work previously proposed in [1-3], targeting an instruction-level energy model to evaluate the energy consumption associated with a program execution on a pipelined VLIW core.

Our *first* goal is the reduction of the complexity of the processor's energy model, without reducing the accuracy of the results. The *second* goal is to show how the energy model can be further simplified by introducing a methodology to automatically cluster the whole Instruction Set with respect to their average energy cost, in order to converge to an highly effective design of experiments for the actual characterization task. The paper describes also the application of the proposed model to a real industrial VLIW core (the Lx Architecture developed by HP Labs and STMicroelectronics), to validate the effectiveness and accuracy of the proposed methodology.

Given a sequence \mathcal{W} of instructions w_n , the model is summarized by the following equation:

$$E(\mathcal{W}) \approx \sum_{1 \leq n \leq N} \left[U(\mathbf{0}|\mathbf{0}) + \sum_{\forall k \in K} \nu(w_n^k | w_{n-1}^k) + m * p * S + l * q * M \right] \quad (1)$$

Where the term $U(\mathbf{0}|\mathbf{0})$ is the base energy cost that represents the energy consumed during an execution of a long instruction constituted entirely by NOPs ($\mathbf{0} = [\text{NOP} \dots \text{NOP}]^T$), $\nu(w_n^k | w_{n-1}^k)$ is the additional energy contribution due to the change of operation on the same lane k , $m(l)$ is the average data (instruction) cache miss length, $p(q)$ is the average probability per stage and per instruction that a data (instruction) cache miss can affect one instruction and S (M) is the average energy consumption of the *entire* processor during these events.

For a k -issue VLIW core, the complexity of the model is now quadratic with respect to the number of operations within the instruction set ($O(K * |ISA|^2)$), while a black-box model has a complexity of $O(|ISA|^K)$. This leads to a major reduction of the characterization effort, since the number of experiments to be done is reduced exponentially.

This model can be further simplified by clustering the operations of the ISA with respect to their average energy consumption to allow a faster and more effective characterization of the core's power consumption, while preserving

the model accuracy. The basic idea of the cluster analysis consists of grouping in the same cluster the operations with a power cost close to each other. The power cost of an operation is defined as the power consumed by the processor when it executes only that operation. Among the various clustering algorithms appeared in literature we have chosen the *k-mean clustering algorithm* to cluster energy values since it requires a lower computational cost with respect to hierarchical clustering.

Once the instructions have been clustered into a set of $C_1 \dots C_j$ clusters, we compute the matrix ν in equation (1) as:

$$\nu(w_n^k | w_{n-1}^k) = \begin{cases} E_i & \text{if } w_n^k = w_{n-1}^k \wedge w_n^k, w_{n-1}^k \in C_i \\ D_{i,j} & \text{if } w_n^k \neq w_{n-1}^k \wedge w_n^k \in C_i, w_{n-1}^k \in C_j \end{cases} \quad (2)$$

The complexity of matrix ν has been reduced from $O(ISA^2)$ to $O(C^2)$, where C is the number of clusters.

The model has been tested against a set of validation benchmarks on the Lx architecture. The validation benchmarks include a set of the Mediabench applications (namely, g721 encoder and decoder, epic encoder and decoder, mpeg2), a set of finite impulse response filters, discrete cosine transforms and matrix elaboration algorithms. The power model has shown, on the validation benchmarks, a mean error of 1.9% and a standard deviation on the error of 5.8%.

Our ongoing work aims at defining a power optimization technique based on the proposed model. The technique consists of a *spatial rescheduling* of the operations within the same long instruction to reduce their instruction power overhead.

References

- [1] L. Benini, D. Bruni, M. Chinosi, C. Silvano, V. Zaccaria, and R. Zafalon, "A power modeling and estimation framework for vliw-based embedded systems," in *Proceedings of International Workshop-Power And Timing Modeling, Optimization and Simulation, PATMOS'01* 26-28 2001, pp. xx-xx.
- [2] M. Sami, D. Sciuto, C. Silvano, and V. Zaccaria, "Instruction level power estimation for embedded VLIW cores," in *Proceedings of the 8th International Workshop on Hardware/Software Co design (CODES-00)* NY, May 3-5 2000, pp. 34-38, ACM.
- [3] M. Sami, D. Sciuto, C. Silvano, and V. Zaccaria, "Power exploration for embedded VLIW architectures," in *Proceedings of the IEEE/ACM International Conference on Computer Aided Design (ICCAD-2000)*. Nov. 5-9, pp. 498-503, IEEE Computer Society Press.