Scalable Computational Tools for Discovery and Design -- Excited State Phenomena in Energy Materials

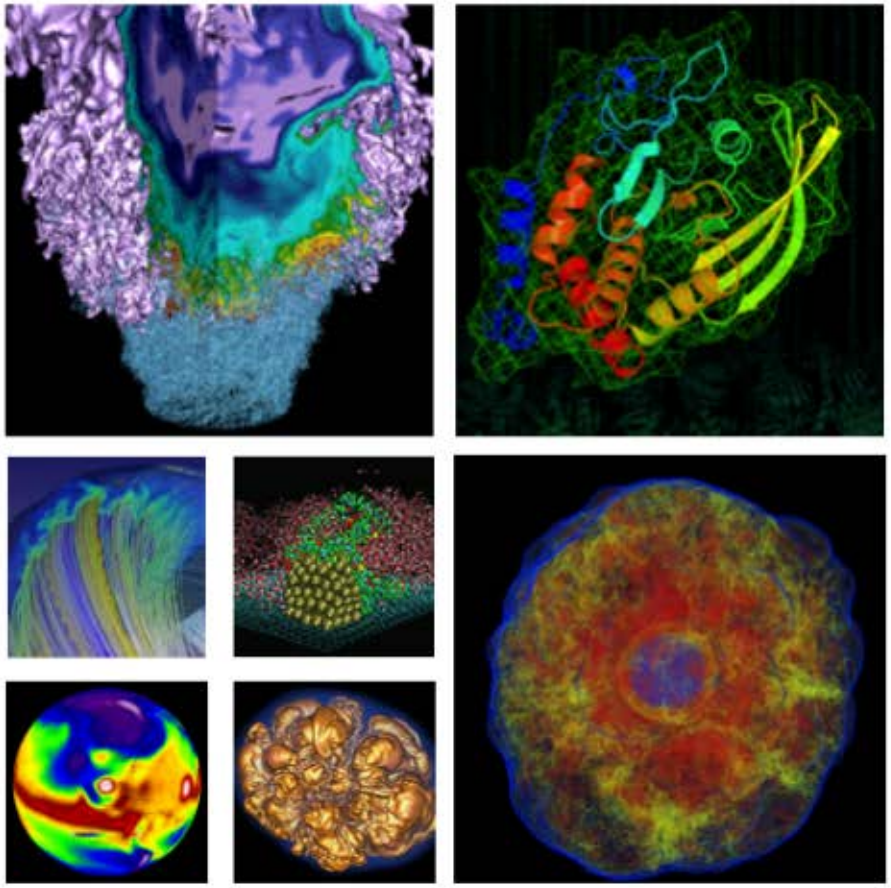

NERSC

Jack Deslippe (Team Rep)

U.S. DEPARTMENT OF ENERGY | Office of Science

THE UNIVERSITY OF TEXAS AT AUSTIN

BERKELEY LAB

PIs

Jim Chelikowsky
UT Austin

Jeff Neaton
LBNL / UC Berkeley

Steven Louie
LBNL / UC Berkeley

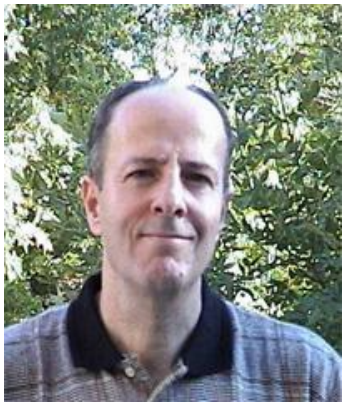

Yousef Saad
U. of Minnesota

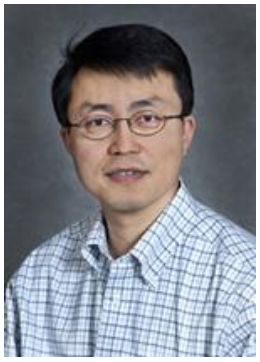

Chao Yang
LBNL

Andrew Canning
LBNL

Alex Demkov UT Austin

Jack Deslippe
LBNL

Two Complementary Codes For Computational Discover and Design

PARSEC

- **Massively parallel real-space DFT code.**
- **Capable of studying systems of 10K atoms.**
- **Implements spectrum slicing approach for parallel eigenstate generation.**

BerkeleyGW

- **Massively parallel excited state (both one- and two-particle) code.**
- **Computes quasiparticle and optical properties of materials of interest to the DOE.**

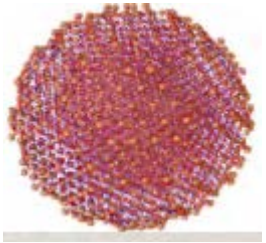
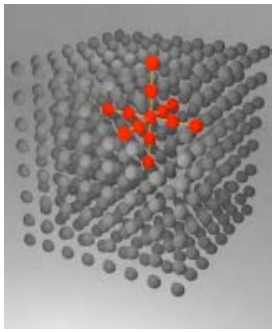
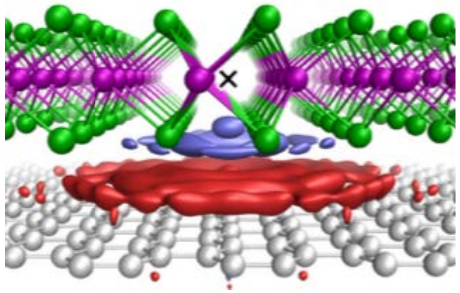
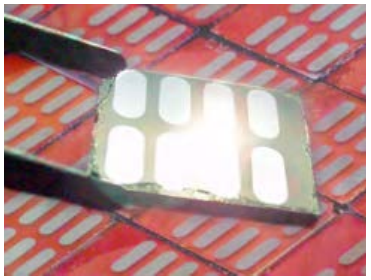

Jack Deslippe

- **NERSC User Services Group (Materials Science / Chemistry Consultant)**
- **NESAP Lead (NERSC's exascale readiness program)**
- **Developer in BerkeleyGW project (SCIDAC Team Member)**
- **My Focus in this presentation will be on BerkeleyGW enhancements over the last couple of years.**

$$[E_{n\mathbf{k}} - H_0(\mathbf{r}) - V_H(\mathbf{r})]\psi_{n\mathbf{k}}(\mathbf{r}) - \int \Sigma(\mathbf{r}, \mathbf{r}', E_{n,\mathbf{k}})\psi_{n\mathbf{k}}(\mathbf{r}')d\mathbf{r}' = 0$$

Materials:

InSb, InAs
Ge
GaSb
Si
InP
GaAs
CdS
AlSb, AlAs
CdSe, CdTe
BP
SiC
C_{60}
GaP
AlP
ZnTe, ZnSe
c-GaN, w-GaN
InS
w-BN, c-BN
diamond
w-AlN
LiCl
Fluorite
LiF

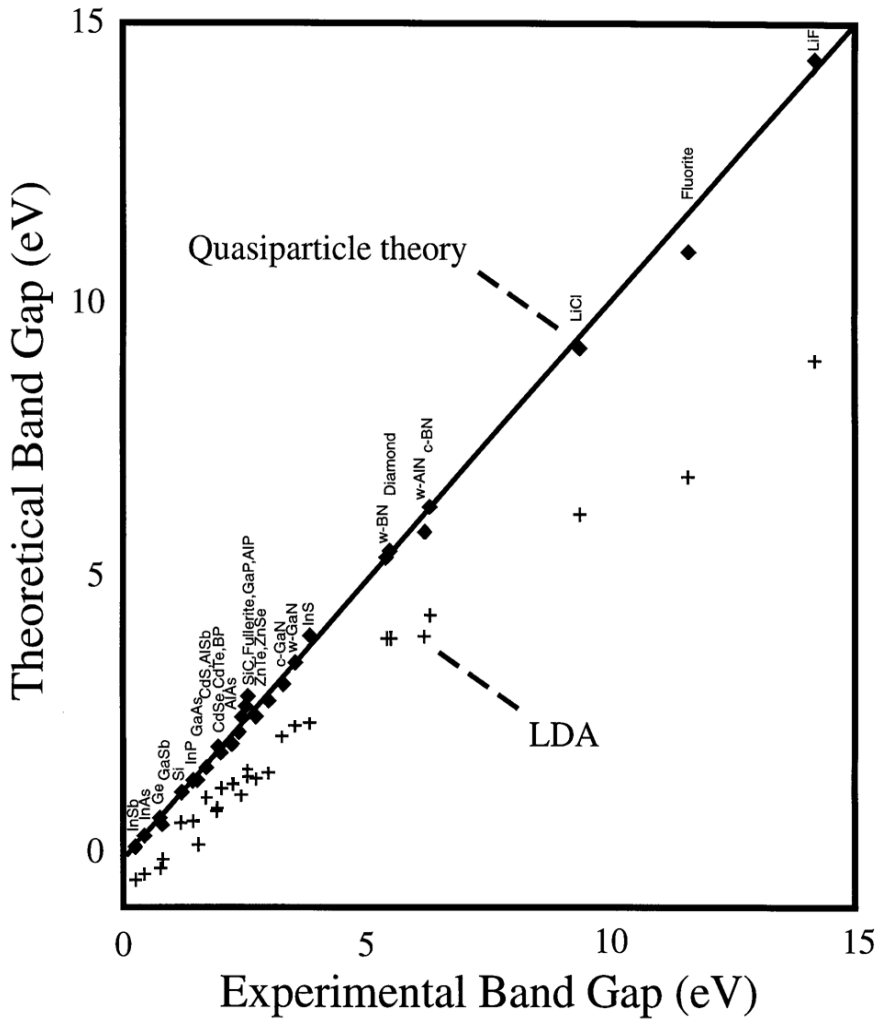

What is GW+BSE

Many-body effects extremely important in **Excited-State properties** of Complex Materials.

Includes screened-interaction for many-body effects

Accurately describes properties important for:

- Photovoltaics
- LEDs
- Junctions / Interfaces
- Defect Energy Levels
- ….

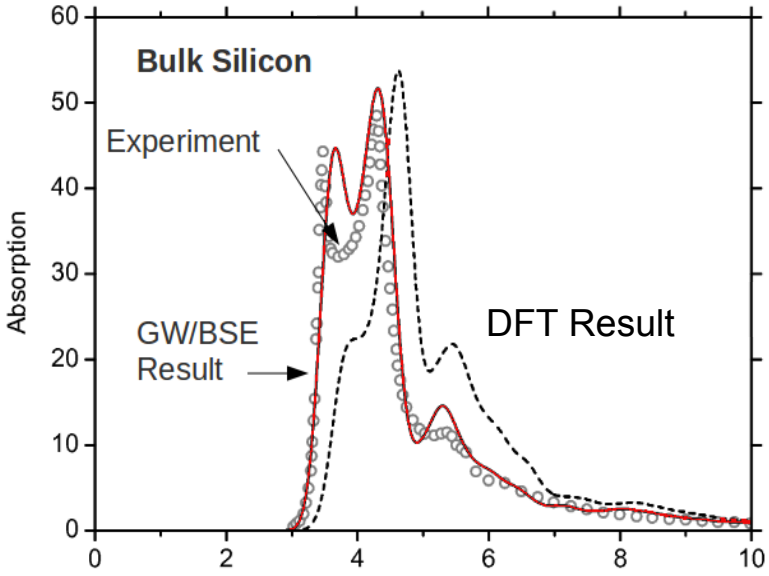

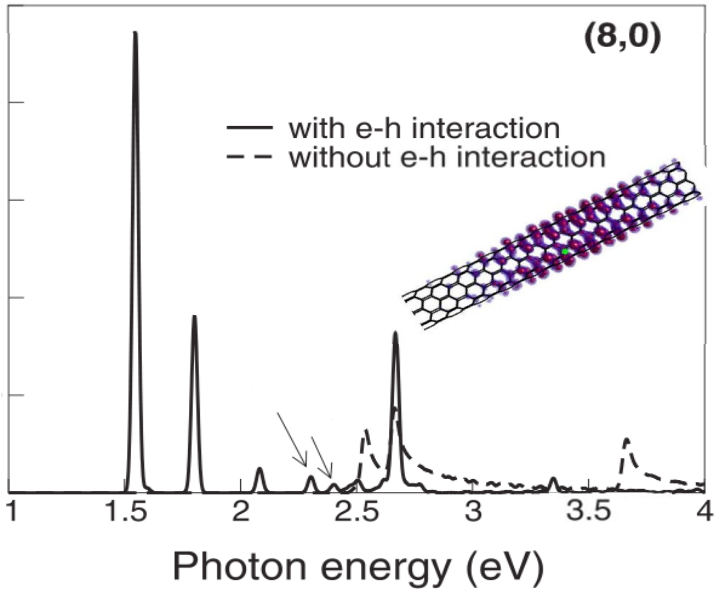

The Good:

Quantitatively accurate for quasiparticle properties in a wide variety of systems.

Accurately describes dielectric screening important in excited state properties.

The Bad:

Prohibitively slow for large systems. Usually thought to cost orders of magnitude more time that DFT.

Memory intensive and scales badly. Exhausted by storage of the dielectric matrix and wavefunctions. Limited ~50 atoms.

The Good:

Quantitatively accurate for quasiparticle properties in a wide variety of systems.

Accurately describes dielectric screening important in excited state properties.

The Bad:

Prohibitively slow for large systems. Usually thought to cost orders of magnitude more time than DFT.

Memory intensive and scales badly. Exhausted by storage of the dielectric matrix and wavefunctions. Limited ~50 atoms.

