Molecular Foundry

BERKELEY LAB

Large-Scale Computation for Excited-State Phenomena in Complex Energy Materials

Jeffrey B. Neaton
Director, Molecular Foundry
Lawrence Berkeley National Laboratory

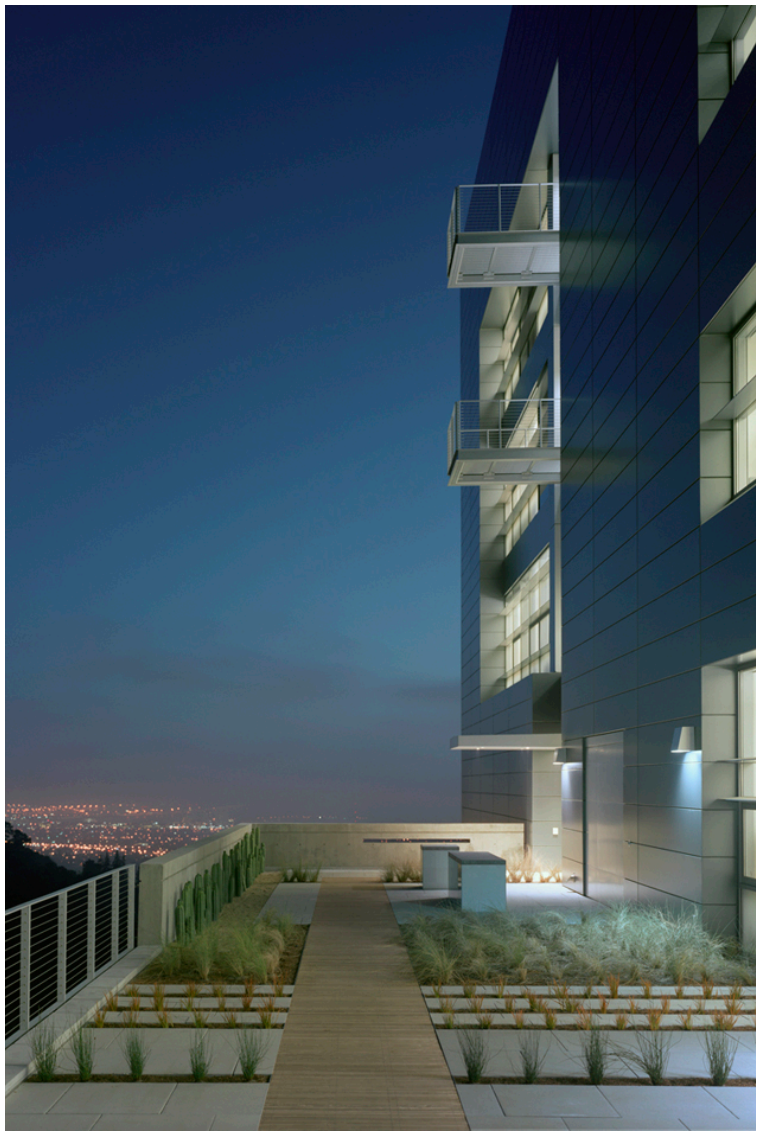

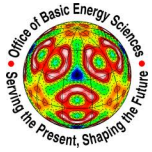

