

José Carlos Cabaleiro, Francisco Fernández Rivera

Title: Increasing data reuse of sparse algebra codes on simultaneous multithreading architectures

Abstract:

In this presentation the problem of the locality of sparse algebra codes on simultaneous multithreading architectures is studied. In this kind of architectures many hardware structures are dynamically shared among the running threads. This puts a lot of stress on the memory hierarchy, and a poor locality, both inter-thread and intra-thread, may become a major bottleneck in the performance of a code. This behavior is even more pronounced when the code is irregular, which is the case of sparse matrix ones. Therefore, techniques that increase the locality of irregular codes on simultaneous multithreading architectures are important to achieve high performance. Our proposal is a data reordering technique specially tuned for this kind of architectures and codes. It is based on a locality model developed in previous works by the authors.

The technique has been tested, first, using a simulator of a simultaneous multithreading architecture, and later, on a real architecture as Intel's Hyper-Threading. Important reductions in the number of cache misses have been achieved, even when the number of running threads grows. On applying the locality improvement technique we also decrease the total execution time and improve the scalability of the code.

Affiliation: Universidad de Santiago de Compostela