BerkeleyGW Usage at NERSC

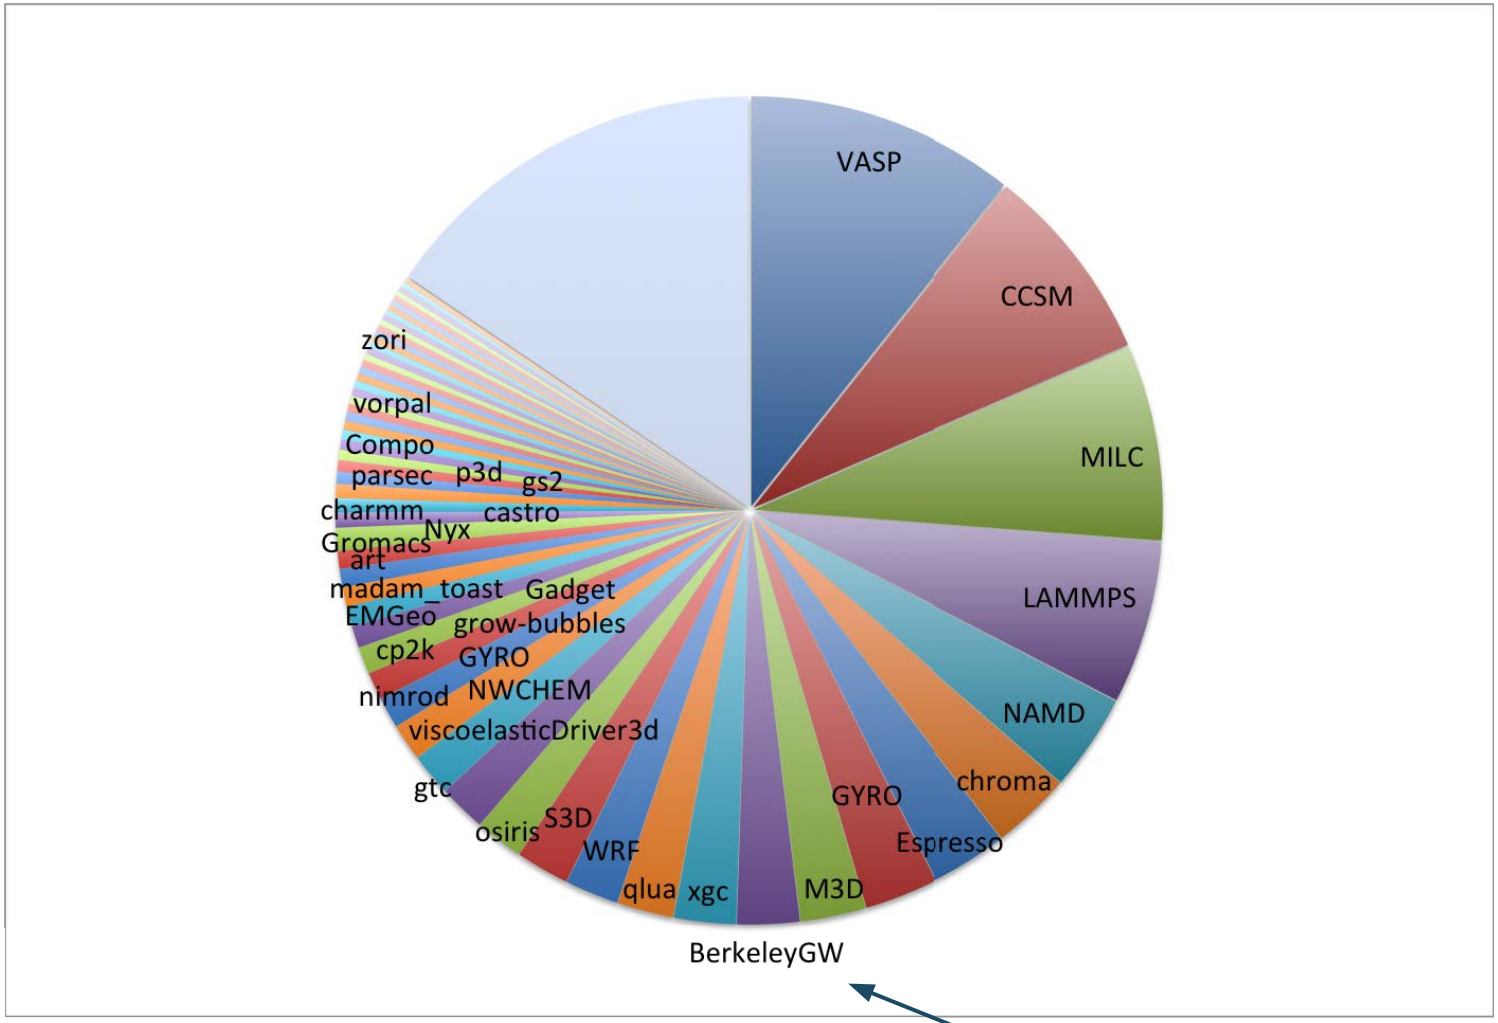

NERSC Code Breakdown 2013

BerkeleyGW Usage at NERSC

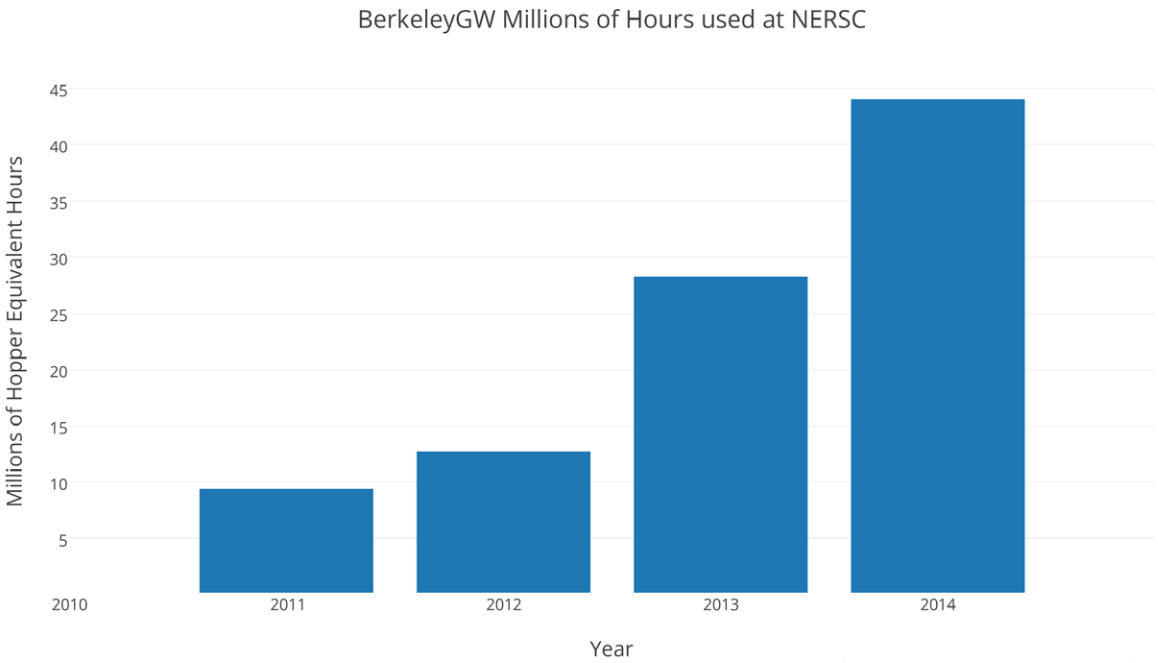

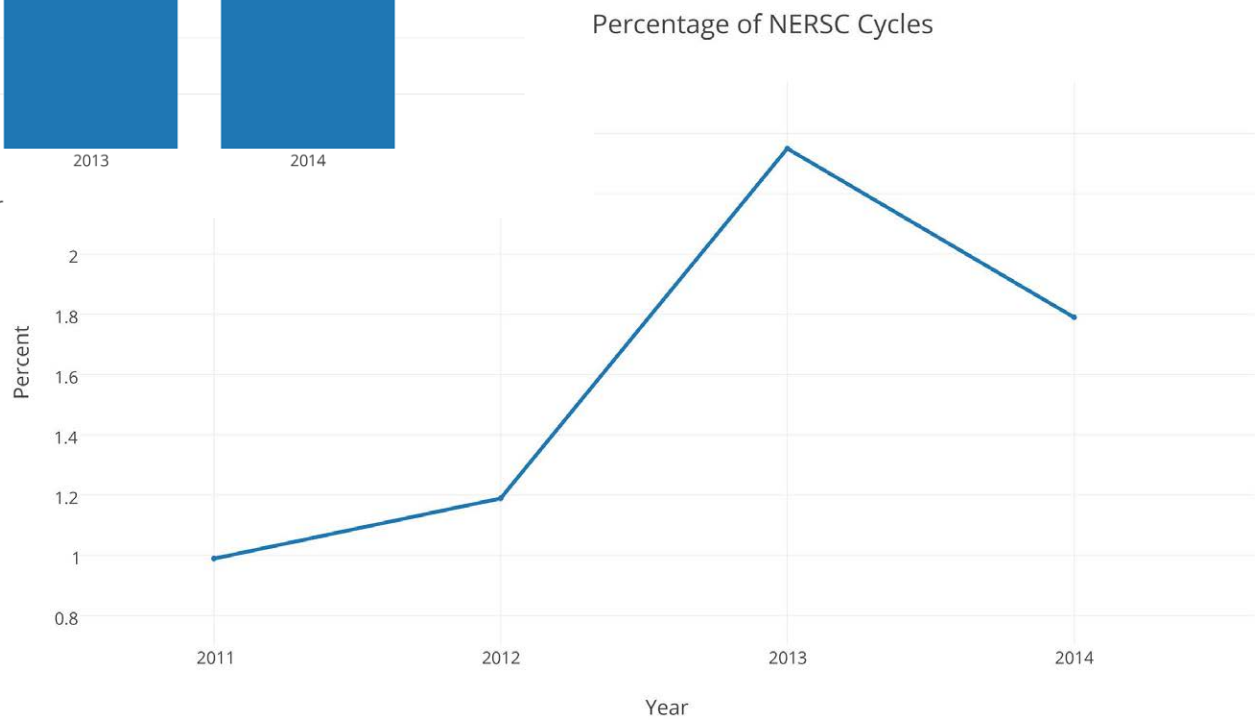

U.S. DEPARTMENT OF ENERGY | Office of Science

Science Stories Using BerkeleyGW

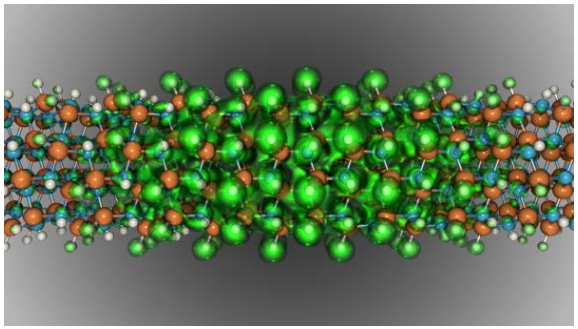

PI: Kioupakis

TO BRIDGE LEDS' GREEN GAP, SCIENTISTS THINK SMALL

PI: Louie

First Ab Initio Method for Characterizing Hot Carriers Could Hold the Key to Future Solar Cell Efficiencies

Science Shorts Lynn Yarris • JULY 17, 2014

Tweet 112 +1 58 Like 78 Share 0 Email 7

One of the major road blocks to the design and development of new, more efficient solar cells may have been cleared. Researchers with the Lawrence Berkeley National Laboratory (Berkeley Lab) have developed the first *ab initio* method – meaning a theoretical model free of adjustable or empirical parameters – for characterizing the properties of "hot carriers" in semiconductors. Hot carriers are electrical charge carriers – electrons and holes – with significantly higher energy than charge carriers at thermal equilibrium.

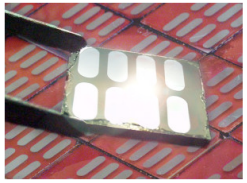

A new and better way to study "hot" carriers in semiconductors, a major source of efficiency loss in solar cells, has been developed by scientists at Berkeley Lab. (Photo by Roy Kaltschmidt)

PI: Wu

Cover Image: Phys. Rev. Lett. Vol. 110, Iss. 1

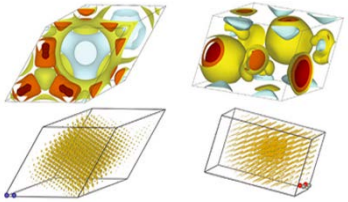

From the article:

Origin of the Variation of Exciton Binding Energy in Semiconductors

Marc Dvorak, Su-Huai Wei, and Zhigang Wu

Phys. Rev. Lett. **110**, 016402 (2013)

PI: Van De Walle / Cohen

IMPORTANT NEW METHOD FOR STUDYING SOLAR MATERIALS

The goal of research in this group in general has been to develop and use first-principles computational methods to understand, predict, and design novel electronic, optoelectronic, photovoltaic, and thermoelectric materials.

Using a NERSC NISE award, researchers were able to compute the phonon-assisted interband optical absorption spectrum of silicon entirely from first principles.

Nearly all commercially-available photovoltaic cells currently depend on this absorption process.

The new method is general enough to to study fundamental physics of other optoelectronic and photovoltaic materials and can address questions that are not accessible by experiment.

Used the BerkeleyGW software written by new NERSC USG consultant Jack Deslippe.

Work done by Jesse Noffsinger, Emmanouil Kioupakis, Chris G. Van de Walle, Steven G. Louie, and Marvin L. Cohen.

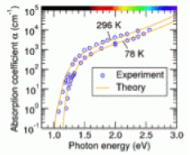

Calculated (solid lines) and experimental (circles) absorption coefficient of silicon in the energy range between the indirect and direct gaps, for two temperatures.

PI: Prendergast

LASER, SUPERCOMPUTER MEASURE SPEEDY ELECTRONS IN SILICON

Simulations at NERSC Help Illuminate Attosecond Laser Experiment Findings

DECEMBER 19, 2014 | Tags: Basic Energy Sciences (BES), Materials Science

Contact: Robert Sanders, rlsanders@berkeley.edu, (510) 643-6998

The entire semiconductor industry, not to mention Silicon Valley, is built on the propensity of electrons in silicon to get kicked out of their atomic shells and start to move through the material. These mobile electrons are routed and switched though transistors, carrying the digital information that characterizes our age.

An international team of physicists and chemists at the University of California, Berkeley, has for the first time taken snapshots of this ephemeral event using attosecond pulses of soft x-ray light lasting only a few billionths of a billionth of a second. The researchers then used supercomputing resources at Lawrence Berkeley National Laboratory's (Berkeley Lab) National Energy Research Scientific Computing Center (NERSC) to help them better understand their findings.

While earlier femtosecond lasers were unable to resolve the jump from the valence shell of the

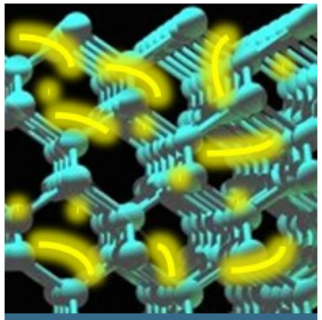

In silicon, electrons attached to atoms in the crystal lattice can be mobilized into the conduction band by light or voltage. Berkeley scientists have taken snapshots of this very brief band-gap jump and timed it at 450 attoseconds. Image: Stephen Leone

U.S. DEPARTMENT OF ENERGY | Office of Science

BERKELEY LAB

BerkeleyGW Workshops

- Emphasized integration with **PARSEC** code
- 3 times the number of applicants than space available! 45 Attendees.
- Survey results show that 100% of attendees found sessions useful of very useful.

BerkeleyGW in the Many-Core Era

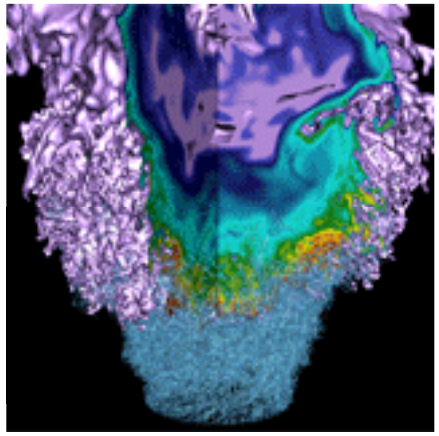

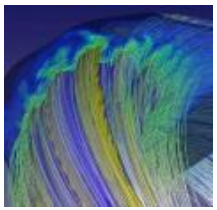

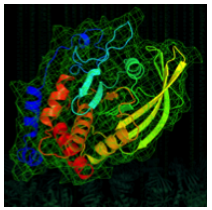

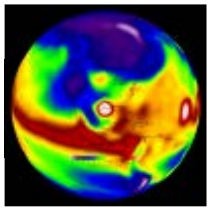

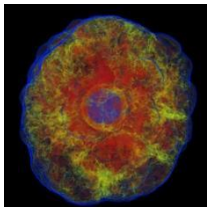

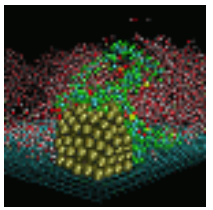

NERSC

U.S. DEPARTMENT OF ENERGY | Office of Science

BERKELEY LAB
Lawrence Berkeley National Laboratory

Cori

- **Cori will begin to transition the workload to more energy efficient architectures**
- **Cray XC system with over 9300 Intel Knights Landing (Xeon-Phi) nodes**
  - Self-hosted, (not an accelerator) manycore processor with over 60 cores per node + high-bandwidth memory
- **Data Intensive Science Support**
  - NVRAM Burst Buffer to accelerate applications

System named after Gerty Cori, Biochemist and first American woman to receive the Nobel prize in science.

What is different about Cori?

Edison (Ivy-Bridge):

- 12 Cores Per CPU
- 24 Virtual Cores Per CPU
- 2.4-3.2 GHz
- Can do 4 Double Precision Operations per Cycle
- 2.5 GB of Memory Per Core
- ~100 GB/s Memory Bandwidth

Cori (Knights-Landing):

- 60+ Physical Cores Per CPU
- 240+ Virtual Cores Per CPU
- Much slower GHz
- Can do 8 Double Precision Operations per Cycle
- < 0.3 GB of Fast Memory Per Core
  < 2 GB of Slow Memory Per Core
- Fast memory has ~ 5x DDR4 bandwidth

Optimization Strategy For Cori

Both **PARSEC** and **BerkeleyGW** are included in top tier of the NERSC Exascale Application Program (NESAP).

- Work with Cray, NERSC, Intel and SUPER staff
- Early access to simulators and hardware.

Strategy:

A. Add OpenMP or other on-node parallelism

B. Effectively use vector Instructions

C. Identify/optimize memory bandwidth hotspots.

Optimization Strategy For Cori

Both **PARSEC** and **BerkeleyGW** are included in top tier of the NERSC Exascale Application Program (NESAP).

- Work with Cray, NERSC, Intel and SUPER staff
- Early access to simulators and hardware.

Strategy:

A. Add OpenMP or other on-node parallelism

B. Effectively use vector Instructions

C. Identify/optimize memory bandwidth hotspots.

```
do i = 1, n
  a(i) = b(i) + c(i)
enddo
```

$$\begin{pmatrix} a_1 \\ \dots \\ a_n \end{pmatrix} = \begin{pmatrix} b_1 \\ \dots \\ b_n \end{pmatrix} + \begin{pmatrix} c_1 \\ \dots \\ c_n \end{pmatrix}$$

Example Use Case For OpenMP in BerkeleyGW and PARSEC

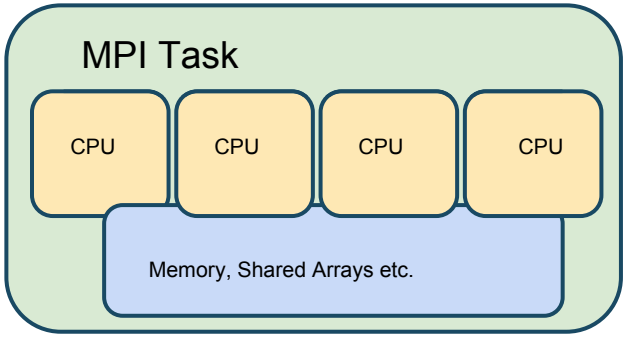

Parallel FFTs involve MPI all-to-all communication (small messages, latency bound).

Reducing the number of MPI tasks in favor OpenMP threads makes large improvement in overall runtime.

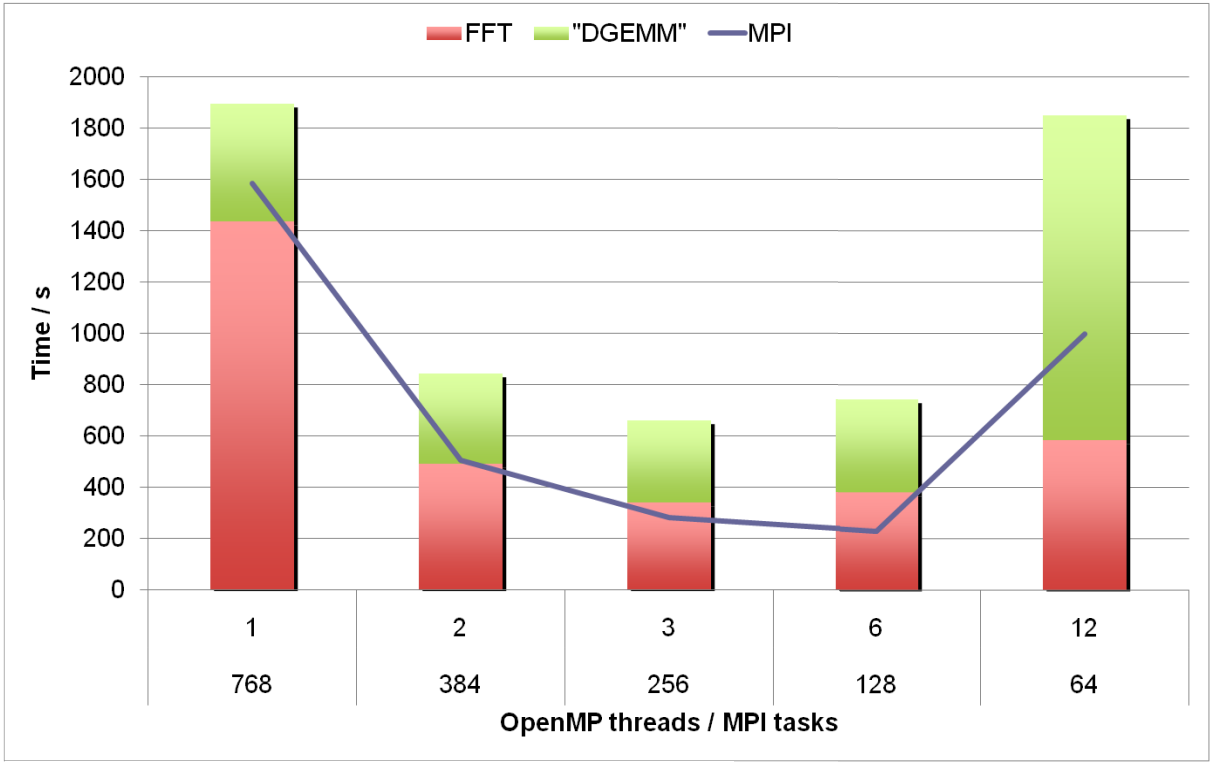

Work by Andrew Canning

BerkeleyGW Self-Energy Calculation

Significant Bottleneck is large matrix reduction like operations. Turning arrays into numbers.

