

Impact of Cache Mapping on Memory Performance of a StrongArm Processor

Zhiyuan Li and Rong Xu

Abstract—Embedded processors such as the Intel StrongARM SA-1110 and the Intel XScale utilize multiple caches at the same level in the cache hierarchy. The main cache and the mini-cache differ in both the size and the associativity. Furthermore, the processors allow programs to specify the cache mapping policy for each virtual page among three choices, i.e. whether to map the page to the main cache, the mini-cache, or neither. In the latter case, the page is marked as noncacheable. In this paper, we investigate the problem of optimal cache mapping, assuming that we can predict the trace of the memory reference in advance. On the theoretical side, we prove that the problem of finding the optimal cache mapping for an arbitrary memory trace is NP-hard. On the experimental side, we present a mapping heuristic and compare the result with the default policy which maps all pages to the main cache. Our measurement shows that, compared to the default policy, the heuristic can reduce the execution time from 1% to 21% for a set of test programs. As a byproduct of performance enhancement, we also save the energy by 4% to 28%.

Index Terms—Handheld devices, horizontally partitioned caches, cache bypass, NP-hardness, performance measurement

ACKNOWLEDGMENTS

This work is sponsored by National Science Foundation through grants CCR-0208760, ACI/ITR-0082834, and CCR-9975309. An earlier draft of the paper, titled “Using Cache Mapping to Improve Memory Performance of Handheld Devices”, was presented to the 4th IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS-2004), March 10 -12, 2004, Austin, Texas, USA. This archival version gives the complete NP-hardness proof for the CACHE-MAPPING problem.