

Abstract Submitted
for the MAR14 Meeting of
The American Physical Society

Optimizing GW for Petascale HPC and Beyond¹ JACK DESLIPPE, ANDREW CANNING, Lawrence Berkeley National Lab, YOUSEF SAAD, University of Minnesota, JAMES CHELIKOWSKY, University of Texas at Austin, STEVEN LOUIE, University of California, Berkeley — The traditional GW-Bethe-Salpeter (BSE) approach has, in practice, been prohibitively expensive on systems with more than 50 atoms. We show that through a combination of methodological and algorithmic improvements, the standard GW-BSE approach can be applied to systems with hundreds of atoms. We will discuss the massively parallel GW-BSE implementation in the BerkeleyGW package (on-top of common DFT packages) including the importance of hybrid MPI-OpenMP parallelism, parallel IO and library performance. We will discuss optimization strategies for and performance on many-core architectures.

¹Support for this work is provided through Scientific Discovery through Advanced Computing (SciDAC) program funded by U.S. Department of Energy, Office of Science, Advanced Scientific Computing Research and Basic Energy Sciences. Grant Number DE-FG02-12ER4

Jack Deslippe
Lawrence Berkeley National Lab

Date submitted: 13 Nov 2013

Electronic form version 1.4