High Performance Solution of Skew-symmetric Eigenvalue Problems with Applications in Solving the Bethe-Salpeter Eigenvalue Problem

Carolin Penke[1] Andreas Marek[2] Peter Benner[3]

A matrix $A \in \mathbb{R}^{n\times n}$ is called skew-symmetric when $A = -A^T$. We present a direct solver for computing the eigenvalues and eigenvectors of A, that is highly efficient and scalable on modern compute architectures.

The symmetric eigenvalue problem, i.e. the case $A = A^T$, has been studied in depth for many years. It lies at the core of many applications in different areas such as electronic structure computations. As the involved matrices easily become extremely large when more complex systems are investigated, parallel algorithms running on supercomputers are necessary. One optimized library that provides the necessary tools is the ELPA library. It contains highly competitive direct solvers for symmetric (and Hermitian) eigenvalue problems running on distributed memory machines such as compute clusters. The skew-symmetric case lacks the ubiquitous presence of its symmetric counterpart and has not received the same extensive treatment. We close this gap by extending the ELPA methodology to the skew-symmetric case.

The motivation is to accelerate the structure-preserving solution of the Bethe-Salpeter eigenvalue problem. The solution of this problem allows a more accurate prediction of optical properties in quantum chemical systems. The resulting matrices are large and dense and call for a parallelizable and scalable algorithm. One particular method relies on efficiently solving a skew-symmetric eigenvalue problem and can be accelerated using the newly implemented method.

[1] https://elpa.mpcdf.mpg.de/

[1]Max Planck Institute for Dynamics of Complex Technical Systems, Computational Methods in Systems and Control Theory

[2]Max Planck Computing and Data Facility

[3]Max Planck Institute for Dynamics of Complex Technical Systems, Computational Methods in Systems and Control Theory