$$\langle n\mathbf{k}|\,\Sigma_{\mathrm{CH}}(E)\,|n'\mathbf{k}\rangle = \frac{1}{2}\sum_{n''}\sum_{\mathbf{q}\mathbf{G}\mathbf{G}'} M^{*}_{n''n}(\mathbf{k},-\mathbf{q},-\mathbf{G})M_{n''n'}(\mathbf{k},-\mathbf{q},-\mathbf{G}') \times \frac{\Omega^{2}_{\mathbf{G}\mathbf{G}'}(\mathbf{q})\,(1-i\tan\phi_{\mathbf{G}\mathbf{G}'}(\mathbf{q}))}{\tilde{\omega}_{\mathbf{G}\mathbf{G}'}(\mathbf{q})\,(E-E_{n''\mathbf{k}-\mathbf{q}}-\tilde{\omega}_{\mathbf{G}\mathbf{G}'}(\mathbf{q}))}\,v(\mathbf{q}+\mathbf{G}')$$

Targeting Intel Xeon Phi Many Core Architecture

1. Target more on-node parallelism. (MPI model already failing users)
2. Ensure key loops/kernels can be vectorized.

Example: Optimization steps for Xeon Phi Coprocessor

Add OpenMP

Refactor to Have 3 Loop Structure:

Outer: MPI
Middle: OpenMP
Inner: Vectorization

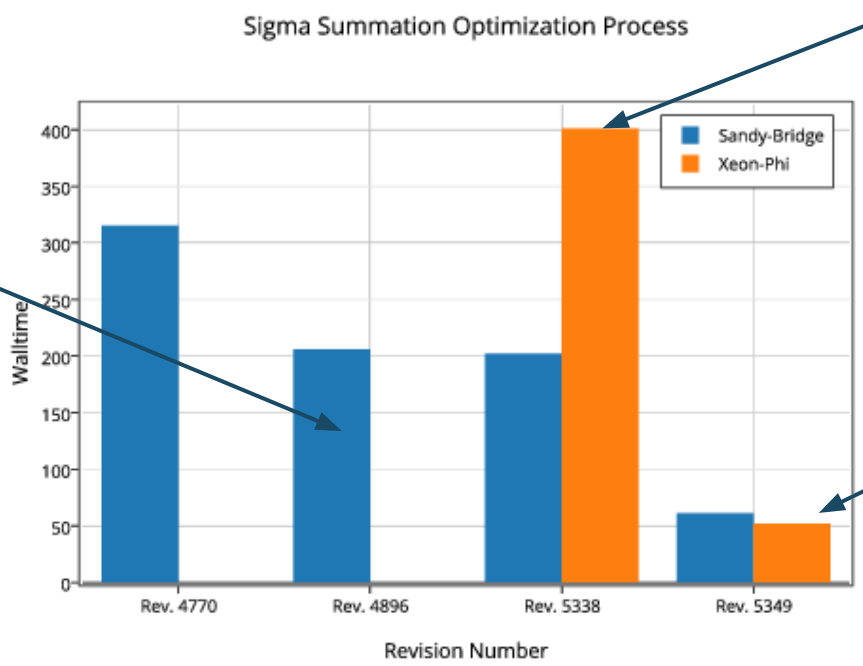

Cache reuse and Ensure Vectorization

Added additional layers of cache-blocking to improve locality. Important on Xeon-Phi which lacks L3.

Bounding Around on the Roofline Model

2013 - Low Operational Intensity

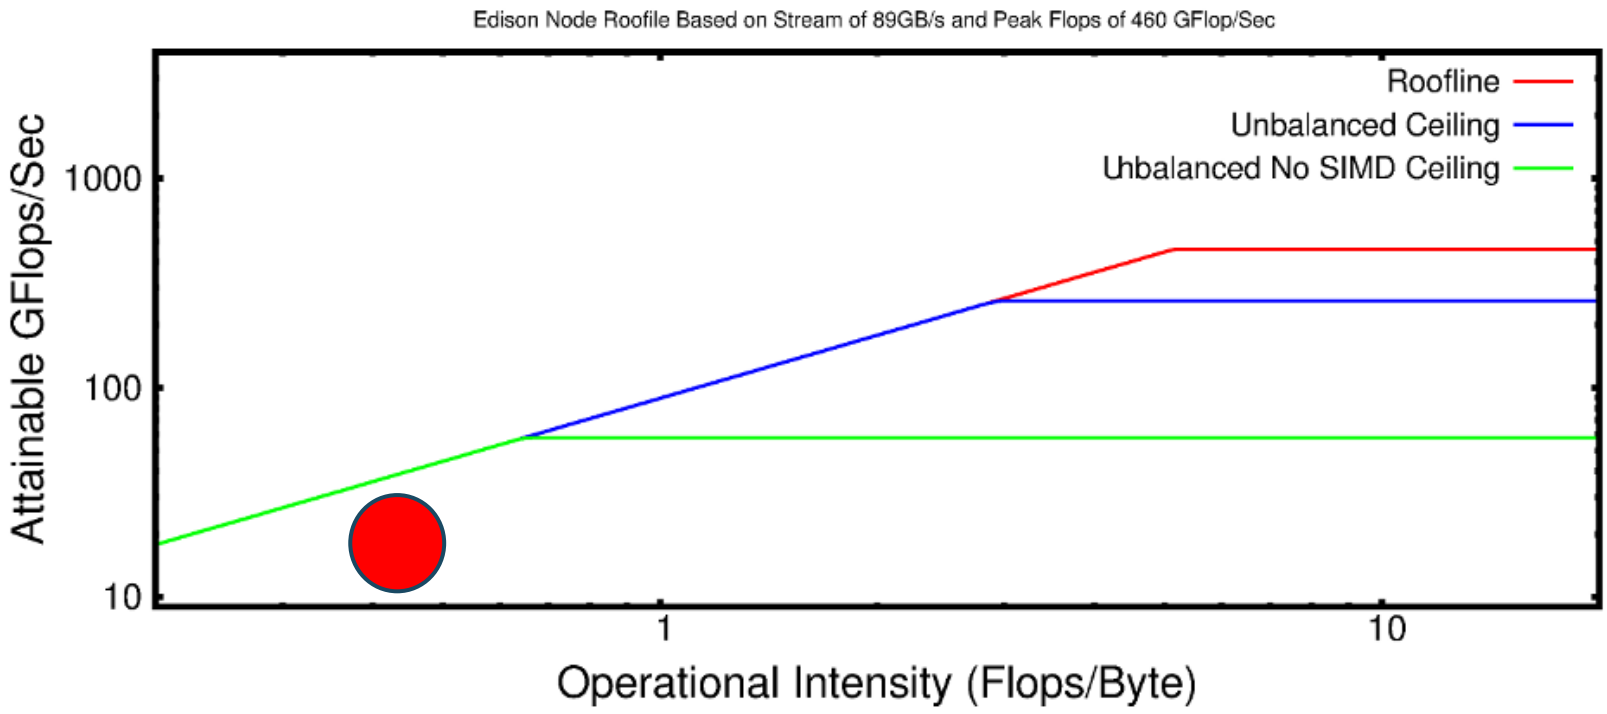

Bounding Around on the Roofline Model

2014 - Refactored loops, improved locality

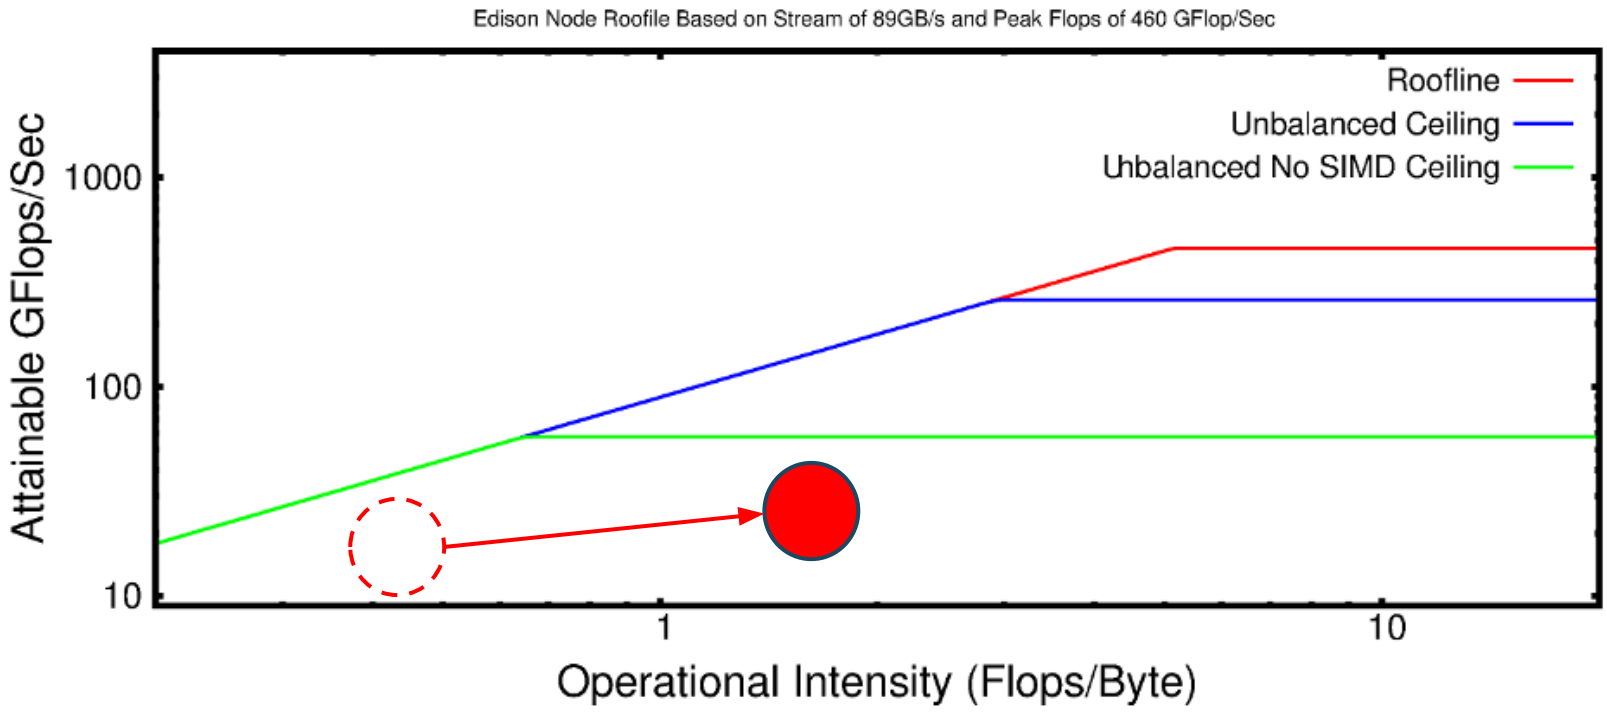

Bounding Around on the Roofline Model

2014 - Vectorized Code

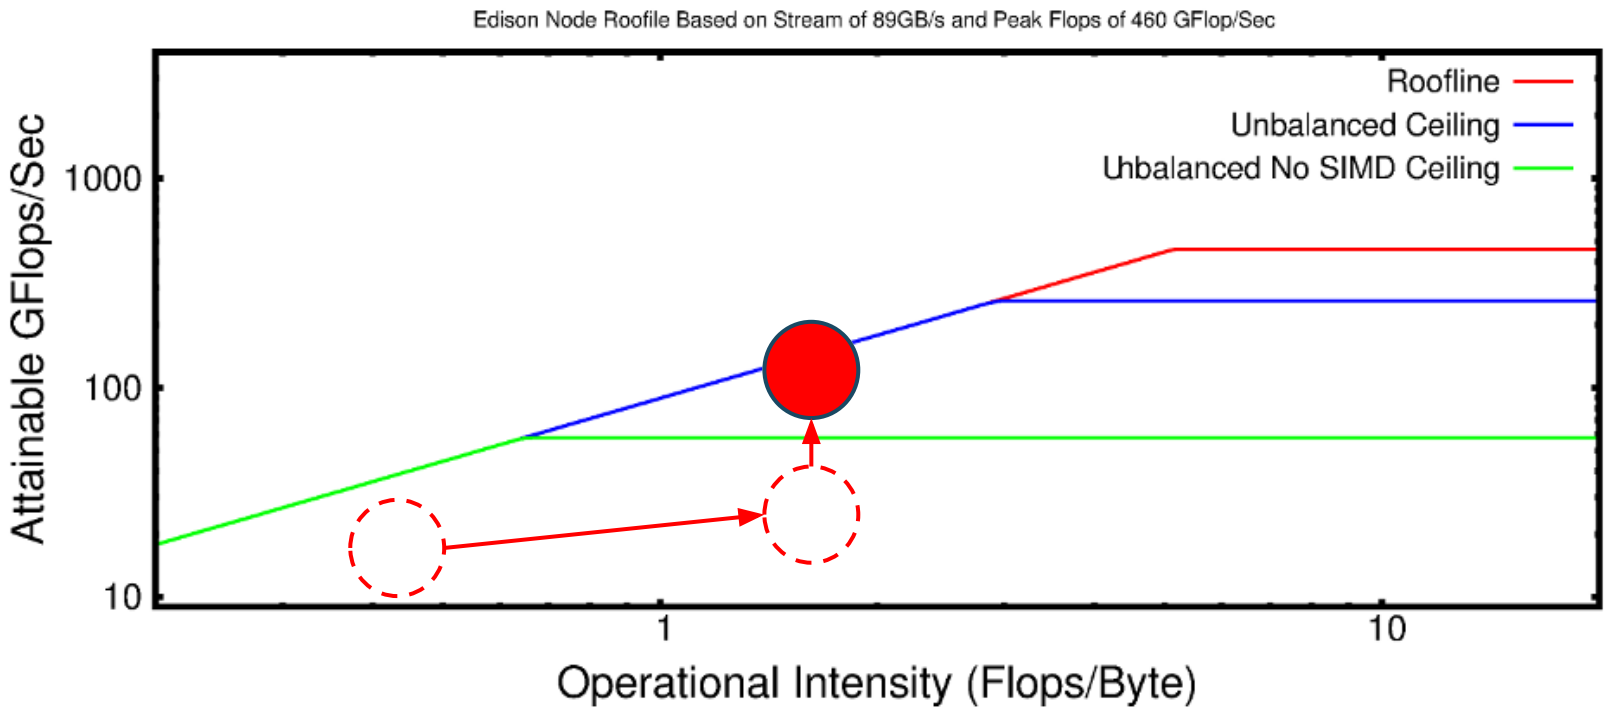

Bounding Around on the Roofline Model

2015 - Cache Blocking

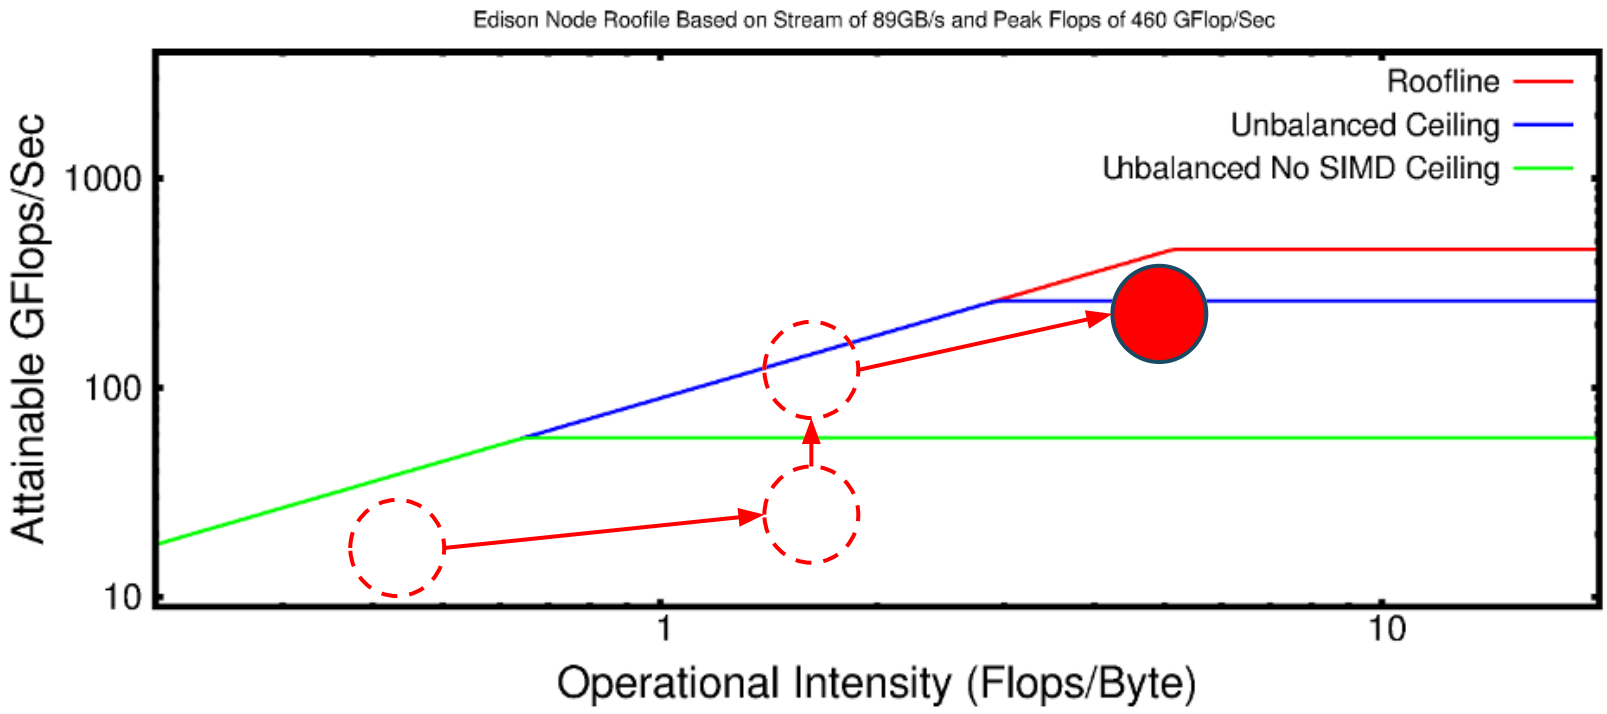

Hybrid MPI-OpenMP Scaling Improvements.

