

Data Prefetching via Off-line Learning

Wong Weng Fai

Abstract—The widely acknowledged performance gap between processors and memory has been the subject of much research. In the Explicitly Parallel Instruction Computing (EPIC) paradigm, the combination of in-order issue and the presence of a large number of parallel function units has further worsen the problem. Prefetching, by hardware, software or a combination of both, has been one of the primary mechanisms to alleviate this problem. In this talk, we will discuss two prefetching mechanisms, one hardware and other software, suitable for implementation in EPIC processors. Both methods rely on the off-line learning of Markovian predictors. In the hardware mechanism, the predictors are loaded into a table that is used by a prefetch engine. We have shown that the method is particularly effective for prefetching into the L2 cache. Our software mechanism which we called predicated prefetch leverages on informing loads. This is used in conjunction with data remapping and offline learning of Markovian predictors. This distinguishes our approach from early software prefetching techniques that only involves static program analysis. Our experiments show that this framework, together with the algorithms used in it, can effectively remove, in the best instance, 30% of the stall cycles due to cache misses. The results also show that the framework performs better than pure hardware stride predictors and has lower bandwidth and instruction overheads than that of pure software approaches.

[Full Text Not Available]