

Architectures Session

Trade-off Analysis of L2 on-chip Cache Architectures for Embedded MPSoCs

M. Sabry¹, M. Ruggiero², D. Atienza¹, and L. Benini²

¹*ESL, EPFL, Lausanne, Switzerland, {mohamed.sabry,david.atienza}@epfl.ch*

²*DEIS, Bologna University, Bologna, Italy, {martino.ruggiero,luca.benini}@unibo.it*

Abstract

On-chip memory organization is one of the most important aspects that can influence the overall system behavior in Multi-processor System-On-Chips. The on-chip memory must provide the required bandwidth to the software tasks, complying with tight constraints in terms of chip area, power consumption and cost. Following the trend set by high-performance processors, high-end embedded cores are moving from single-level on chip caches to a two-level on-chip cache hierarchy. Even if there is general consensus on L1 private caches, for L2 there is still not a dominant architectural paradigm.

Using the same architectural approach for these two domains (hi performance and embedded systems) could not be the best solution, since they are subject to very different constraints. Embedded MPSoC system architecture has always been quite different from the general-purpose one. It usually uses simpler processor microarchitectures, well-matched to the characteristics of the targeted applications. Simplicity essentially derives from precise knowledge of the specific application focus and from the several constraints that are present in this domain.

This work presents a virtual platform for design space exploration focused on in-depth analysis of L2 cache architectures for MPSoCs in terms of area, power, system performance and predictability trade-offs. We developed a set of parameterizable models for L2 caches and integrated them in an accurate virtual platform environment specifically designed for Embedded MPSoC design space explorations. Our enhanced virtual platform is highly modular and capable of simulating at cycle-accurate level an entire MPSoC, including cores, L1 and L2 caches, L3 memories and system buses. We instantiated several L2 cache architectural templates: private, shared, or hybrid. We integrated a cache power, timing, and area estimation tool; CACTI. We also imported an embedded benchmarking suite; MiBench.

Benchmarking results show that the selection of the optimal one is not trivial. The private L2 cache organization performs better than the shared one. However, the hybrid L2 cache achieves better execution time and energy consumption in presence of private and shared data (such as communicating tasks allocated on different cores), but at the cost of a larger area.