

---

# Large-Scale GW Calculations on Pre-Exascale HPC Systems

**Mauro Del Ben**

Lawrence Berkeley National Laboratory  
Berkeley, CA 94720, USA  
mdelben@lbl.gov

**Andrew Canning**

Lawrence Berkeley National Laboratory

**Karthik Raman, Ruchira Sasanka**

Intel Corporation

**Steven G. Louie**

Materials Sciences Division, Lawrence  
Berkeley National Laboratory & University of  
California at Berkeley

**Felipe H. da Jornada**

Materials Sciences Division, Lawrence  
Berkeley National Laboratory & University of  
California at Berkeley

**Nathan Wichmann**

Cray

**Chao Yang**

Lawrence Berkeley National Laboratory

**Jack Deslippe**

NERSC, Lawrence Berkeley National  
Laboratory

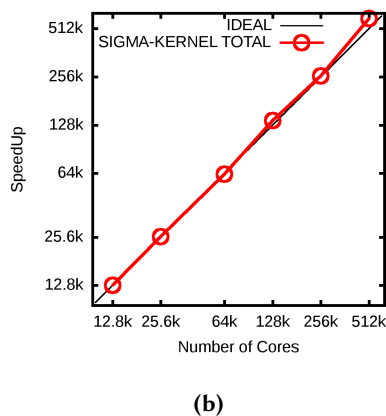
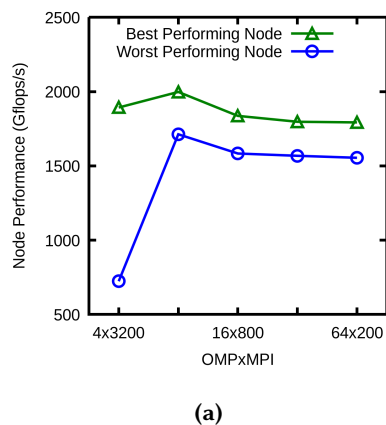
**ABSTRACT**

The accurate determination of excitation spectra of materials, such as the electronic band gap, is critical for the design of novel devices, including photovoltaics, transistors, batteries, and LEDs. Many-body perturbation-theory methods, and the *ab initio* GW approach [7, 8] in particular, have emerged over the last decades as the gold standard for computing these quantities.[9] However, the *ab initio* GW formalism is often limited to systems of at most 100 atoms due to its computational complexity. For this reason, it is important to show that, through the use of novel algorithms and an optimal HPC implementation, GW calculations containing 1000's of atoms are achievable on leadership-class HPC systems, such as the Cori system[2], a Cray XC40, Xeon-Phi powered system, at NERSC[10].

---

SC'17, Denver, Colorado USA

© 2017 ACM. This is the author's version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in *Proceedings of The International Conference for High Performance Computing, Networking, Storage and Analysis, November 2017*, <http://dx.doi.org/10.1145/nnnnnnn.nnnnnnn>.



**Figure 1:** For the BerkeleyGW-sigma code shown are (a) the performance per node measured as a function of the ratios of OpenMP threads and MPI tasks employing a total of 200 nodes and (b) the parallel speed-up with respect to 200 nodes.

We demonstrate this capability in this poster utilizing a highly-tuned version of the BerkeleyGW software package[1, 4] for the Cori system.

In particular we show that the GW approach, within a traditional sum-over-states framework, is well suited for exascale and pre-exascale HPC systems. The version of BerkeleyGW that we optimized is capable of scaling well to the full-Cori system at NERSC, using a high fraction of peak performance (> 11 PFLOPs) and achieving excellent time to solution (see figure 1). Additionally, there are multiple orders of magnitude more parallelism to exploit at the system level, node level and vector level. The ability to compute at near peak system performance ultimately derives from the fact that we can and have cast the computationally intense work at the node level as large ZGEMM operations. Optimal performance additionally required a new communication scheme that replaced MPI\_BCAST and MPI\_REDUCE statements with a ring-based (nearest neighbor) communication scheme that additionally allows overlapping with computation. Time to solution can be further reduced by approximations that reduce the plane-wave basis size, via low-rank approximations,[11–14] and compress the empty-orbital space required for summations.[3, 5, 6] The combination of methodological improvements as well as many-core HPC optimization allow us to perform large-scale GW calculations without sacrificing accuracy or convergence. We expect GW calculations to continue to grow in use as exascale-like systems with many-core processors and GPU acceleration gain traction, and as physicists and material scientists turn their attention to the study of novel materials requiring the use of methods that are more accurate and predictive for excitation spectra.

