Efficient Communication for the Distributed Matrix Assembly towards Massive Parallel HPC

R.-P. Mundani[1] **and E. Rank**[2]

Abstract: Finite element methods (FEM) are frequently used within a plenitude of various applications from different disciplines such as science and engineering. Typically both time and memory consuming, FEM is a predestinated candidate for high-performance computing as it specifically benefits from parallelisation. While parts of FEM – such as computing element stiffness matrices – are easily deployed on parallel architectures without huge effort, for others, the system matrix assembly for instance, it is already difficult to leverage compute clusters with more than 10,000 cores.

In particular, the parallel matrix assembly is to be solved efficiently on shared-memory architectures ([1], e.g.), but those parallelisation strategies – sufficient hardware supposed – neither scale nor can be transferred to distributed memory architectures for being executed on some thousand cores. Therefore, different approaches are necessary which typically entail high communication advent. To tackle the communication problem, we proposed a strategy which is related to the idea of bitonic sorting networks for the inter processes communication. Bitonic sorting networks are known to sort an input sequence of N elements in $\Theta(\text{ld}N)$ time (see [2] for more details). Due to the usage of a bitonic data flow pattern communication can be organised in a very efficient way, allowing a distributed matrix assembly that shows excellent scalability behaviour.

In the upcoming paper, we will address the advantages but also shortcomings of this approach, its parallel implementation using MPI, and its benefits in order to reduce the communication complexity. We will further present measurements carried out on the Shaheen supercomputer at King Abdullah University of Science and Technology (KAUST), Saudi Arabia. Even this communication strategy is only a small piece of the big puzzle, it is nevertheless an important step into the right direction of being ready for the exascale challenge.

References

[1] M. N. de Rezende and J. B. de Paiva: *A parallel algorithm for stiffness matrix assembly in a shared memory environment*, Computers and Structures 76 (5), 2000, pp. 593–602.

[2] K. E. Batcher: *Sorting networks and their applications*, Proceedings of the AFIPS Spring Joint Computer Conference, Thomas Book Company, 1968, pp. 307–314.

[1,2] Chair for Computation in Engineering
Faculty for Civil Engineering and Surveying
Technische Universität München
Arcisstrasse 21, 80333 München, Germany
mundani@tum.de, rank@bv.tum.de