Introduction

Per Stenström[1] and David Whalley[2]

[1] Chalmers University of Technology, Sweden
[2] Florida State University, U.S.A.

In January 2007, the second edition in the series of International Conferences on High-Performance Embedded Architectures and Compilers (HiPEAC'2007) was held in Ghent, Belgium. We were fortunate to attract around 70 submissions of which only 19 were selected for presentation. Among these, we asked the authors of the five most highly rated contributions to make extended versions of them. They all accepted to do that and their articles appear in this section of the second volume.

The first article by Keramidas, Xekalakis, and Kaxiras focuses on the increased power consumption in set-associative caches. They present a novel approach to reduce dynamic power that leverages on the previously proposed cache decay approach that has been shown to reduce static (or leakage) power.

In the second article by Magarajan, Gupta, and Krishnaswamy the focus is on techniques to encrypt data in memory to preserve data integrity. The problem with previous techniques is that the decryption latency ends up on the critical memory access path. Especially in embedded processors, caches are small and it is difficult to hide the decryption latency. The authors propose a compiler-based strategy that manages to reduce the impact of the decryption time significantly.

The third article by Kluyskens and Eeckhout focuses on detailed architectural simulation techniques. It is well-known that they are inefficient and a remedy to the problem is to use sampling. When using sampling, one has to warm up memory structures such as caches and branch predictors. This paper introduces a novel technique called Branch History Matching for efficient warmup of branch predictors.

The fourth article by Bhadauria, McKee, Singh, and Tyson focuses on static power consumption in large caches. They introduce a reuse-distance drowsy cache mechanism that is simple as well as effective in reducing the static power in caches.

Finally, in the fifth paper by Hu, Jimenez, and Kremer, the focus is on a methodology to make accurate and fast estimations of the power consumption in a program. The authors note that while detailed power simulation is slow and inaccurate, real power measurements result in huge amounts of data. To this end, they present an infrastructure that can identify the program phases that best characterize the power consumption profile in the program.

We do hope that you learn a lot and get inspiration by this excellent excerpt from the second HiPEAC conference.

P. Stenström (Ed.): Transactions on HiPEAC II, LNCS 5470, p. 3, 2009.