New Materials & Phenomena to Rebalance the Carbon Cycle

Oxide interfaces & epitaxial heterostructures

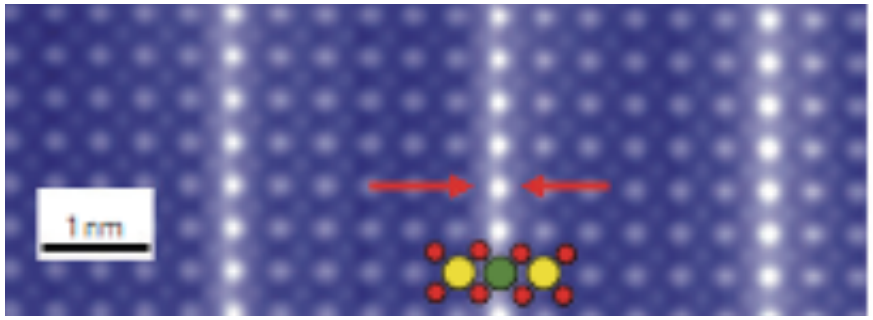

Organic crystals & polymers

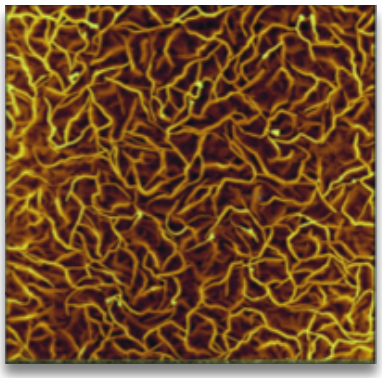

Molecular junctions & hybrid assemblies