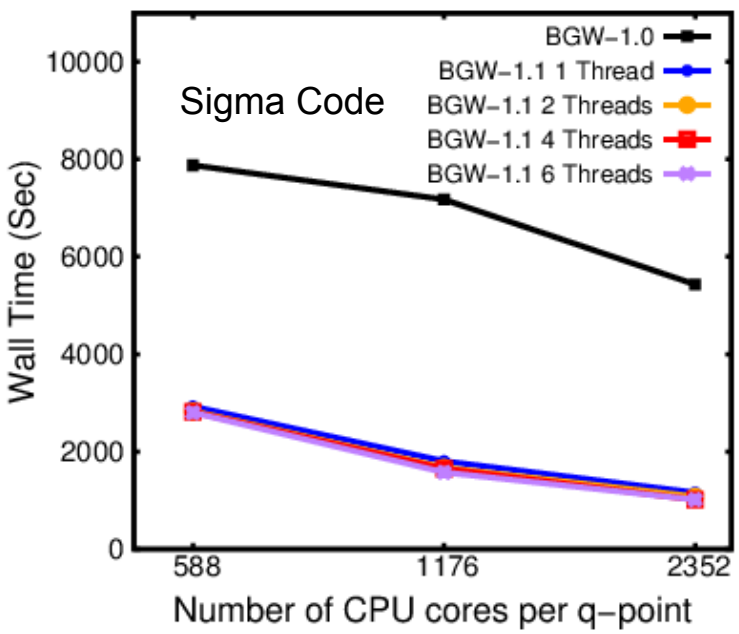

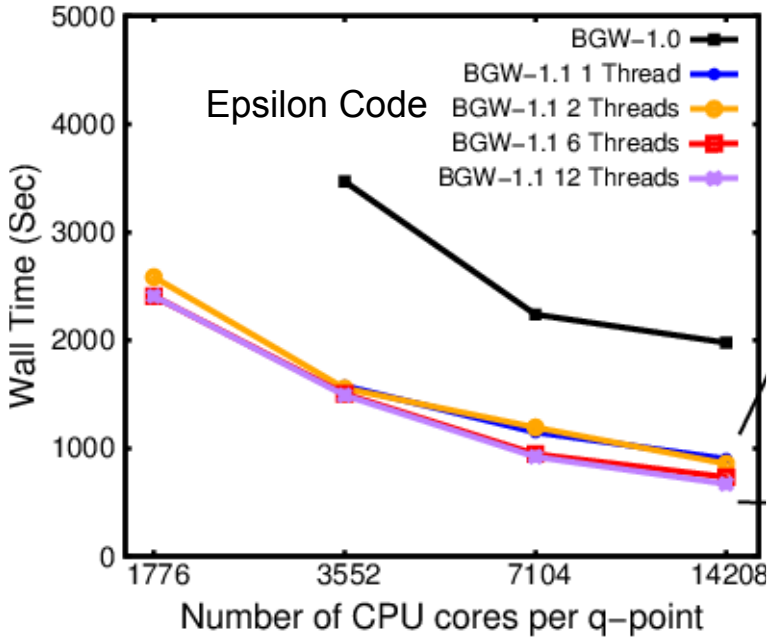

Science/Method Advances

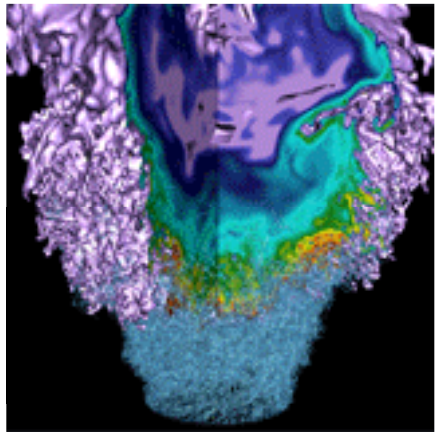

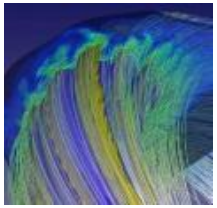

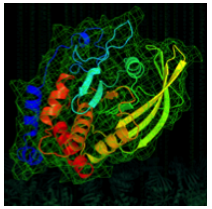

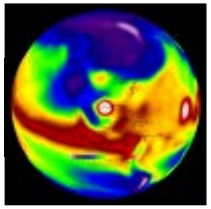

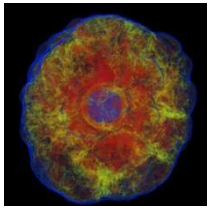

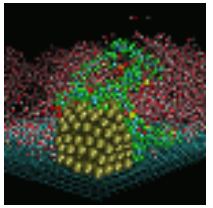

NERSC

U.S. DEPARTMENT OF ENERGY | Office of Science

BERKELEY LAB
Lawrence Berkeley National Laboratory

New Features

- Exciton Bandstructures (Finite Center of Mass Momentum Excitons)
- Full-Frequency Calculations With ~2X GPP Cost (Hilbert Transform and Contour Deformation Approach)
- Parallelization Over Frequencies
- Parallel IO and Transferable File Format
- Parallel reduced size FFTs
- **Vastly Improved K-Point Convergence**
- Support for PARSEC Input, Abinit Input, RMG Input
- Support for Static COHSEX Starting Point
- Empty State Requirement Reduction
- **Full BSE Calculations in Parallel**
- **Accurate/efficient GPP models for Informatics**

- Over 2500 Commits since BerkeleyGW 1.0. Many performance improvements, bug fixes and new features

Improving K-Point Integration

How to efficiently capture long wavelength features in the dielectric matrix?

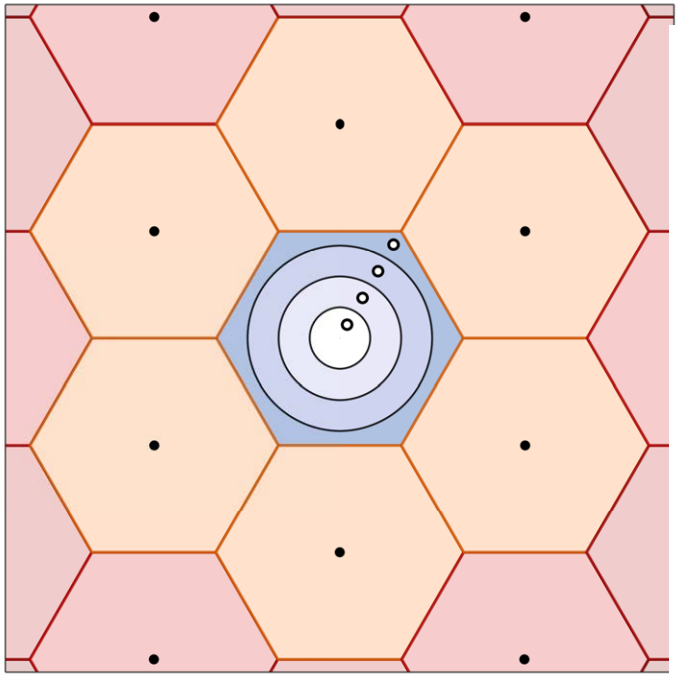

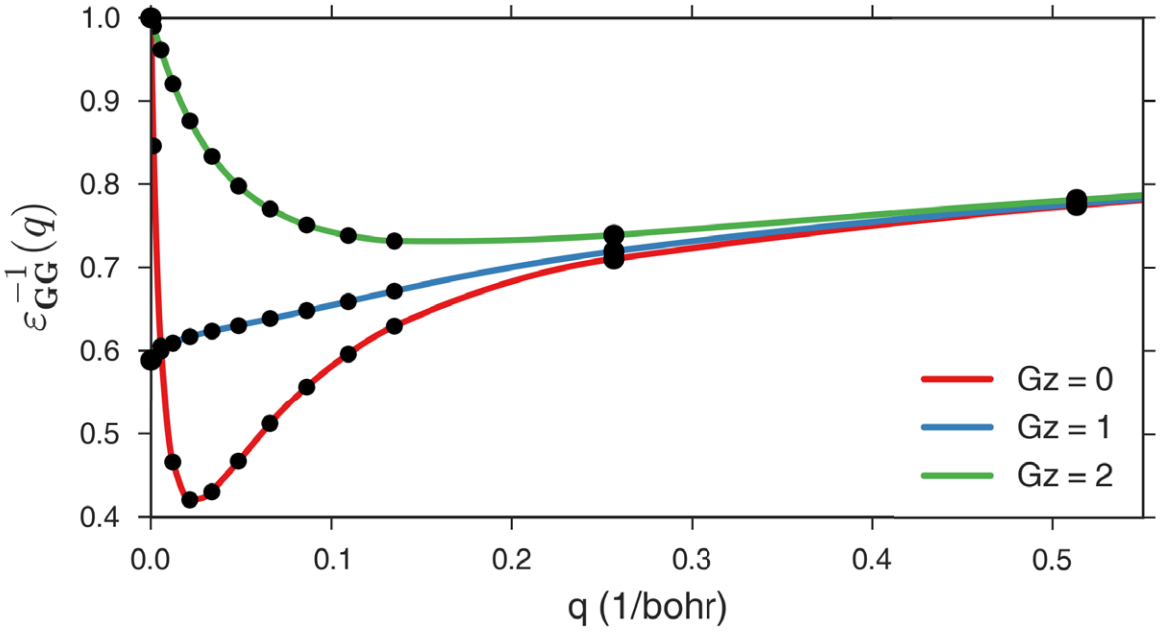

Our solution: subsampling method - a hybrid sampling of the Brillouin Zone.

Improving K-Point Integration

Subsampling

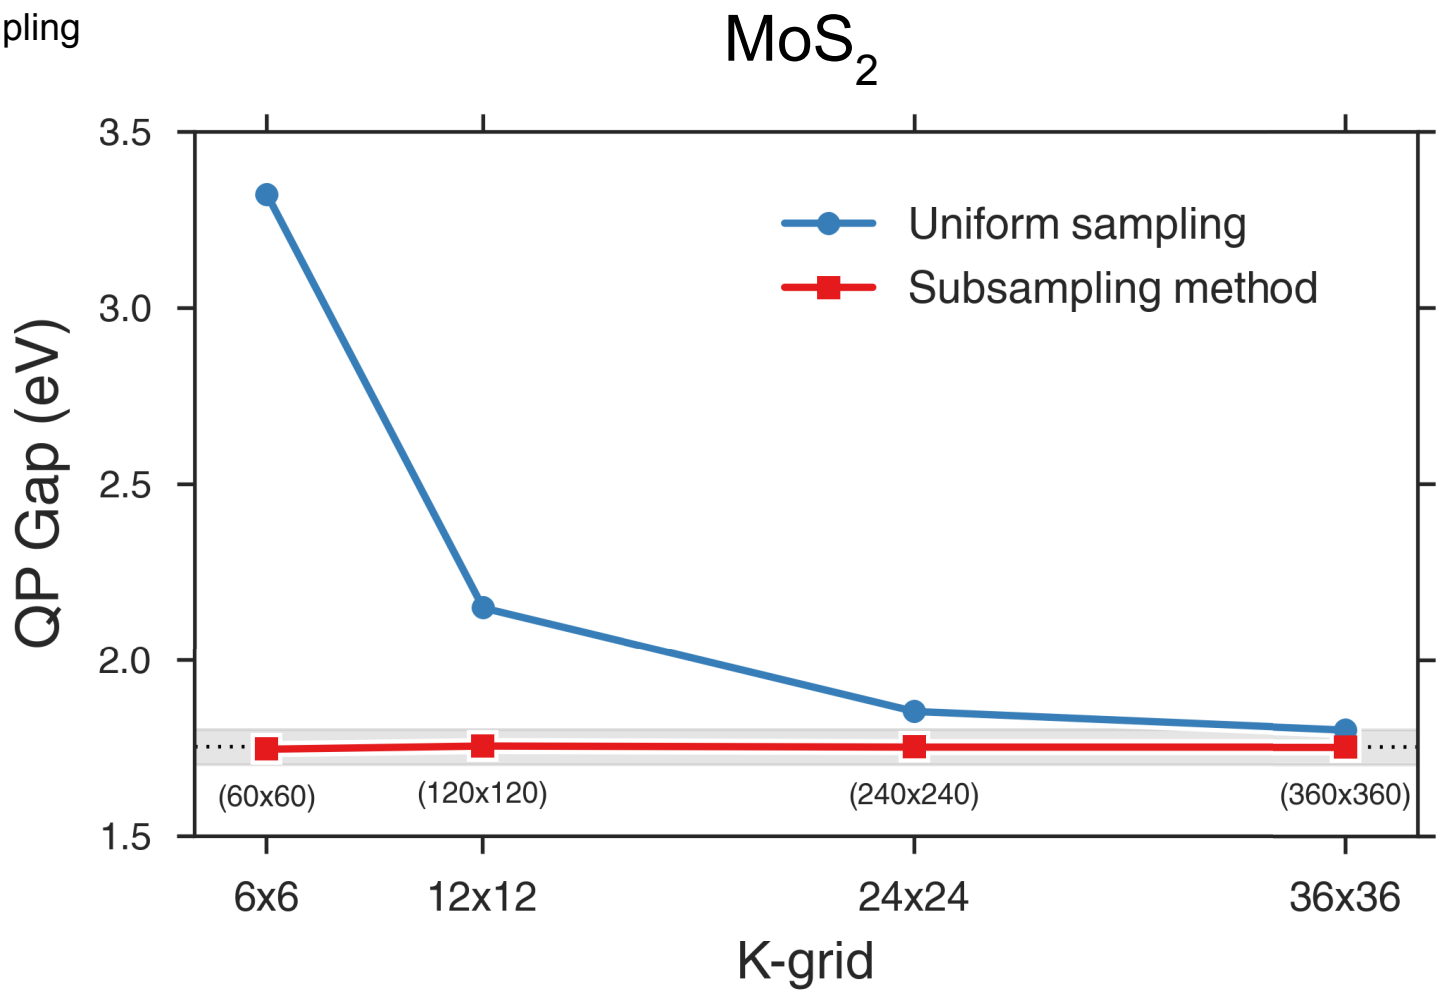

Improving K-Point Integration

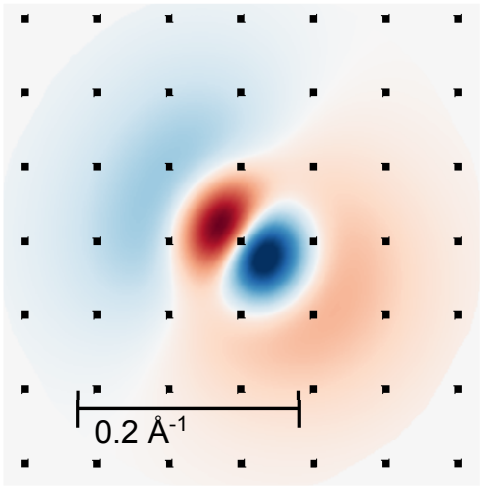

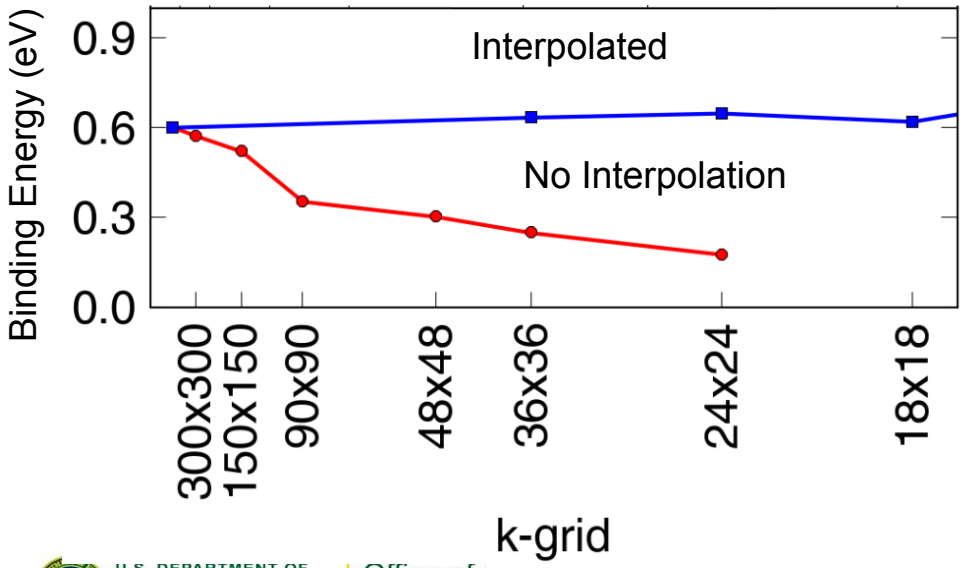

Improved cluster interpolation method for BSE

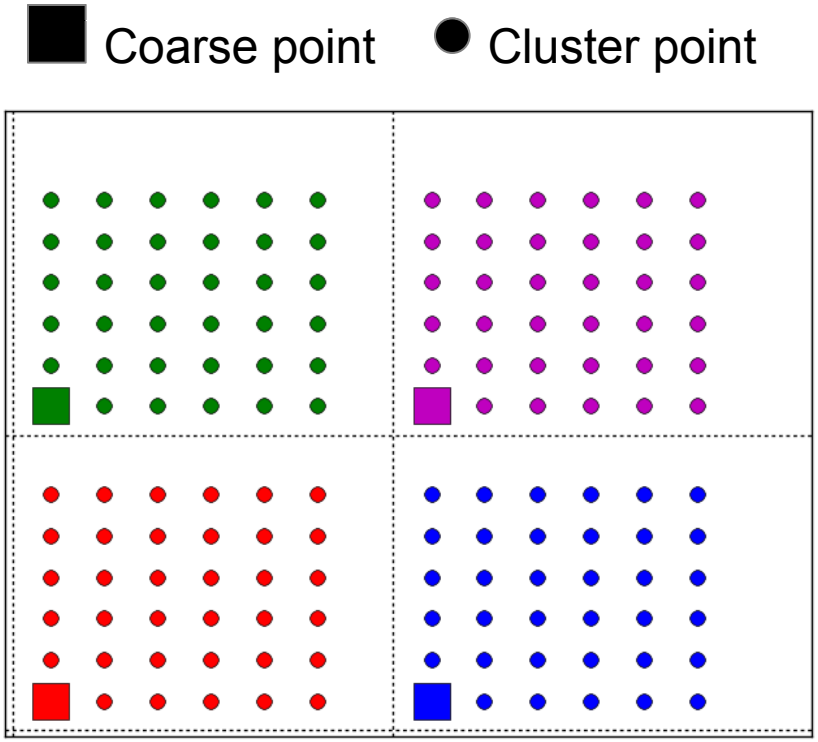

Efficient Full BSE Calculations

The Full BSE Hamiltonian is complex, non Hermitian

$$H_{\mathrm{BSE}} = \begin{bmatrix} R & C \\ -C^* & -R^* \end{bmatrix}$$

We need a parallel diagonalization routine for matrix of this form.

Want cost to be similar to the Tamm-Dancoff approximation:

$$H_{\mathrm{BSE}}^{\mathrm{TDA}} = [R]$$

Efficient Full BSE Calculations

- Challenge: diagonalize non-Hermitian matrix (2x the rank):
  - Efficiently
  - In parallel
  - Preserving structure of the solutions

- Solution: **new solver** written by Meiyue Shao and Chao Yang.

