A Distributed and Parallel Asynchronous Unite and Conquer Method to Solve Large Scale Non-Hermitian Linear Systems

Serge G. Petiton [1]

Parallel preconditioned Krylov subspace methods are commonly used for solving large sparse linear systems. Facing the development of extreme scale platforms, the minimization of synchronous global communication becomes very critical to obtain good efficiency and scalability. We highlight the recent development of the Unite and Conquer GMRES/Least Squares-ERAM (UCGLE) method, a distributed and parallel solver for large scale non-Hermitian linear systems, which aims to exploit the modern computer architectures.

UCGLE combines the distributed and parallel algorithms to optimize the communication and accelerate the convergence using a least squares polynomial preconditioner. The sequential version of this preconditioner was presented by Yousef Saad in 1987. UCGLE composes three computation components: 1) GMRES Component; 2) Least Squares Preconditioning Component; 3) ERAM Component, implemented with asynchronous communication. Thus the preconditioning part is independent from the solver part. UCGLE minimizes the residuals, using the eigenvalues computed asynchronously in parallel. These approximated eigenvalues can be reused, and improved during the resolution of different systems, to enhance the reusability and fault tolerance. UCGLE is well-suited to complex heterogeneous platforms, which puts the emphasis on multi-grain, multi-level memory, reducing synchronization and promoting asynchronization.

UCGLE was implemented based on PETSc and SLEPc, adding the basic sending and receiving functions to enable the asynchronous communication. This allows us to make a fair comparison between UCGLE and other preconditioners provided by PETSc.

Using a software engine, which distributes parallel implementations of Krylov methods on a cluster of hundreds computing nodes with accelerators GPUs, we experiment our method with generated large-scale sparse matrices, especially some with eigenvalues known, compare it with basic preconditioners, evaluate the importance of parameters such as the preconditioner polynomial degree, and required eigenvalues number, also the influence of the spectrum distribution in complex plane. The reusability is discussed by resolving a sequence of linear systems with a same matrix but different right-hand sides. We also study the performance and scalability on the platform with/without GPU accelerators. We obtained better convergence and performance than using some classical preconditioners. In conclusion, we propose future improvements and research perspectives.

[1] University Lille 1, France (Serge.Petiton@univ-lille1.fr)