

Implementation of an additional parallelization level for an High Performance Computing algorithm

Bachelor's thesis, Master's thesis, Student project
 Contact: Michael Rippl, Chair of Scientific Computing, ripplm@in.tum.de

Keywords: HPC, MPI, Load Balancing, Parallel Numerics, Numerical Linear Algebra, Eigenvalue problem


Abstract

For the transformation of the Generalized Eigenvalue Problem $Ax = \lambda Bx$ to the Standard Eigenvalue Problem $Ax = \lambda x$ we work on a novel algorithm. The algorithm offers the possibility to add an additional level of parallelism and modify the matrix starting from top and bottom concurrently. The required changes have to be identified and implemented.

Global Picture

The solution of dense symmetric Eigenvalue problems is a crucial step in many simulations in science and engineering. Often solving a series of eigenvalue problems is the most expensive step in this simulations. Therefore powerful and highly scalable parallel algorithms are needed for this task.

Example: Electronic Structure Theory



- Prediction of material properties from atomic scale upwards
- Simulations are run on Supercomputers



Schrödinger equation

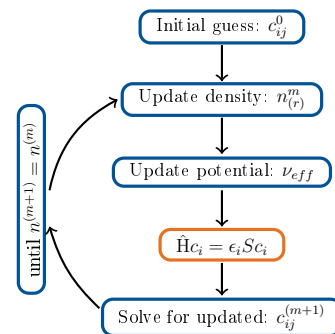
$$\hat{\mathcal{H}}\Phi_i = E_i\Phi_i$$

↓ Approx. ↓
 ↓ Discret. ↓

$$\hat{H}c_i = \epsilon_i S c_i$$

↓

Eigenproblem to solve in Iteration:



The ELPA library and the context

Together with our collaborators from other universities and several Max-Planck-institutes we are working on ELPA [1, 2, 3], a highly scalable library for solving the Eigenvalue problem for dense symmetric matrices.

To solve the Eigenproblem we follow a two step approach: Transform the full matrix to banded form, transform the banded matrix to a tridiagonal matrix and solve the Eigenproblem on the tridiagonal matrix. Additionally, two backtransformations follow to get the correct Eigenvectors.

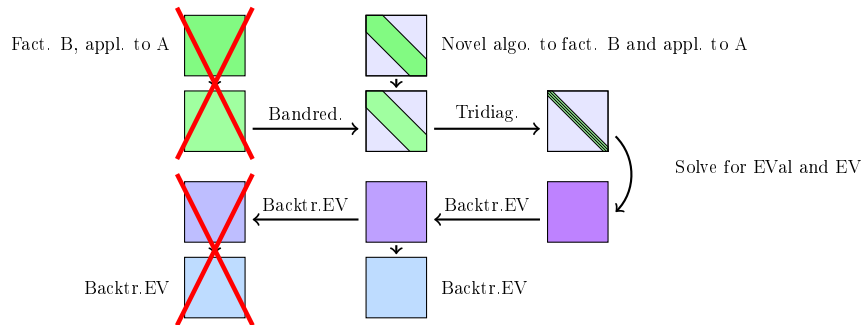


Figure 1: Algorithmic path of ELPA using the novel algorithm

However, in Electronic Structure Theory, a generalized Eigenvalue problem with banded matrices has to be solved. For this, we are working on a novel algorithm which will allow us to circumvent the expensive transformation to a banded matrix (see Figure 1).

This algorithm has the nice feature that it can be run from top and bottom of the matrix independently. Therefore, we can group the processes now in two groups: processes solely for the upper half and processes solely for the lower half of the matrix.

The data is distributed in a blockcyclic way over the processes beforehand. Thus, the data has also to be moved according to the new distribution of processes.

In the middle of the matrix both process groups meet and it has to be ensured that the lower half group works first on the data. This has to be considered when thinking about distributing the processes to the lower and upper half.

Tasks

- Identify and implement the necessary modification to the algorithm
- Develop a method for the required data transfer in the parallel data distribution
- Run performance measurements and compare the approach with the version without the additional parallelization layer

Requirements

- Good programming skills
- Basic skills in parallel programming
- Knowledge in HPC and Computer Architecture
- Ideally background in Linear Algebra

Bibliography

- [1] Elpa project website. <http://elpa-aeo.mpcdf.mpg.de/>.
- [2] Thomas Auckenthaler. *Highly scalable eigensolvers for petaflop applications*. PhD thesis, PhD thesis. Fakultät für Informatik, Technische Universität München, 2012.
- [3] Andreas Marek, Volker Blum, Rainer Johanni, Ville Havu, Bruno Lang, Thomas Auckenthaler, Alexander Heinecke, Hans-Joachim Bungartz, and Hermann Lederer. The elpa library: scalable parallel eigenvalue solutions for electronic structure theory and computational science. *Journal of Physics: Condensed Matter*, 26(21):213201, 2014.