- Si (real matrix), nmat=24 000, #PEs=128

	TDA	Our solver	Generic solver
Time (s)	78.163	243.481 (3.1x)	1198.535 (15.3 x)

- Naphthalene (complex matrix), nmat=8 000, #PEs=72

	TDA	Our solver	Generic solver
Time (s)	41.088	259.309 (6.3x)	593.593 (14.4x)

Meiyue Shao, Felipe Homrich da Jornada, Chao Yang, Jack Deslippe, Steven G. Louie. Submitted to Journal Elsevier Linear Algebra and its Applications (arXiv preprint arXiv:1501.03830)

Example: Benzene

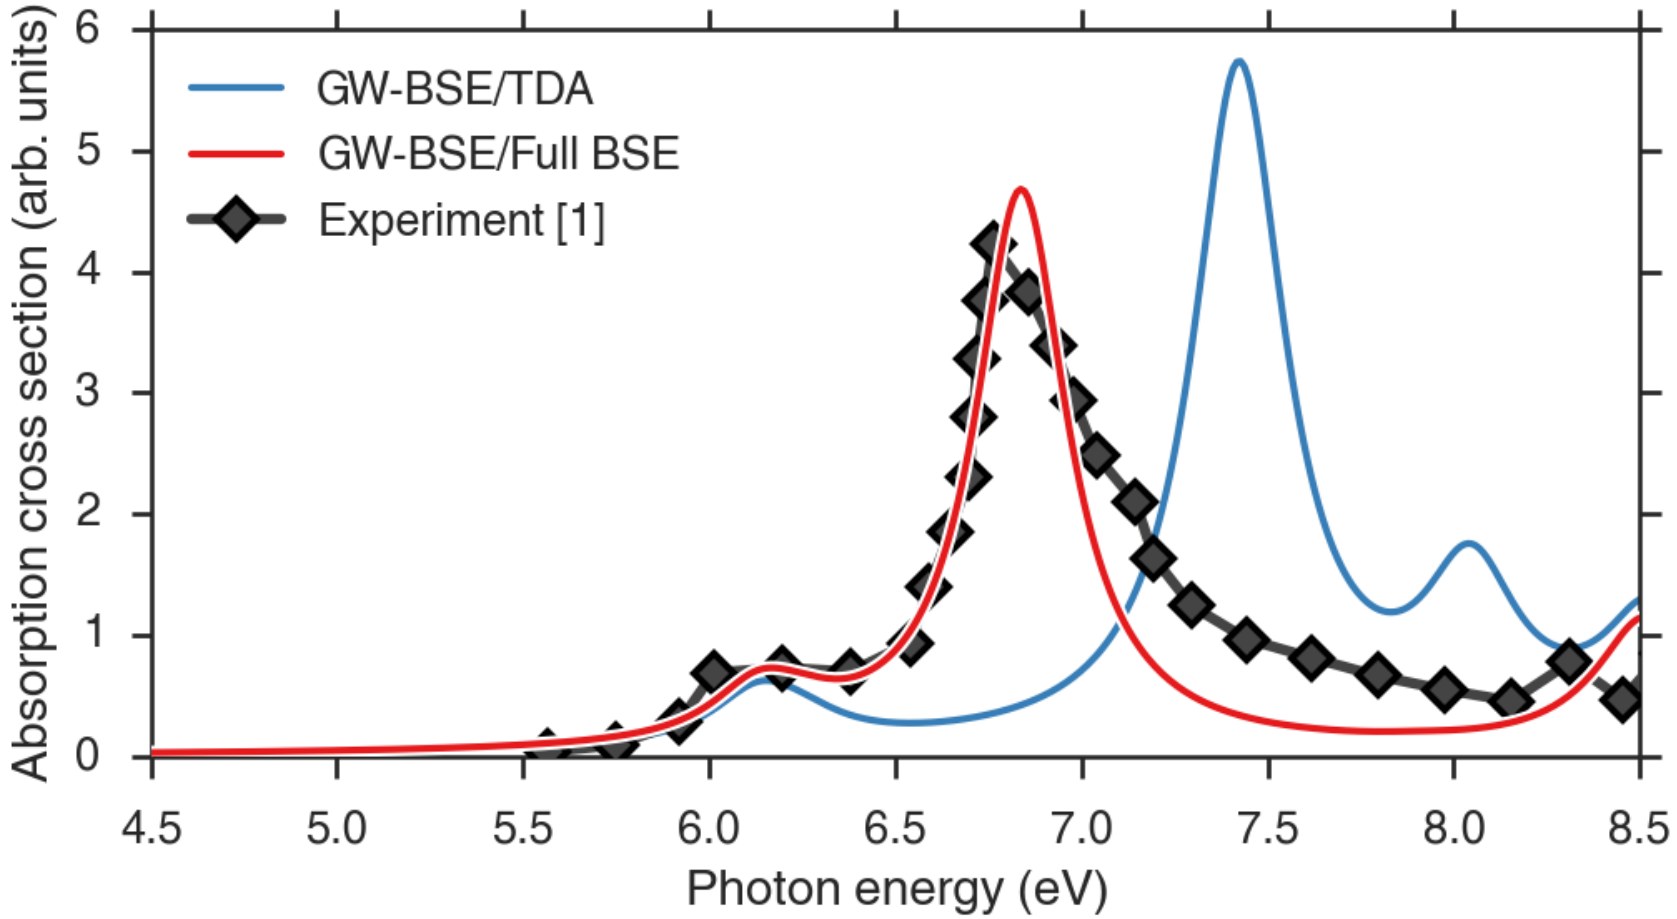

GW - Towards Informatics

- Closed-shell, small molecules
- Characterized experimentally
- Include a variety of elements

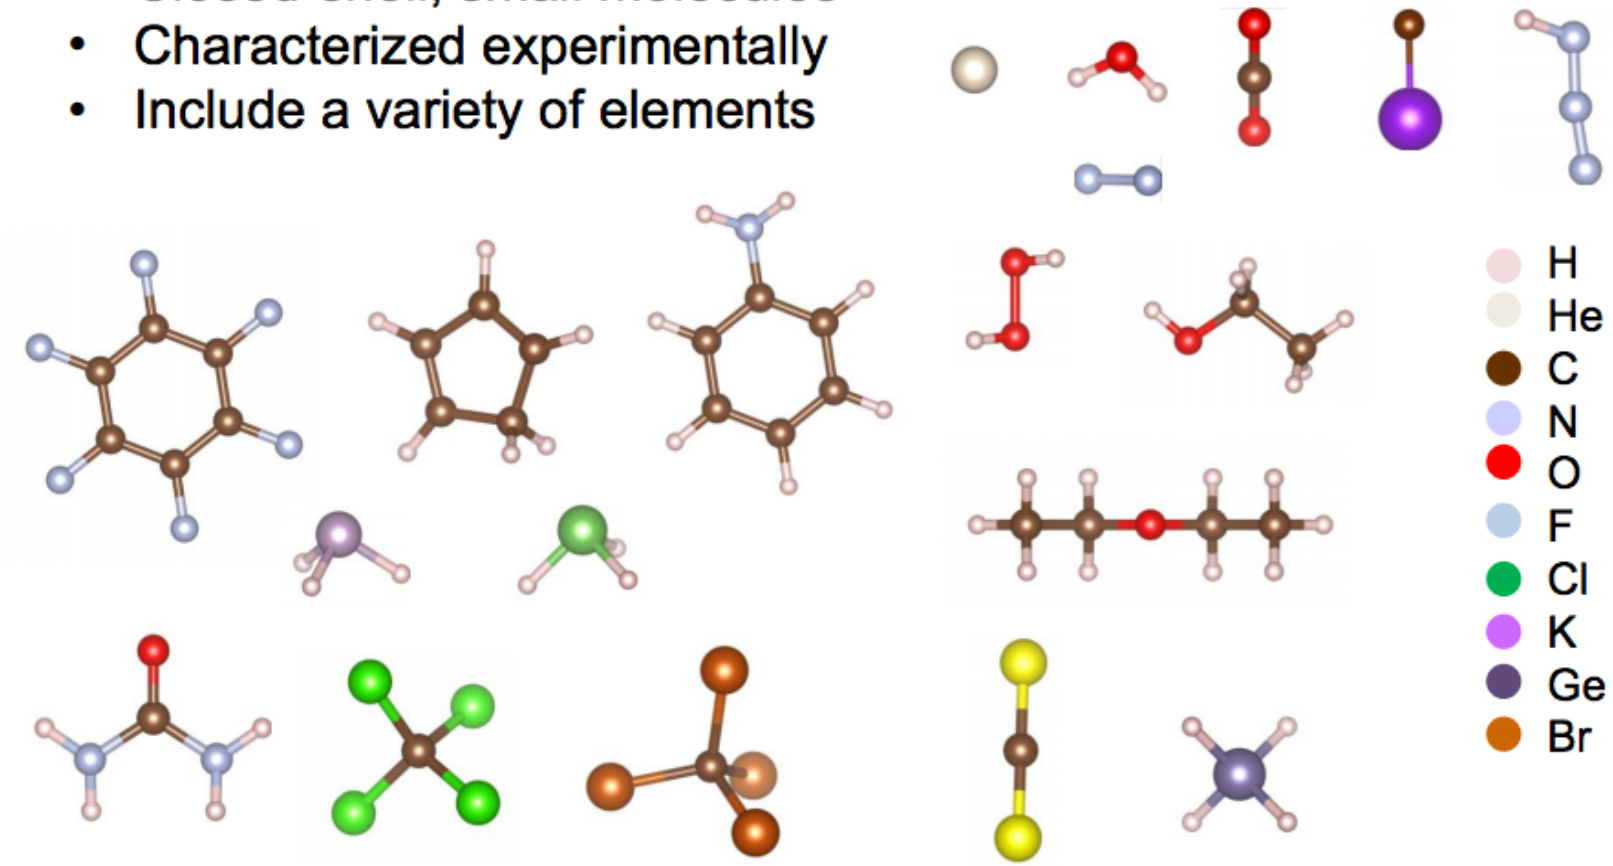

GW - Towards Informatics

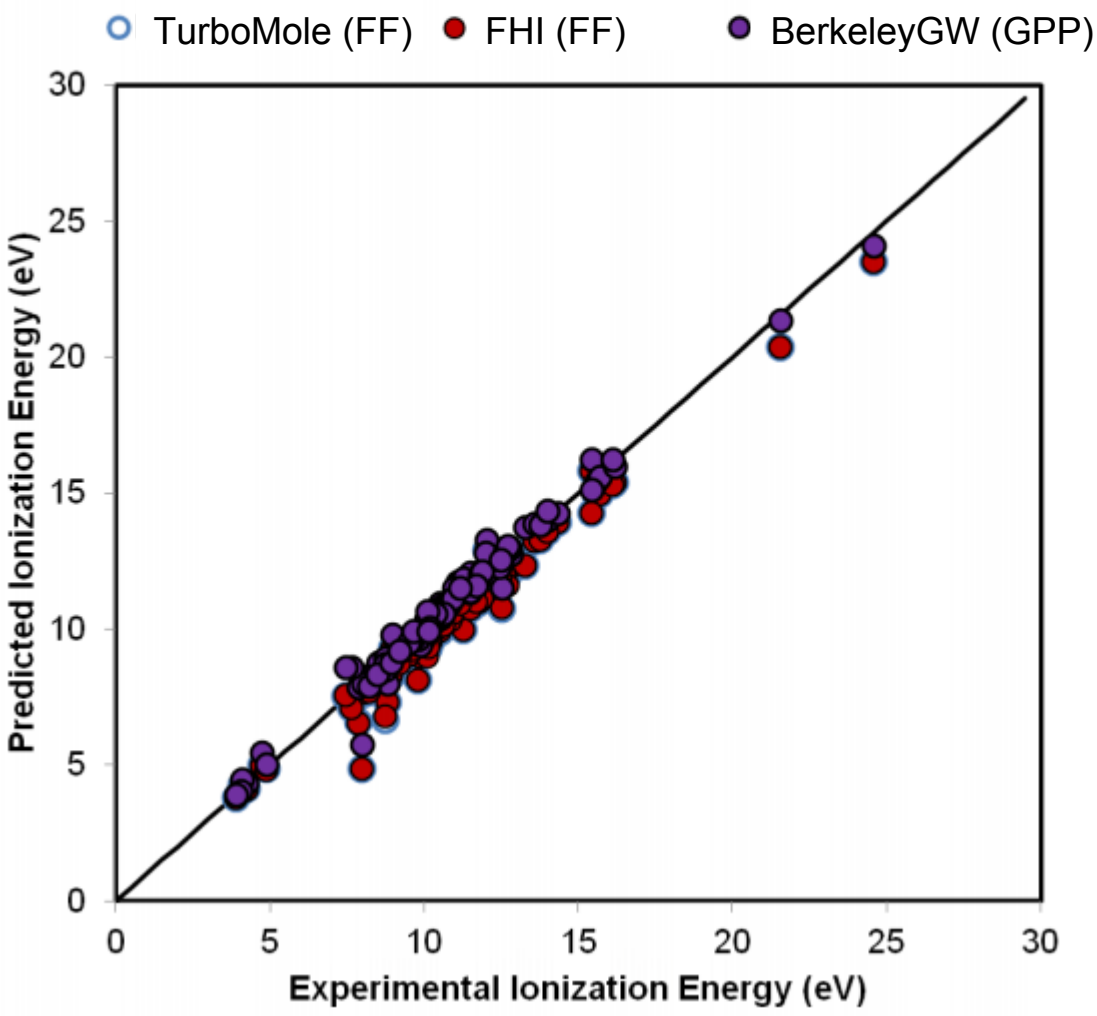

Michiel van Setten, … Jack Deslippe, Steven Louie, Chao Yang, Jeffrey Neaton, Ferdinand… , " Submitted Journal of Chemical Theory and Computation

GW - Towards Informatics

Full-Frequency Calculations within G0W0
Systematically underestimate band gaps.

GPP gives generally very good energies for informatics based approaches.

Poles in FF G0W0 ~

E_v_lda - E_gap
E_c_lda + E_gap

Poles in FF GW0 ~

E_v_QP - E_gap
E_c_QP - E_gap

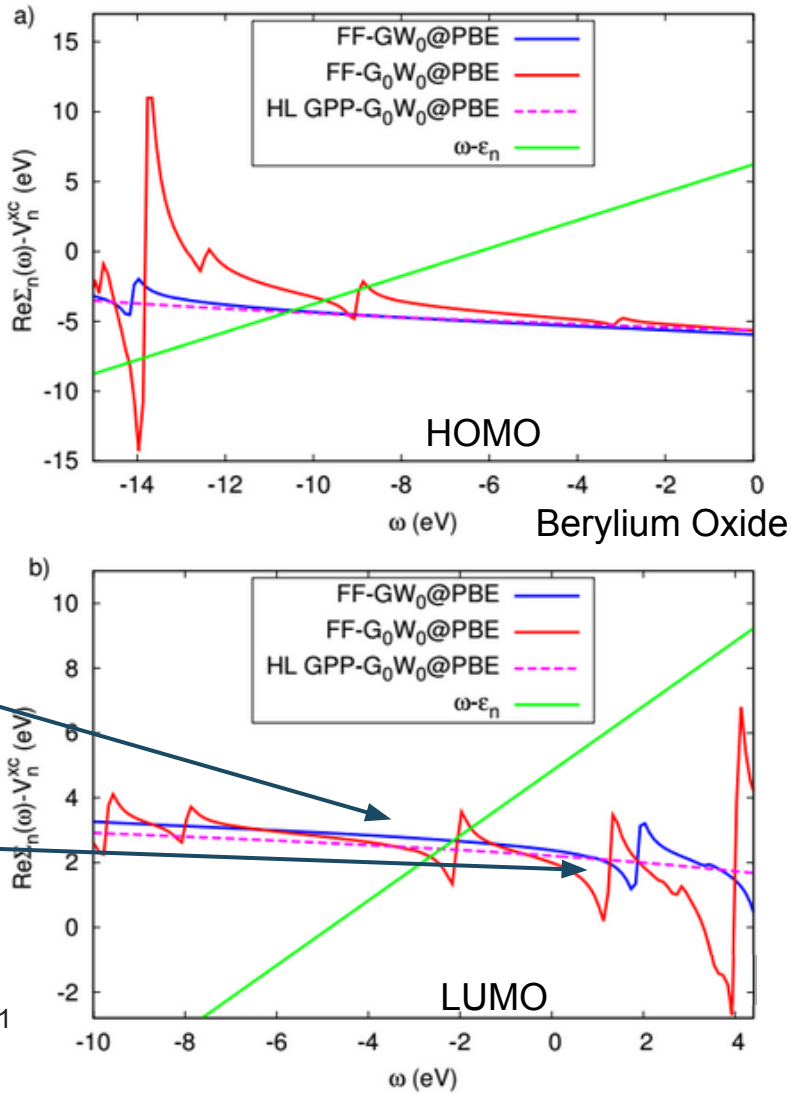

Lischner, Johannes … Jack Deslippe, J.B. Neaton, S.G. Louie. *Physical Review B* 90.11 (2014): 115130.

Conclusions

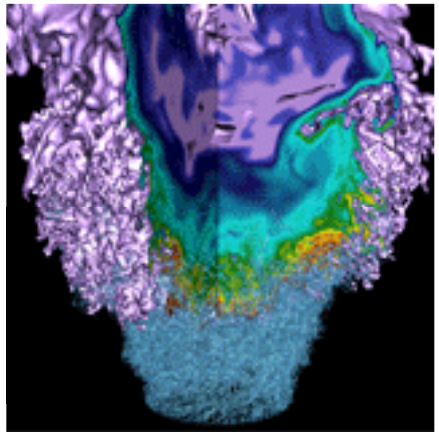

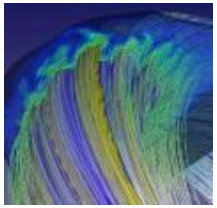

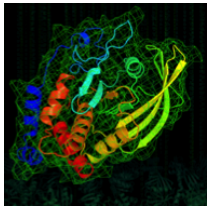

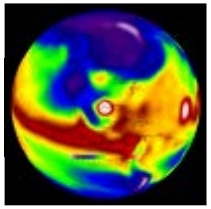

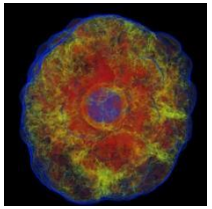

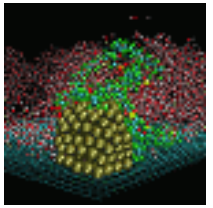

Conclusions:

BerkeleyGW routinely run on systems with hundreds of atoms.

