Architectural Support for Reducing Parallel Processing Overhead in an Embedded Multiprocessor

In Proceedings of the 2010 IEEE/IFIP International Conference on Embedded and Ubiquitous Computing (EUC '10). IEEE Computer Society, Washington, DC, USA, 47-52. DOI=10.1109/EUC.2010.17 http://dx.doi.org/10.1109/EUC.2010.17

Jian Wang, Joar Sohl, and Dake Liu
Division of Computer Engineering
Department of Electrical Engineering

Abstract

The host-multi-SIMD chip multiprocessor (CMP) architecture has been proved to be an efficient architecture for high performance signal processing which explores both task level parallelism by multi-core processing and data level parallelism by SIMD processors. Different from the cache-based memory subsystem in most general purpose processors, this architecture uses on-chip scratchpad memory (SPM) as processor local data buffer and allows software to explicitly control the data movements in the memory hierarchy. This SPM-based solution is more efficient for predictable signal processing in embedded systems where data access patterns are known at design time. The predictable performance is especially important for real time signal processing. According to Amdahl¡¯s law, the nonparallelizable part of an algorithm has critical impact on the overall performance. Implementing an algorithm in a parallel platform usually produces control and communication overhead which is not parallelizable. This paper presents the architectural support in an embedded multiprocessor platform to maximally reduce the parallel processing overhead. The effectiveness of these architecture designs in boosting parallel performance is evaluated by an implementation example of 64x64 complex matrix multiplication. The result shows that the parallel processing overhead is reduced from 369% to 28%.