Hasan Metin Aktulga
Solving Large-scale Eigenvalue Problems in Nuclear Configuration Interaction Calculations

Lawrence Berkeley National Lab
1 Cyclotron Rd
MS 50F-1650
Berkeley
CA 94720
hmaktulga@lbl.gov
Chao Yang
Esmond G. Ng
Pieter Maris
James P. Vary

The structures of nuclei can be analyzed by computing the eigenvalues and eigenvectors of a Hamiltonian matrix H constructed by what is known in quantum mechanics as the configuration interaction (CI) technique.In most cases, we are interested in the lowest few eigenvalues of H. In some applications, we may be interested in a relatively large number of low energy states (eigenvalues) with a prescribed total angular momentum, J. This type of calculation, which we will refer to as a total-J calculation, is particularly useful for investigating, among others, nuclear level densities as a function of J and excitation energy and for evaluating scattering amplitudes as a function of channel J. To achieve this goal, we perform a simultaneous diagonalization of H and J^2, where J^2 is the total angular momentum square operator.

The Hamiltonian and the angular momentum matrices constructed from the configuration interaction approximation are sparse. However, their dimensions can be extremely large. Therefore, the matrix eigenvalue problem must be solved by iterative methods.

In this talk, we discuss how to efficiently solve this type of eigenvalue problem on multi-core machines with tens of thousands of processing units. In particular, we will examine efficient strategies for distributing matrices to achieve good load balance, the communication overhead associated with orthogonalization and efficient ways to perform sparse matrix vector multiplications.

In addition to discussing these issues in the context of the Lanczos algorithm which is currently used in several nuclear CI codes, we will also examine how we may benefit from the use of a block iterative solver such as a polynomial accelerated subspace iteration or the locally optimal block conjugate gradient (LOBPCG) method, especially when the Hamiltonian is so large that it cannot be completely stored in the main memory.