Order 1000 atoms possible with DOE HPC resources like Edison and Mira.

BerkeleyGW is ahead of the pack of GW codes.

- Foundry users find Full-Frequency calculation of MgO takes hours with BerkeleyGW, weeks with abinit
- Yambo celebrates passing the 1000 CPU core in 2014, BerkeleyGW commonly run on 10-100x that scale.
- VASP GW limited in size by memory requirements of (G, G') matrices.

C60 Pentacene Interface under investigation by Neaton Group

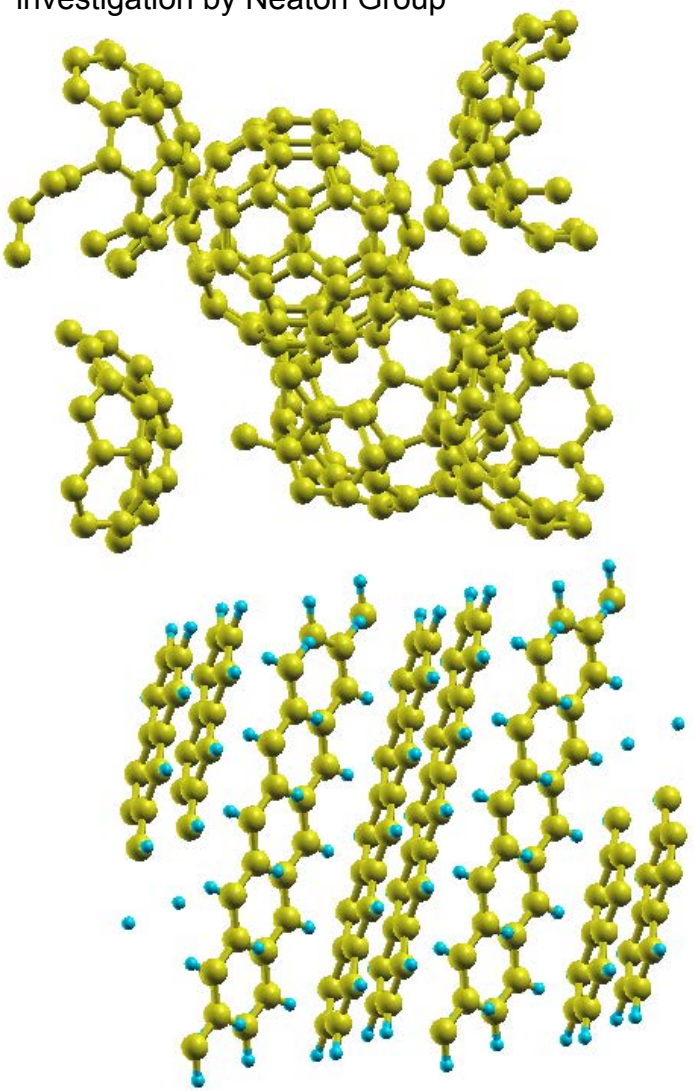

Our Posters

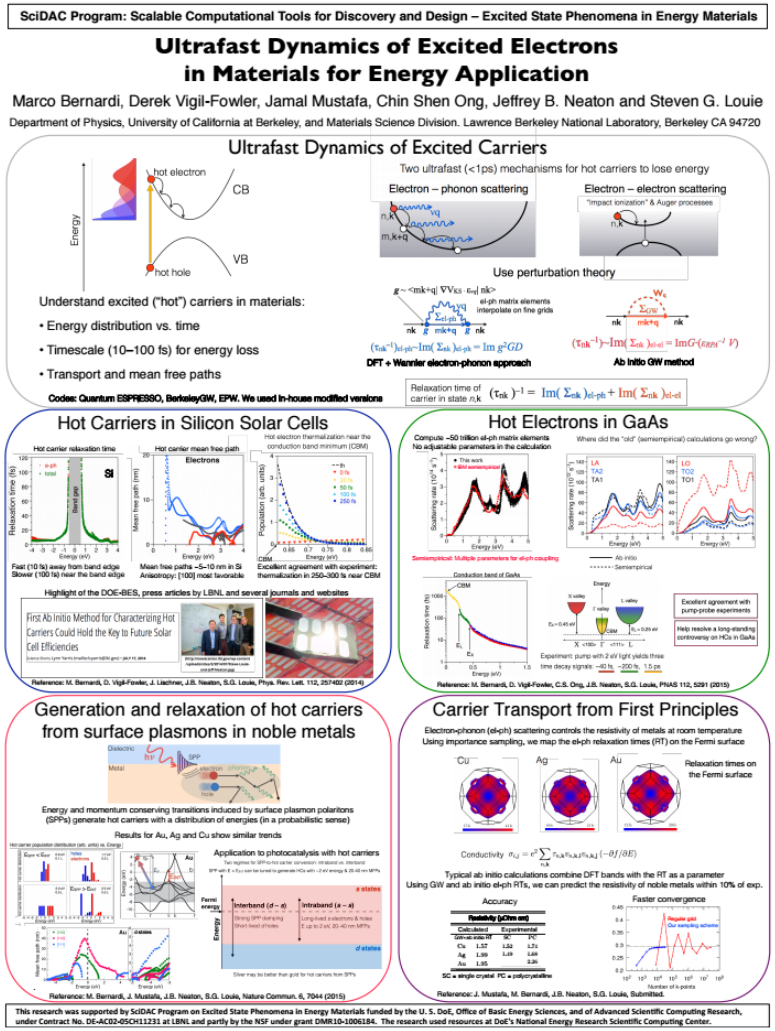

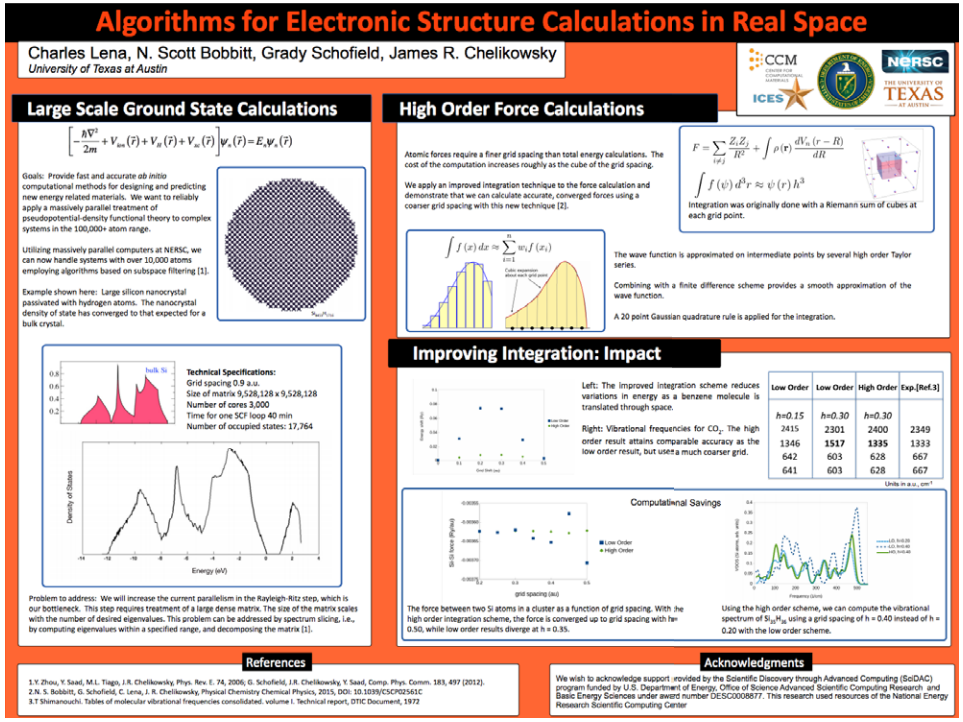

U.S. DEPARTMENT OF ENERGY | Office of Science

BERKELEY LAB
Lawrence Berkeley National Laboratory

Acknowledgements

This work could not have been done without SCIDAC!

SCIDAC Program on Excited State Phenomena in Energy Materials funded by the U. S. Department of Energy, Office of Basic Energy Sciences and of Advanced Scientific Computing Research, under Contract No. DE-AC02-05CH11231 at Lawrence Berkeley National Laboratory and under Award No. DESC0008877 at University of Texas, Austin

Extra Slides

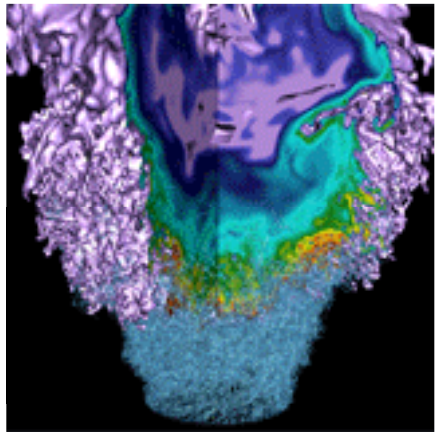

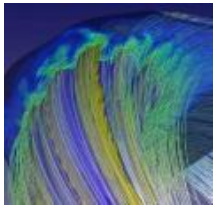

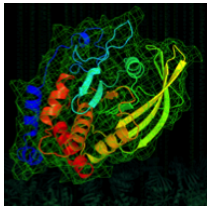

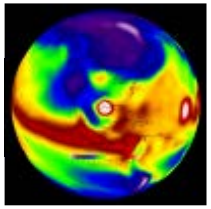

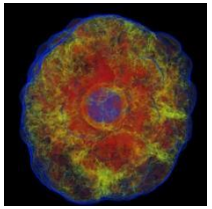

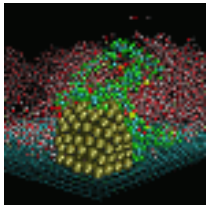

NERSC

For a typical GW calculation, the LDA starting point is sufficient:

$$E_n^{QP} \approx \langle \psi_n^{MF} | H_{Hartree} | \psi_n^{MF} \rangle + \langle \psi_n^{MF} | \Sigma | \psi_n^{MF} \rangle$$

Notable exceptions - Silane:

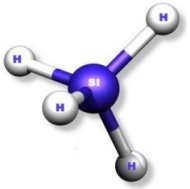

M. Rohlfing and S.G. Louie Phys. Rev. B 62 4927 (2000).

Using the Static Limit to Improve GW

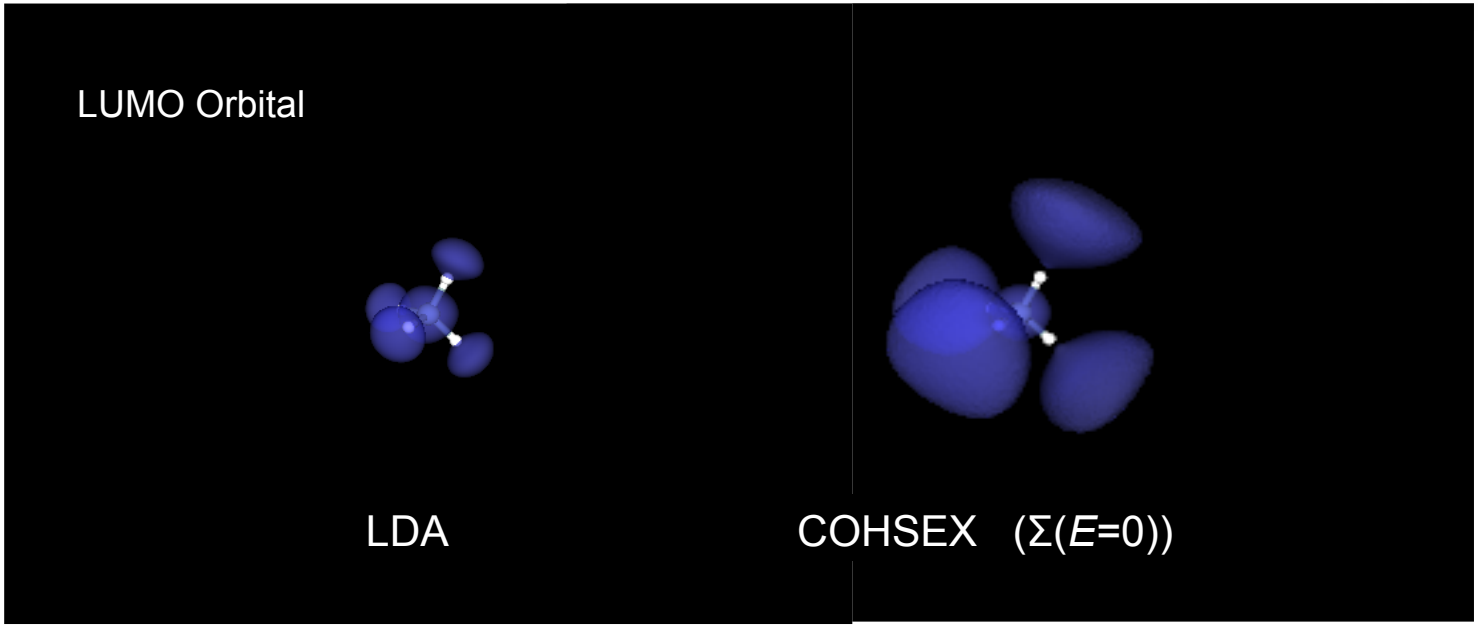

	LDA	LDA+GW	CSX	CSX+GW
HOMO	-8.52	-12.80	-13.2	-12.80
LUMO	-0.465	1.02	0.1	0.29
QP gap	8.06	13.82	13.3	13.10

Jain, M., Deslippe, J...., Chelikowsky, J. R., & Louie, S. G. (2014). *Physical Review B*,*90*(11), 115148.

Using the Static Limit to Improve GW

$$\langle n\mathbf{k}|\Sigma^{\infty}_{\mathrm{CH}}(\mathbf{r},\mathbf{r}';E)\,|n'\mathbf{k}\rangle =$$

$$\langle n\mathbf{k}|\,\Sigma^{N}_{\mathrm{CH}}(\mathbf{r},\mathbf{r}';E)\,|n'\mathbf{k}\rangle + \frac{1}{2}\Big(\,\langle n\mathbf{k}|\,\Sigma^{Coh/\infty}_{\mathrm{CH}}(\mathbf{r},\mathbf{r}')\,|n'\mathbf{k}\rangle - \langle n\mathbf{k}|\,\Sigma^{Coh/N}_{\mathrm{CH}}(\mathbf{r},\mathbf{r}')\,|n'\mathbf{k}\rangle\,\Big).$$

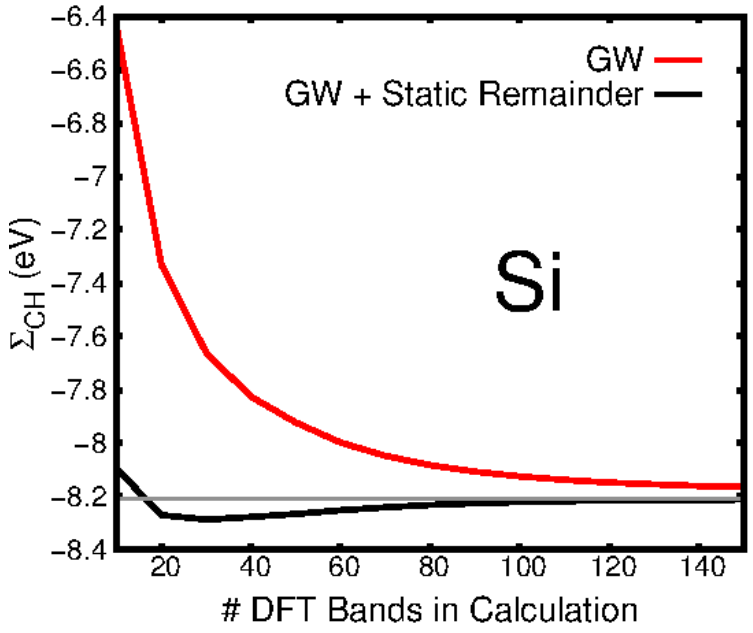

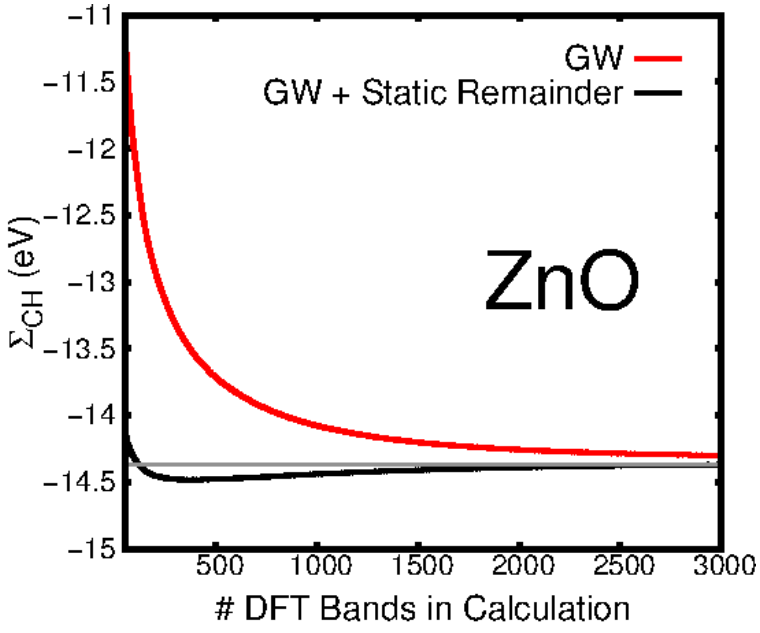

Deslippe, Jack, … S.G. Louie et al. *Physical Review B* 87.16 (2013): 165124. First done in work of Tiago and Chelikowsky.

Finite Momentum Excitons

- **Q**→0
  - Optical absorption ~ Im ε_M
- **Q** ≠0
  - Optically inactive
  - Energy loss ~ -Im ε^{-1}
  - Important for exciton dynamics, relaxation, ...

