

## **A Data Cache with Dynamic Mapping**

Paolo D'Alberto, Alexandru Nicolau and Alexander Veidenbaum

To appear at the *16th Workshop on Languages and Compilers for Parallel Computing (LCPC03)*, College Station, TX, 2-4 October 2003

### **Abstract**

Dynamic Mapping is an approach to cope with a loss of performance due to cache interference and to improve performance predictability of blocked algorithms for modern architectures. An example is matrix multiply: tiling matrix multiply for a data cache of 16KB using optimal tiles size achieves an average data-cache miss rate of 3%, but with peaks of 16% due to interference.

Dynamic Mapping is a software-hardware approach for which the mapping in cache is determined at compile time, by manipulating the address used by the data cache. The reduction in the cache misses translates into a 2-fold speed-up for matrix multiply and FFT by eliminating data-cache miss spikes.

Dynamic mapping has the same goal as other proposed approaches, but it determines the cache mapping before issuing a load. It uses the computational power of the processor - instead of the memory controller or the data cache mapping - and it has no effect on the access time in memory and cache. It is an approach combining several concepts, such as non-standard cache mapping functions and data layout reorganization and, potentially, without any overhead.