Combining Building Blocks for Parallel Multilevel Matrix Multiplication

Sascha Hunold and Thomas Rauber
Department of Mathematics and Physics
University of Bayreuth, Germany
{hunold,rauber}@uni-bayreuth.de

Gudula Rünger
Department of Computer Science
Chemnitz University of Technology, Germany
ruenger@informatik.tu-chemnitz.de

Keywords: Strassen's algorithm, matrix multiplication, multiprocessor tasks, task parallelism

Extended Abstract

Matrix-matrix multiplication is one of the core computations in many algorithms from scientific computing or numerical analysis and many efficient realizations have been invented over the years, including many parallel ones. The current trend to use clusters of PCs or SMPs for scientific computing suggests to revisit matrix-matrix multiplication and investigate efficiency and scalability of different versions on clusters. In this talk we present parallel algorithms for matrix-matrix multiplication which are built up from several algorithms in a multilevel structure.

On a single processor ATLAS or PHiPAC create very efficient implementations by adjusting the computation order to the specific memory hierarchy and by exploiting functional parallelism of the processor (SSE2). Parallel approaches include many methods based on decomposition like Cannon's algorithm, or the algorithm of Fox. Efficient implementation variants of the latter are SUMMA or PUMMA. Matrix-matrix multiplication by Strassen or Strassen-Winograd benefits from a reduced number of operations but require a special schedule for a parallel implementation.

In the context of clusters of SMPs, mixed programming models like mixed task and data parallelism are important since efficiency and scalability can be improved by using multiprocessor tasks (M-tasks). Task parallel implementations or mixed matrix-matrix multiplications have already been proposed in literature. One possibility of parallelizing Strassen's algorithm is to distribute the seven intermediate results onto a group of processors of size 7^i, preferably in a ring or torus configuration. Other approaches include to mix the common Fox BMR method (broadcast multiply roll) with Strassen's algorithm. Another mixed parallel algorithm that exploits the complexity reduction of Strassen's algorithm on the top level and combines it with the performance of ScaLAPACK on the bottom layer.

In this talk we consider matrix-matrix multiplication on clusters of PCs or SMPs and investigate the performance of several parallel hierarchical algorithms and their realizations. In particular, we investigate new multilevel algorithms with different building blocks, including well-known parallel algorithms like Strassen multiplication as well as new algorithms that have been designed to exploit the memory hierarchy of recent microprocessors by an increased locality of memory references. We show that a suitable combination of the building blocks can lead to significant performance improvements compared with an execution of the algorithms in isolation. The building blocks are expressed as M-tasks to exploit mixed task and data parallelism. The combination of the M-tasks is performed with the Tlib library.

The composed methods use the Strassen algorithm on the upper level to create coarse-grained M-tasks that are assigned to disjoint processor groups. For the intermediate level efficient variants like PDGEMM from ScaLAPACK and a new hierarchical decomposition-based algorithm tpMM are used. For the lowest level we use BLAS or ATLAS for the multiplication of smaller blocks on single processors. It is a crucial point to choose a good schedule for Strassen and to pick the right cutoff of the hierarchical decomposition for the coupling of the different levels. Depending on the cluster platform or parallel machine different strategies lead to the most efficient implementation. We show that a suitable combination of different algorithms that are used as building blocks can be used to obtain multilevel algorithms which lead to significant performance improvements compared to an execution of the algorithms in isolation. In particular, it is possible to obtain algorithms that are nearly twice as fast on some cluster platforms as the original PDGEMM method from ScaLAPACK. An important point for the construction of the multilevel algorithms lies in the issue, which building blocks should be combined in which order and at which group size the building blocks are assembled. The evaluation shows that a combination of Strassen's method at the top level with special communication-optimized algorithms on the intermediate level and ATLAS or DGEMM at the bottom level leads to the best results.