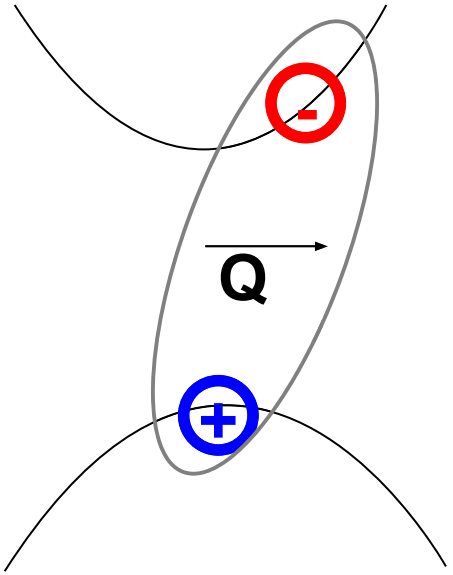

Finite Momentum Excitons

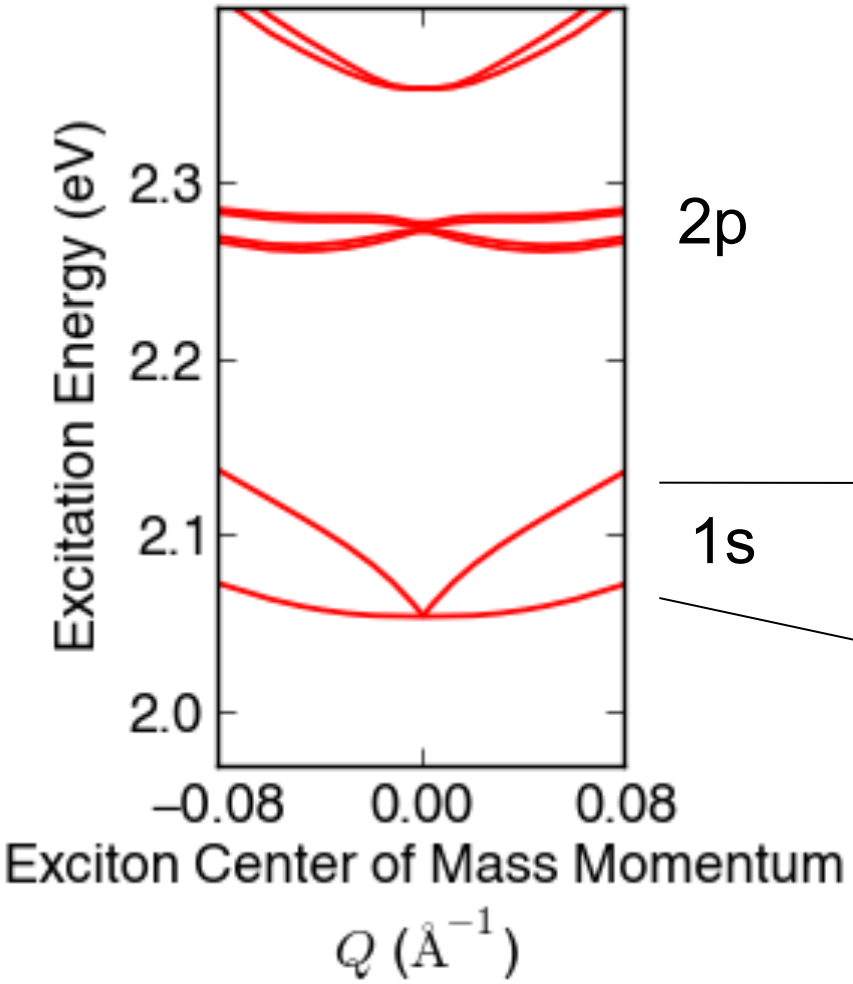

- Ab-initio exciton bandstructure of MoS2
- Access novel physics:
  - Nonanalytic dispersion due to 2D Coulomb interaction
  - Valley quantum phase with valley pseudospin winding number = 2

→ Couples to longitudinal electric field

→ Couples to transverse electric field

D.Y. Qiu, T. Cao, and S.G. Louie, arxiv:1311.0963[cond-mat.mtrl-sci] (2015).

GW Parallelism

GW is an ideal case for Many-Core / Exascale. Many levels of parallelism can be exploited. Ideal for many-core.

$$\epsilon^{r/a}_{\mathbf{GG'}}(\mathbf{q},E) = \delta_{\mathbf{GG'}} - v(\mathbf{q}+\mathbf{G}) \times \sum_{n}^{occ}\sum_{n'}^{emp}\sum_{\mathbf{k}} M^{*}_{nn'}(\mathbf{k},\mathbf{q},\mathbf{G}) M_{nn'}(\mathbf{k},\mathbf{q},\mathbf{G'}) \times \frac{1}{2}\left[\frac{1}{E_{n\mathbf{k}+\mathbf{q}} - E_{n'\mathbf{k}} - E \mp i\delta} + \frac{1}{E_{n\mathbf{k}+\mathbf{q}} - E_{n'\mathbf{k}} + E \pm i\delta}\right]$$

Inner Dimension of Hybrid MPI-OpenMP ZGEMM (OpenMP new to 1.2)
Outer Dimension of Hybrid MPI-OpenMP ZGEMM (OpenMP new to 1.2)
MPI Group Level Parallelization over Frequencies (New to 1.2)
Trivially parallelization over q points

Subsampling

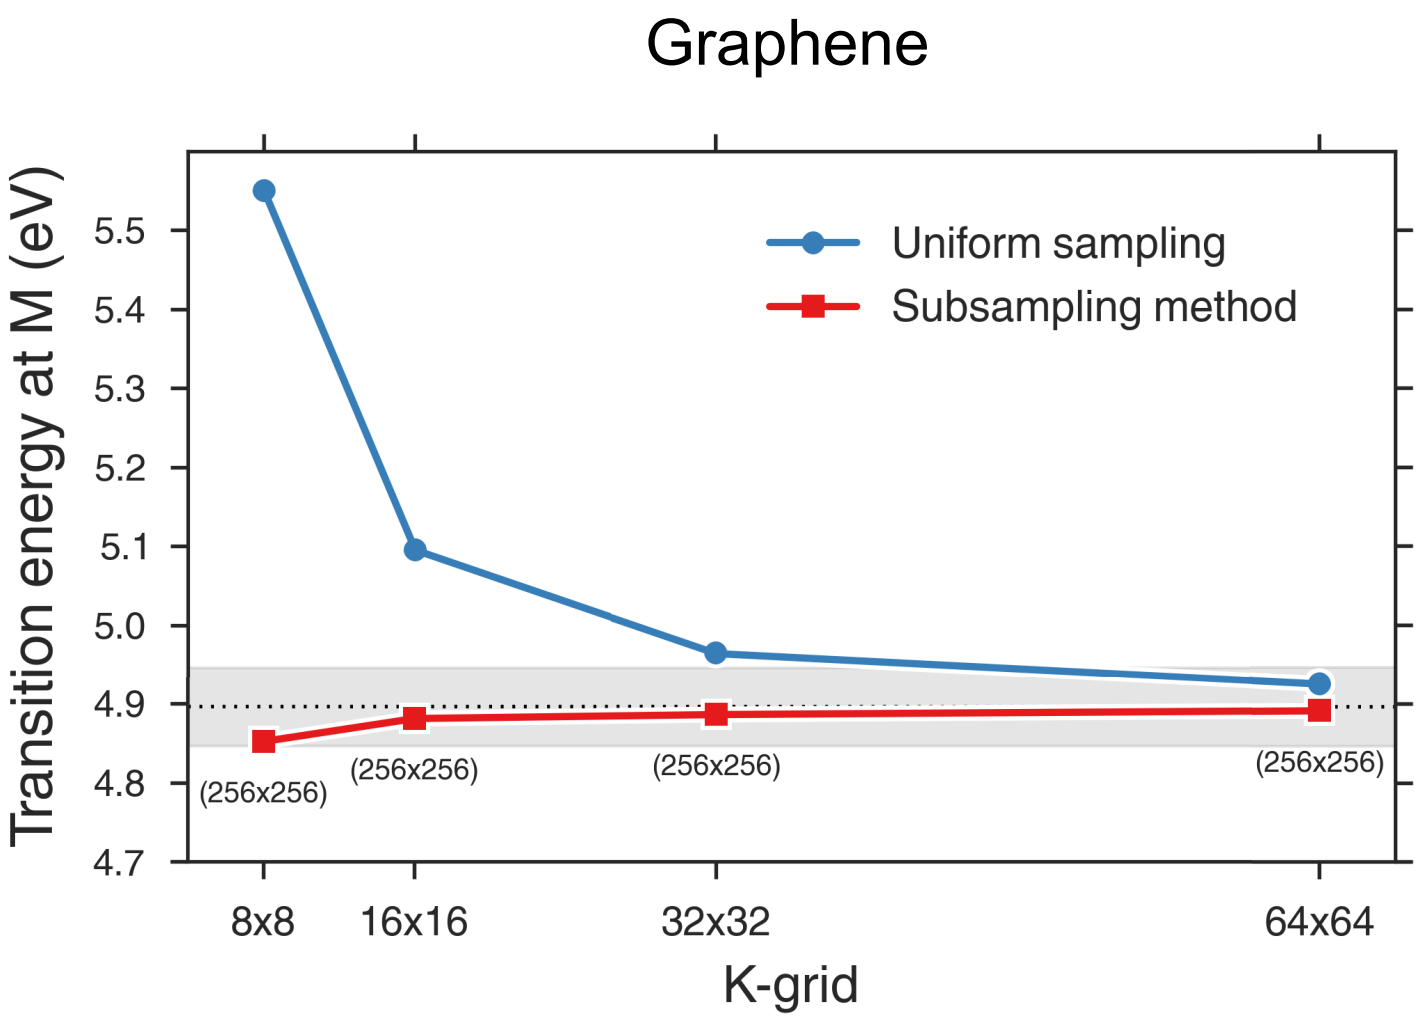

Simplified Final Loop Structure

```
!$OMP DO reduction(+:achtemp)
  do my_igp = 1, ngpown

    ...

    do iw=1,3

      scht=0D0
      wxt = wx_array(iw)

      do ig = 1, ncouls

        !if (abs(wtilde_array(ig,my_igp) * eps(ig,my_igp)) .lt. TOL) cycle

        wdiff = wxt - wtilde_array(ig,my_igp)
        delw = wtilde_array(ig,my_igp) / wdiff

        ...

        scha(ig) = mygpvar1 * aqsntemp(ig) * delw * eps(ig,my_igp)

        scht = scht + scha(ig)

      enddo ! loop over g

      sch_array(iw) = sch_array(iw) + 0.5D0*scht

    enddo

    achtemp(:) = achtemp(:) + sch_array(:) * vcoul(my_igp)

  enddo
```

ngpown typically in 100's to 1000s. Good for many threads.

Original inner loop. Too small to vectorize!

ncouls typically in 1000s - 10,000s. Good for vectorization. Don't have to worry much about memory. alignment.

Attempt to save work breaks vectorization and makes code slower.

1. Compute via *nxn'* FFTs (N^3 Step. Big Prefactor.):

$$M_{nn'}(\mathbf{k},\mathbf{q},\mathbf{G}) = \langle n\mathbf{k}+\mathbf{q}|\, e^{i(\mathbf{q}+\mathbf{G})\cdot\mathbf{r}}\, |n'\mathbf{k}\rangle$$

$$M_{nn'}(\mathbf{k},\mathbf{q},\{\mathbf{G}\}) = FFT^{-1}\left(\phi_{n,\mathbf{k}+\mathbf{q}}(\mathbf{r}) * \phi^*_{n',\mathbf{k}}(\mathbf{r})\right)$$

2. Compute sum via large ZGEMM (N^4 Step. Small Prefactor. All to All Communication Done):

$$\chi_{\mathbf{GG'}}(\mathbf{q};0) = \mathbf{M}(\mathbf{G},\mathbf{q},(n,n',\mathbf{k}))\cdot\mathbf{M}^T(\mathbf{G}',\mathbf{q}(n,n',\mathbf{k}))$$

Where, $\mathbf{M}(\mathbf{G},\mathbf{q},(n,n',\mathbf{k})) = M_{nn'}(\mathbf{k},\mathbf{q},\mathbf{G})\cdot\dfrac{1}{\sqrt{E_{n\mathbf{k}+\mathbf{q}}-E_{n'\mathbf{k}}}}$

3. Matrix Inversion. ScaLAPACK

MPI Scaling of Epsilon Code:

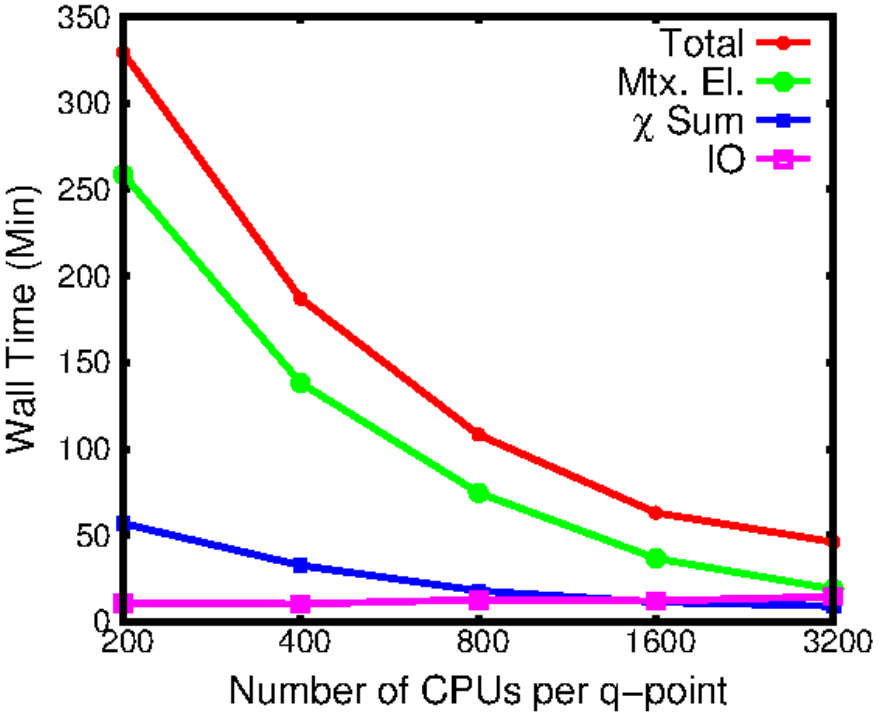

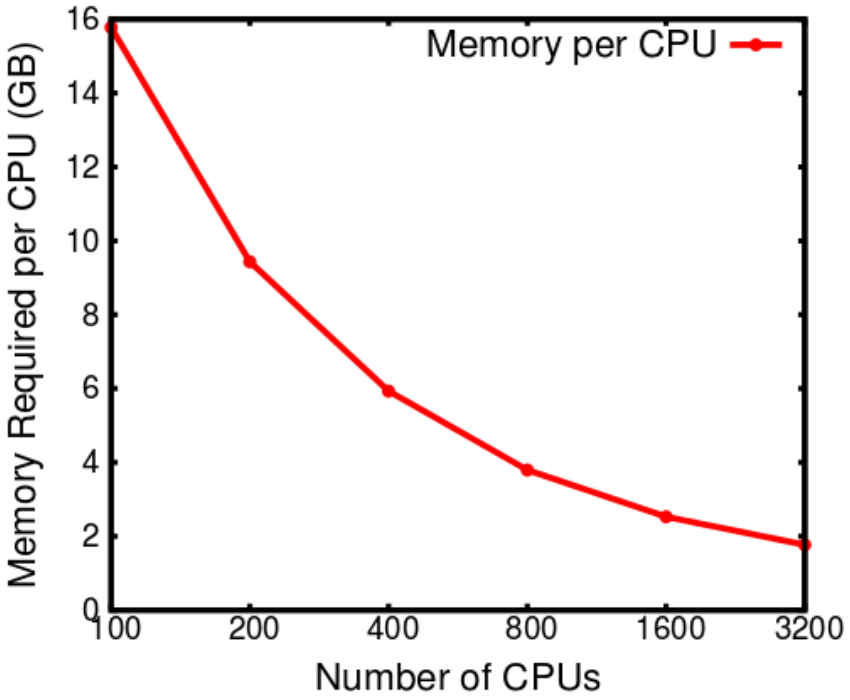

(Sigma GPP Option)

4. Manual loop reductions to compute sum for self-energy.

N^3 x <number of bands of interest>

$$\langle n\mathbf{k}|\,\Sigma_{\mathrm{SX}}(E)\,|n'\mathbf{k}\rangle = -\sum_{n''}^{\mathrm{occ}}\sum_{\mathbf{q}\mathbf{G}\mathbf{G}'} M^*_{n''n}(\mathbf{k},-\mathbf{q},-\mathbf{G})M_{n''n'}(\mathbf{k},-\mathbf{q},-\mathbf{G}') \times\left[\delta_{\mathbf{G}\mathbf{G}'}+\frac{\Omega^2_{\mathbf{G}\mathbf{G}'}(\mathbf{q})\,(1-i\tan\phi_{\mathbf{G}\mathbf{G}'}(\mathbf{q}))}{(E-E_{n''\mathbf{k}-\mathbf{q}})^2-\tilde{\omega}^2_{\mathbf{G}\mathbf{G}'}(\mathbf{q})}\right] v(\mathbf{q}+\mathbf{G}')$$

$$\langle n\mathbf{k}|\,\Sigma_{\mathrm{CH}}(E)\,|n'\mathbf{k}\rangle = \frac{1}{2}\sum_{n''}\sum_{\mathbf{q}\mathbf{G}\mathbf{G}'} M^*_{n''n}(\mathbf{k},-\mathbf{q},-\mathbf{G})M_{n''n'}(\mathbf{k},-\mathbf{q},-\mathbf{G}') \times\frac{\Omega^2_{\mathbf{G}\mathbf{G}'}(\mathbf{q})\,(1-i\tan\phi_{\mathbf{G}\mathbf{G}'}(\mathbf{q}))}{\tilde{\omega}_{\mathbf{G}\mathbf{G}'}(\mathbf{q})\,(E-E_{n''\mathbf{k}-\mathbf{q}}-\tilde{\omega}_{\mathbf{G}\mathbf{G}'}(\mathbf{q}))}\,v(\mathbf{q}+\mathbf{G}')$$

Things that prevent vectorization in your code

Compilers want to “vectorize” your loops whenever possible. But sometimes they get stumped. Here are a few things that prevent your code from vectorizing:

Loop dependency:

```
do i = 1, n
    a(i) = a(i-1) + b(i)
enddo
```

Task forking:

```
do i = 1, n
    if (a(i) < x) cycle
    if (a(i) > x) …
enddo
```

Application Readiness

- **NERSC partnering with selected projects (~20) to help prepare application codes for Cori**

- **The program will provide:**
  - early access to NERSC-8 hardware and testbed systems
  - special vendor (Cray + Intel) training and optimization sessions
  - NERSC Staff support and training

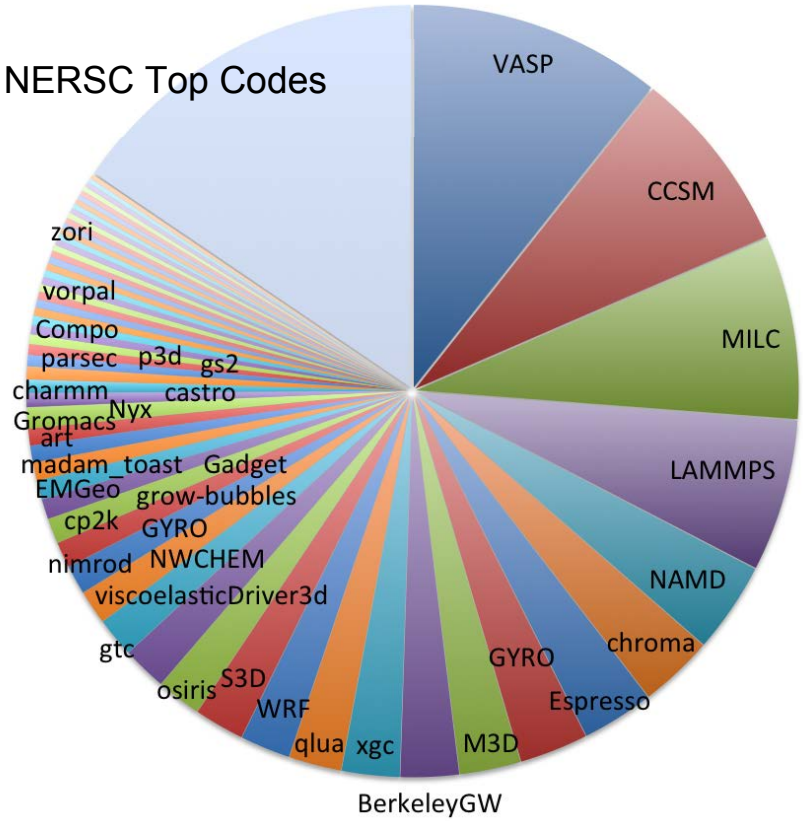

Things that prevent vectorization in your code

Example From NERSC User Group Hackathon - (Astrophysics Transport Code)

```
for (many iterations) {
   … many flops …
   et = exp(outcome1)
   tt = pow(outcome2,3)
   IN = IN * et +tt
}
```

```
for (many iterations) {
   … many flops …
   et(i) = exp(outcome1)
   tt(i) = pow(outcome2,
3)
}
for (many iterations) {
    IN = IN * et(i) + tt(i)
}
```

30% speed up for entire application!