**KEYWORDS**

Large-scale GW calculations, BerkeleyGW, many-core HPC optimization, Cori, NERSC.

**ACM Reference format:**

Mauro Del Ben, Felipe H. da Jornada, Andrew Canning, Nathan Wichmann, Karthik Raman, Ruchira Sasanka, Chao Yang, Steven G. Louie, and Jack Deslippe. 2017. Large-Scale GW Calculations on Pre-Exascale HPC Systems. In *Proceedings of The International Conference for High Performance Computing, Networking, Storage and Analysis, Denver, Colorado USA, November 2017 (SC'17)*, 3 pages. DOI: 10.1145/nnnnnnn.nnnnnnn

**ACKNOWLEDGMENTS**

This work was supported by the Center for Computational Study of Excited-State Phenomena in Energy Materials at the Lawrence Berkeley National Laboratory, which is funded by the U.S. Department of Energy, Office of Science, Basic Energy Sciences, Materials Sciences and Engineering Division under Contract No. DE-AC02-05CH11231, as part of the Computational Materials Sciences Program.

## REFERENCES

- [1] BerkeleyGW 2017. <http://www.berkeleygw.org>. (2017).
- [2] CORI-NERSC 2017. <http://www.nersc.gov/systems/cori/>. (2017).
- [3] Jack Deslippe, Georgy Samsonidze, Manish Jain, Marvin L Cohen, and Steven G Louie. 2013. Coulomb-hole summations and energies for G W calculations with limited number of empty orbitals: A modified static remainder approach. *Physical Review B* 87, 16 (2013), 165124.
- [4] Jack Deslippe, Georgy Samsonidze, David A Strubbe, Manish Jain, Marvin L Cohen, and Steven G Louie. 2012. BerkeleyGW: A massively parallel computer package for the calculation of the quasiparticle and optical properties of materials and nanostructures. *Computer Physics Communications* 183, 6 (2012), 1269–1289.
- [5] Weiwei Gao, Weiyi Xia, Xiang Gao, and Peihong Zhang. 2016. Speeding up GW Calculations to Meet the Challenge of Large Scale Quasiparticle Predictions. *Scientific Reports* 6 (2016).
- [6] Linda Hung, H Felipe, Jaime Souto-Casares, James R Chelikowsky, Steven G Louie, and Serdar Ögüt. 2016. Excitation spectra of aromatic molecules within a real-space G W-BSE formalism: Role of self-consistency and vertex corrections. *Physical Review B* 94, 8 (2016), 085125.
- [7] Mark Hybertsen and Steven G Louie. 1985. First-principles theory of quasiparticles: calculation of band gaps in semiconductors and insulators. *Physical review letters* 55, 13 (1985), 1418.
- [8] Mark S Hybertsen and Steven G Louie. 1986. Electron correlation in semiconductors and insulators: Band gaps and quasiparticle energies. *Physical Review B* 34, 8 (1986), 5390.
- [9] Manish Jain, James R Chelikowsky, and Steven G Louie. 2011. Reliability of hybrid functionals in predicting band gaps. *Physical review letters* 107, 21 (2011), 216806.
- [10] NERSC 2017. <http://www.nersc.gov>. (2017).
- [11] Huy-Viet Nguyen, T Anh Pham, Dario Rocca, and Giulia Galli. 2012. Improving accuracy and efficiency of calculations of photoemission spectra within the many-body perturbation theory. *Physical Review B* 85, 8 (2012), 081101.
- [12] T Anh Pham, Huy-Viet Nguyen, Dario Rocca, and Giulia Galli. 2013. GW calculations using the spectral decomposition of the dielectric matrix: Verification, validation, and comparison of methods. *Physical Review B* 87, 15 (2013), 155148.
- [13] Hugh F. Wilson, François Gygi, and Giulia Galli. 2008. Efficient iterative method for calculations of dielectric matrices. *Physical Review B* 78 (2008), 113303.
- [14] Hugh F. Wilson, Deyu Lu, François Gygi, and Giulia Galli. 2009. Iterative calculations of dielectric eigenvalue spectra. *Physical Review B* 79 (2009), 245106.