Project Information

Project title: The ChASE library for large Hermitian eigenvalue problems

Name and institution of principal investigator: Andre Schleife, UIUC

Names and affiliations of co-PIs and collaborators: Edoardo Di Napoli, JSC, FZ Juelich

Corresponding author name and contact information (email): Andre Schleife, schleife@illinois.edu

Executive summary:

We study the modern implementation of one of the oldest subspace iteration methods: the Chebyshev Accelerated Subspace Eigensolver (ChASE), recently developed at JSC. In this project, we implement this algorithm into an existing code for the computation of excitonic optical spectra based on the Bethe-Salpeter equation. This approach relies on a large Hamiltonian and we are interested in the lowest eigenstates. ChASE is tailored to compute a fraction not exceeding 20% of the extremal spectrum of dense Hermitian eigenproblems. We then thoroughly characterize the parallel scaling and the timing of the new implementation and we apply it to computation of excitonic properties for In_2O_3 and ZnO materials.

Description of research activities and results:

Key Challenges:

Accurate predictions for the interaction of light with matter are crucial, for instance, to design novel optoelectronic or photovoltaic materials. Such an understanding also provides access to their fundamental properties, e.g. in the case of doped systems. Atomistic first-principles methods that are capable of computing absorption spectra including excitonic effects rely on solving the Bethe-Salpeter equation for the optical polarization function. While this allows computational materials scientists to predict optical properties of real materials, this approach is computationally expensive. For modern, complex oxide materials such as indium oxide resulting matrices have ranks reaching up to 400,000. Exact diagonalization of such matrices is prohibitively expensive, and oftentimes the lowest eigenvalues are of greatest interest, since those determine the lowest exciton binding energies. Efficient computational techniques, that parallelize very well and, ideally, benefit from GPU usage, are needed to deal with such large Hamiltonians.

Why it matters:

In order to predict optical properties of a given material, accurate simulations are necessary. The quantities that enter these simulations lead to a large computational cost, in order to achieve this accuracy. This crucially helps to understand existing materials and, for instance, we have shown the influence of free electrons that arise due to doping in organo-metal halide perovskites, illustrating that the line shape of their optical-absorption onset is very robust against free-

electron doping. We also have shown the influence of strain on the electron-electron interaction, that affects optical properties of strained carbon nanotubes that can be used as strain sensors.

At the same time, it is our goal to achieve materials design and, again, a high level of accuracy is necessary. To mitigate the high computational cost, we are performing this project, which makes our code faster and, thus, more broadly applicable, either to more complex materials, or a larger number of materials to be investigated.

Why Blue Waters:

There are several reasons to use Blue Waters: First, the machine has large numbers of nodes/cores, which is what we needed for our scaling tests. In addition, these large numbers of nodes/cores are also important when applying the resulting code to actual materials science questions. In addition, the Blue Waters team (in particular Victor Anisimov and Roland Haass) was very helpful in rewriting the reading routing of our code, making it orders of magnitude more efficient on parallel file systems, such as the one used on Blue Waters. Finally, we envision testing with GPUs in the future and Blue Waters is again a machine that provides us with these resources.

Accomplishments:

In this project we tested and customized ChASE, so as to facilitate the computation of the desired lowest eigenpairs of large dense eigenproblems on hybrid architectures. This enables conclusions with implications for electron-hole separation in solar cell absorbers and the overall optical properties of a material in the vicinity of the absorption onset. We successfully implemented the code and applied it to weak as well as strong-scaling test cases. The code is now very close to be used for production runs in the future and replace our much older implementation.

In strong-scalability tests, we choose one certain BSE eigenvalue problem, such that the size fit into one node. Using both the ChASE solver, and the legacy CG solver, we obtain the solution of the eigenvalue problem, and examine the change in computer time when we increase the number of computer nodes. We used an excitonic Hamiltonian with a rank of 41,252 and in Fig. 1 we show the results of the timings. The new implementation is significantly faster and also scales better. We then applied this code to ZnO and computed exciton binding energies with unprecedented accuracy.

List of publications associated with this work:

[1] J. Winkelmann, E. Di Napoli, A. Schleife, X. Zhang: Solving BSE equation on thousands of GPU cores; in preparation

Invited talk at APS March Meeting, 2019.

Plan for next year:

We plan to wrap up our ongoing JLESC collaboration with Edoardo Di Napoli at Juelich.

We also plan to simulate ultrafast electron dynamics in ceramic materials, such as diamond or SiC, that are subject to electron irradiation. The goal of this study is to compute the number of electrons that are excited out of their ground states into excited states upon impact of an energetic electron. The dependence of this quantity on the incoming electron's kinetic energy will be explored. These simulations will be carried out on hundreds of nodes, using our massively parallel real-time TDDFT implementation. We expect to transfer small amounts of data in and hundreds of Gigabyte out of the Blue Waters machine. Input and output requirements are modest for all these calculations.

Resources requested: We request a total of 184k node hours and the default storage allocation in the project directory as well as the Nearline storage.

Estimated use: Q1: 30 %, Q2: 30 %, Q3: 40 %

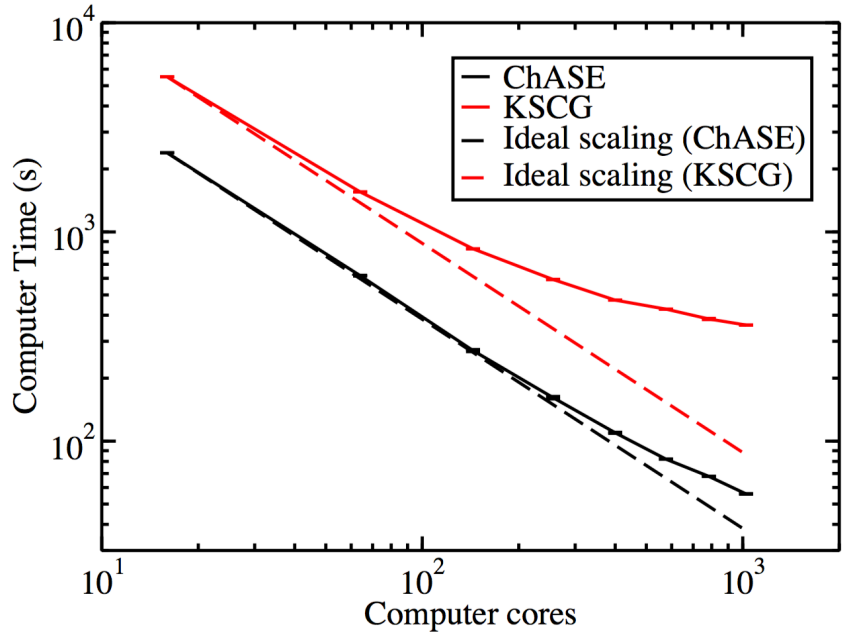

Fig. 1: Strong scalability test for the solver for the BSE matrix. The computer times for each MPI rank is the average over 8 runs on the same compute nodes. Ideal scaling is indicated as dashed lines.