Failure of the MPI-Only Programming Model in BerkeleyGW

★ Big systems require more memory. Cost scales as N_{atm}^2 to store the data.

★ In an MPI GW implementation, in practice, to avoid communication, data is duplicated and **each MPI task has a memory overhead**.

★ On Edison, users sometimes forced to use 1 of 24 available cores, in order to provide MPI tasks with enough memory. **90% of the computing capability is lost**.

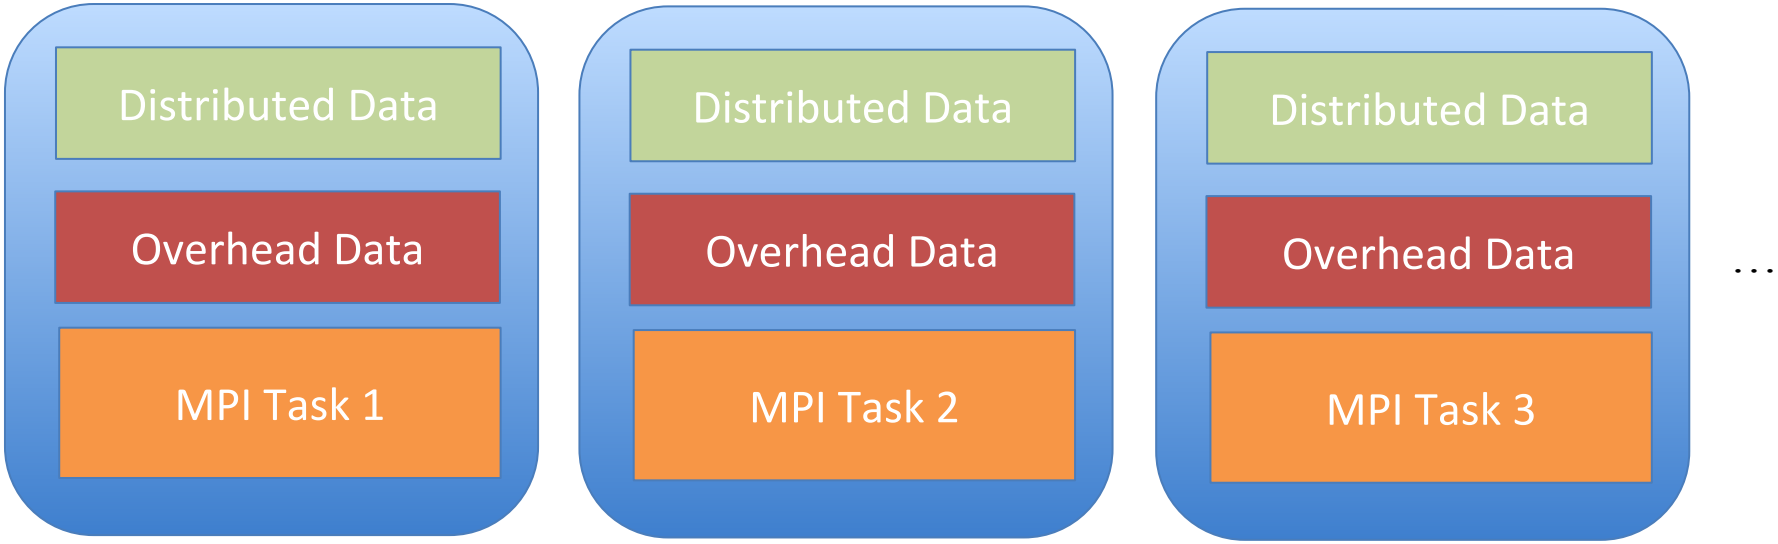

What is GW+BSE

$$|N,S\rangle = \sum_{v}^{hole} \sum_{c}^{elec} A_{vc}^{S} a_{v}^{+} b_{c}^{+} |N,0\rangle + \ldots$$

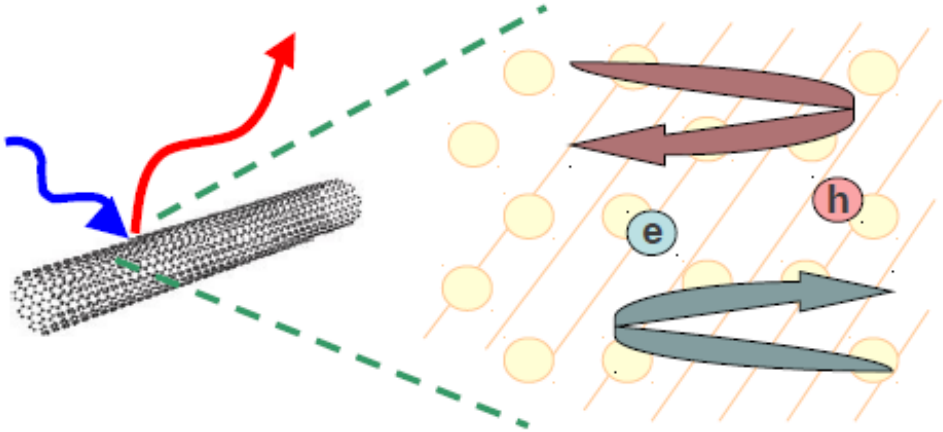

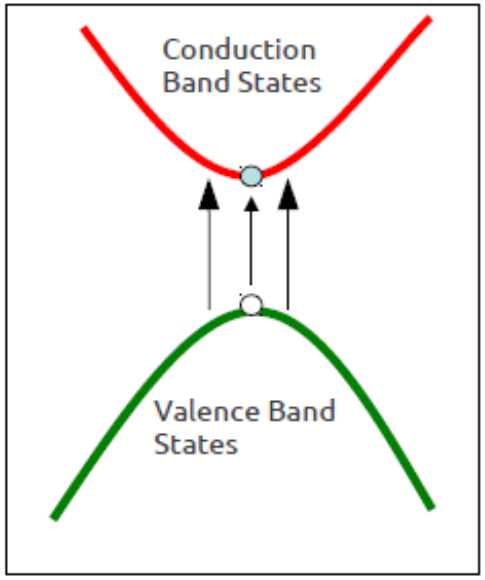

$$\left(E_{ck}^{QP} - E_{vk}^{QP}\right) A^{S}{}_{vck} + \sum_{k'v'c'} \langle vck | K^{eh} | v'c'k' \rangle A^{S}{}_{v'c'k'} = \Omega^{S} A^{S}{}_{vck}$$

Epsilon/Sigma Improvements

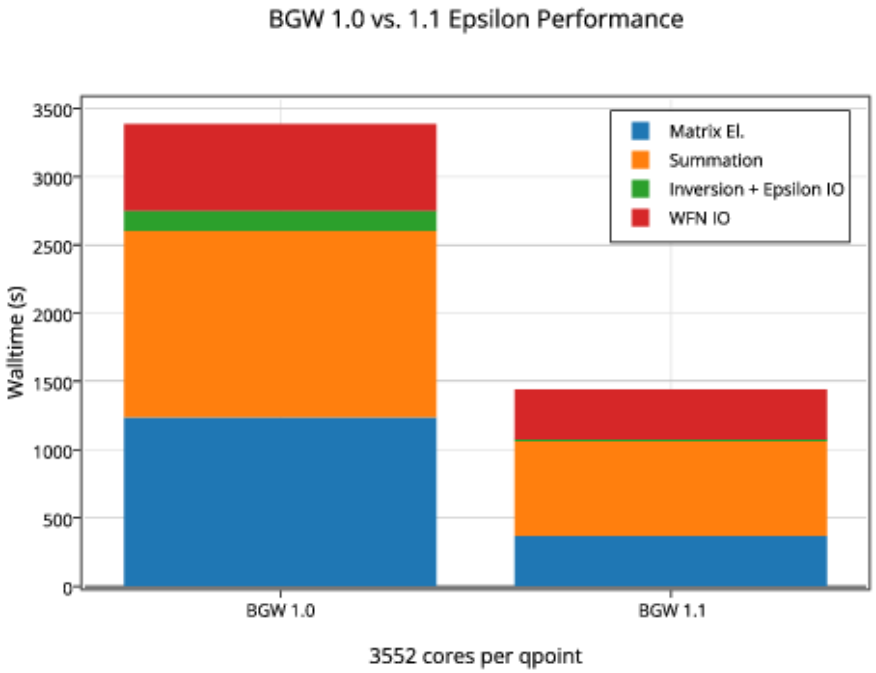

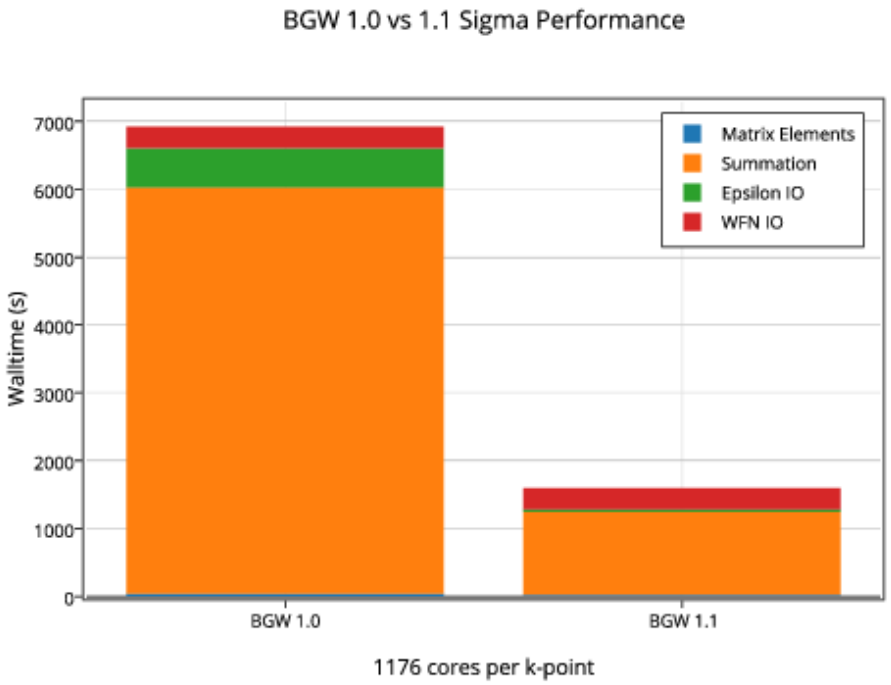

- Performance improvements from:
  - Parallel IO (HDF5)
  - Vectorization
  - Memory-locality improvements
  - FFT Size/Performance Improvements

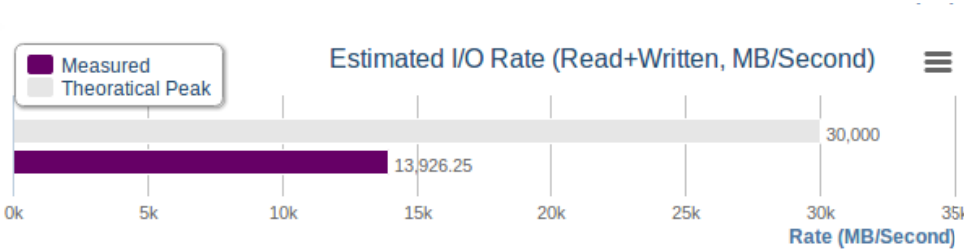

NESAP Participation

Both **PARSEC** and **BerkeleyGW** are included in the top tier of the NERSC Exascale Science Application Program

- **Early access to hardware**
  - Early access on the full Cori system
- **Technical deep dives**
  - Access to Cray and Intel staff on-site
  - Multi-day deep dive ('dungeon' session) with Intel staff at Oregon Campus
- **User Training Sessions**
  - From NERSC, Cray and Intel staff on OpenMP, vectorization, application profiling
  - Knights Landing architectural briefings from Intel

NESAP Advances with Cray, Intel and SUPER

- Worked with Cray, Intel and SUPER* to identify further bottlenecks in BerkeleyGW kernels.

- Added additional layers of cache-blocking to improve locality. Important on Xeon-Phi which lacks L3.

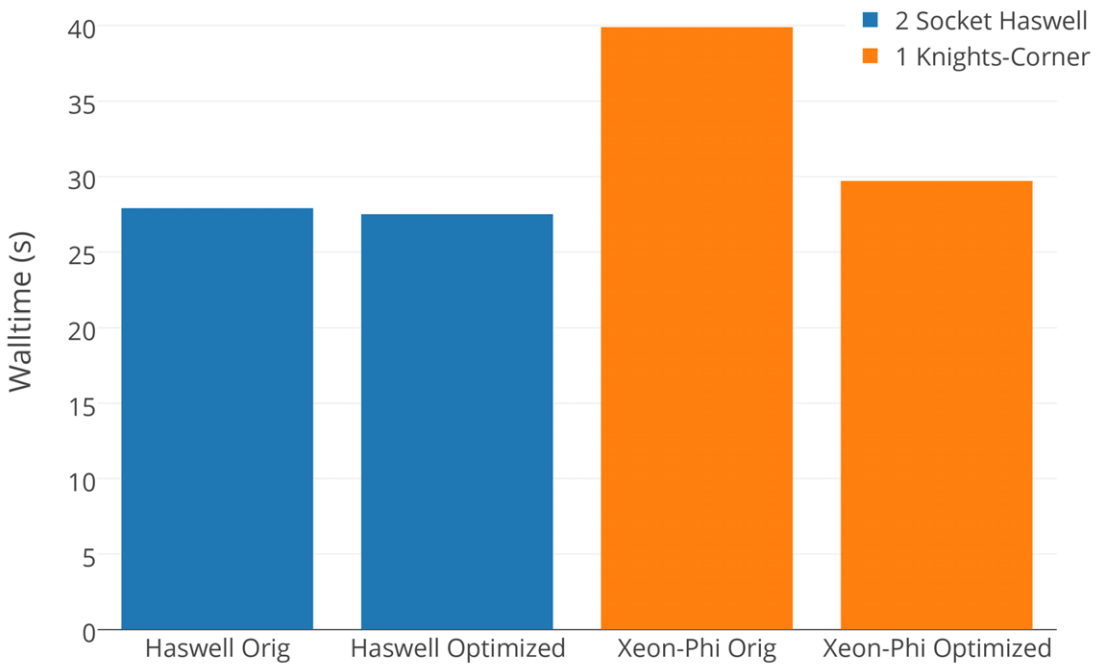

* Collaboration with SUPER SciDAC Institute (J. Tilson, K. Huck, R.Fowler)

Memory Bandwidth Optimizations

What to do if your code is memory bandwidth bound?

1. Try to improve memory locality, cache reuse

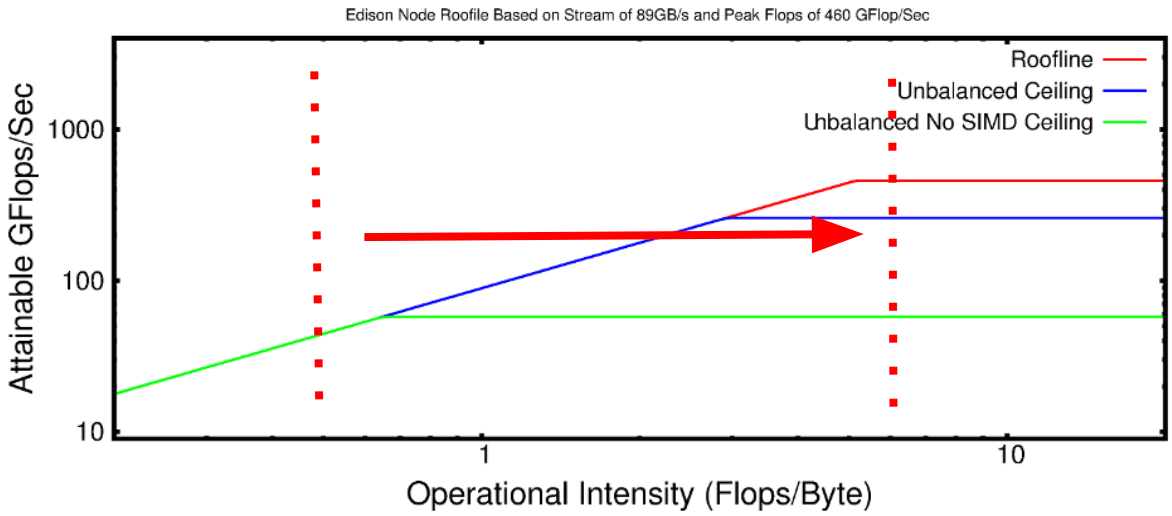

2. Identify arrays leading to high bandwidth usage and make sure they are/will-be allocated in HBM on Cori.