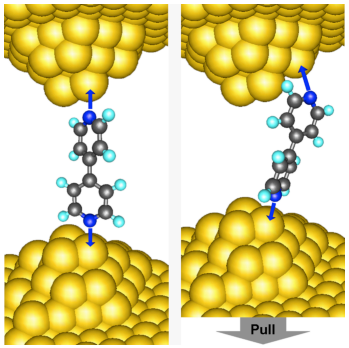

Nanoporous metal-organic frameworks

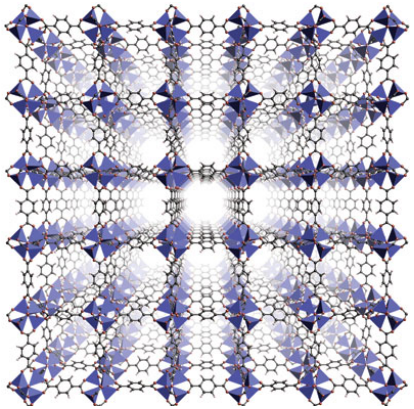

New Materials & Phenomena to Rebalance the Carbon Cycle

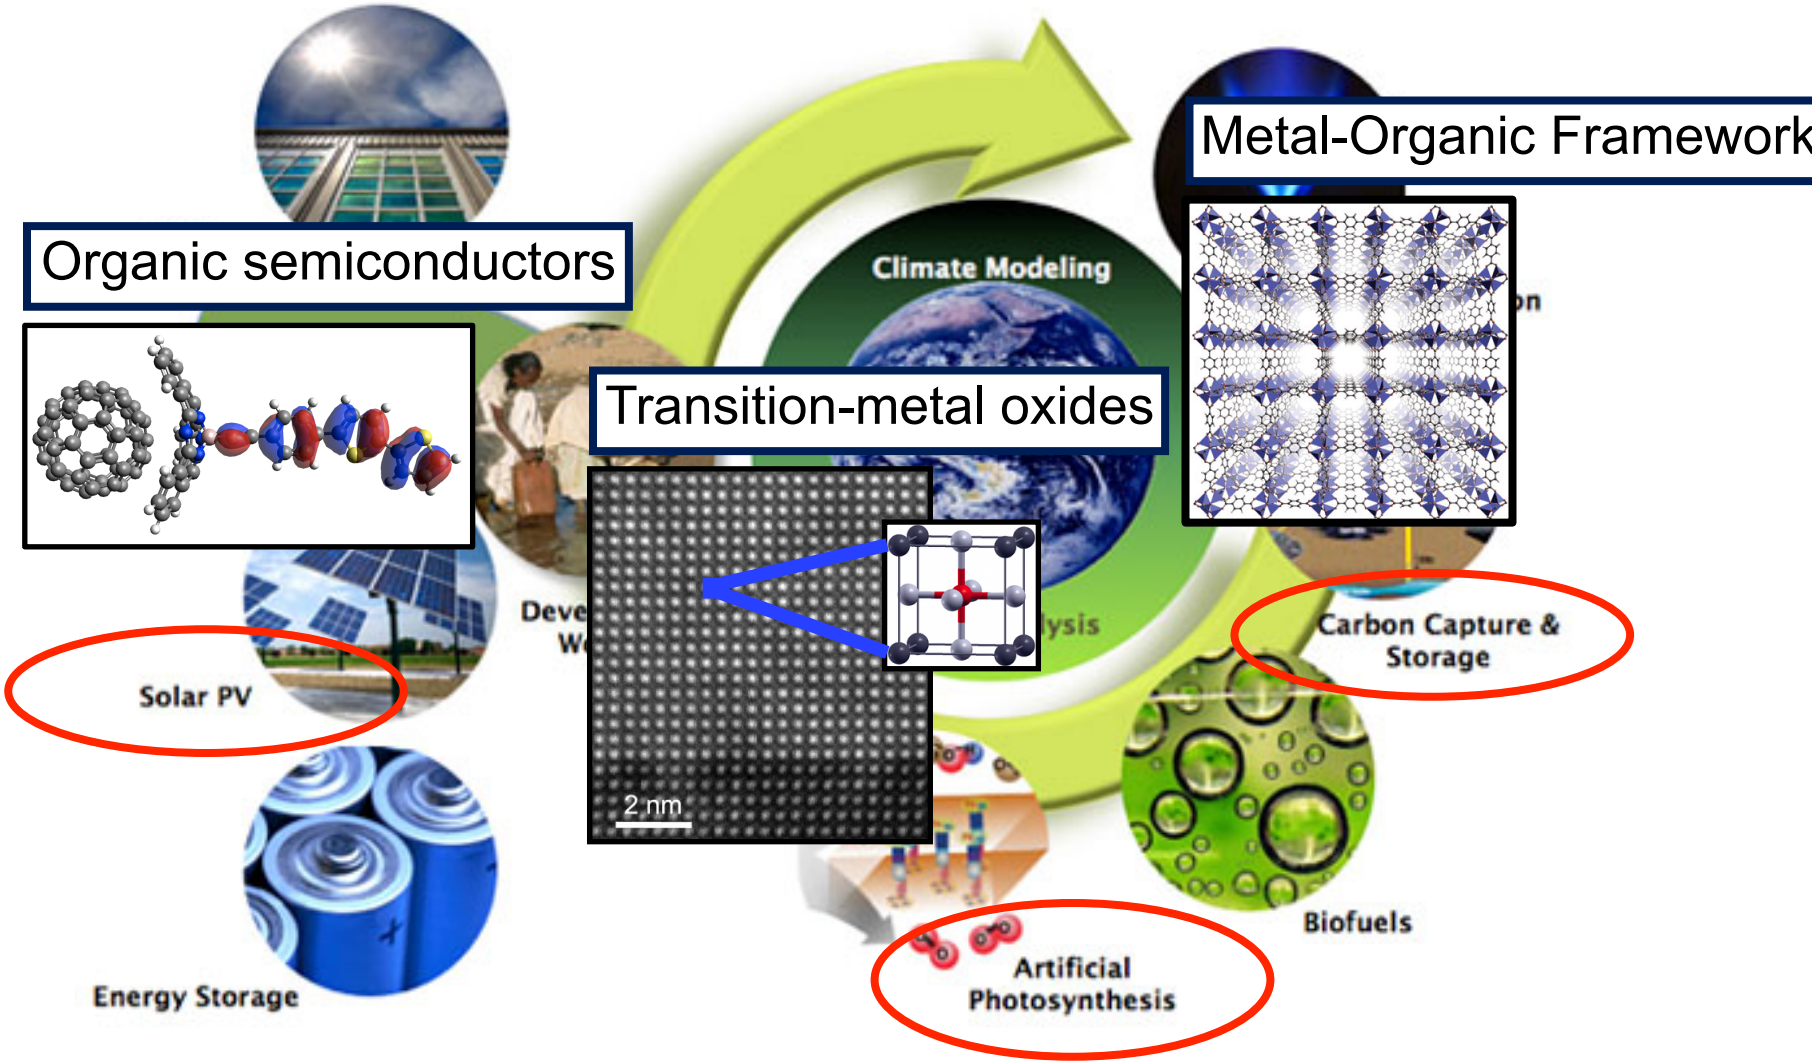

Excited States at Functional Interfaces

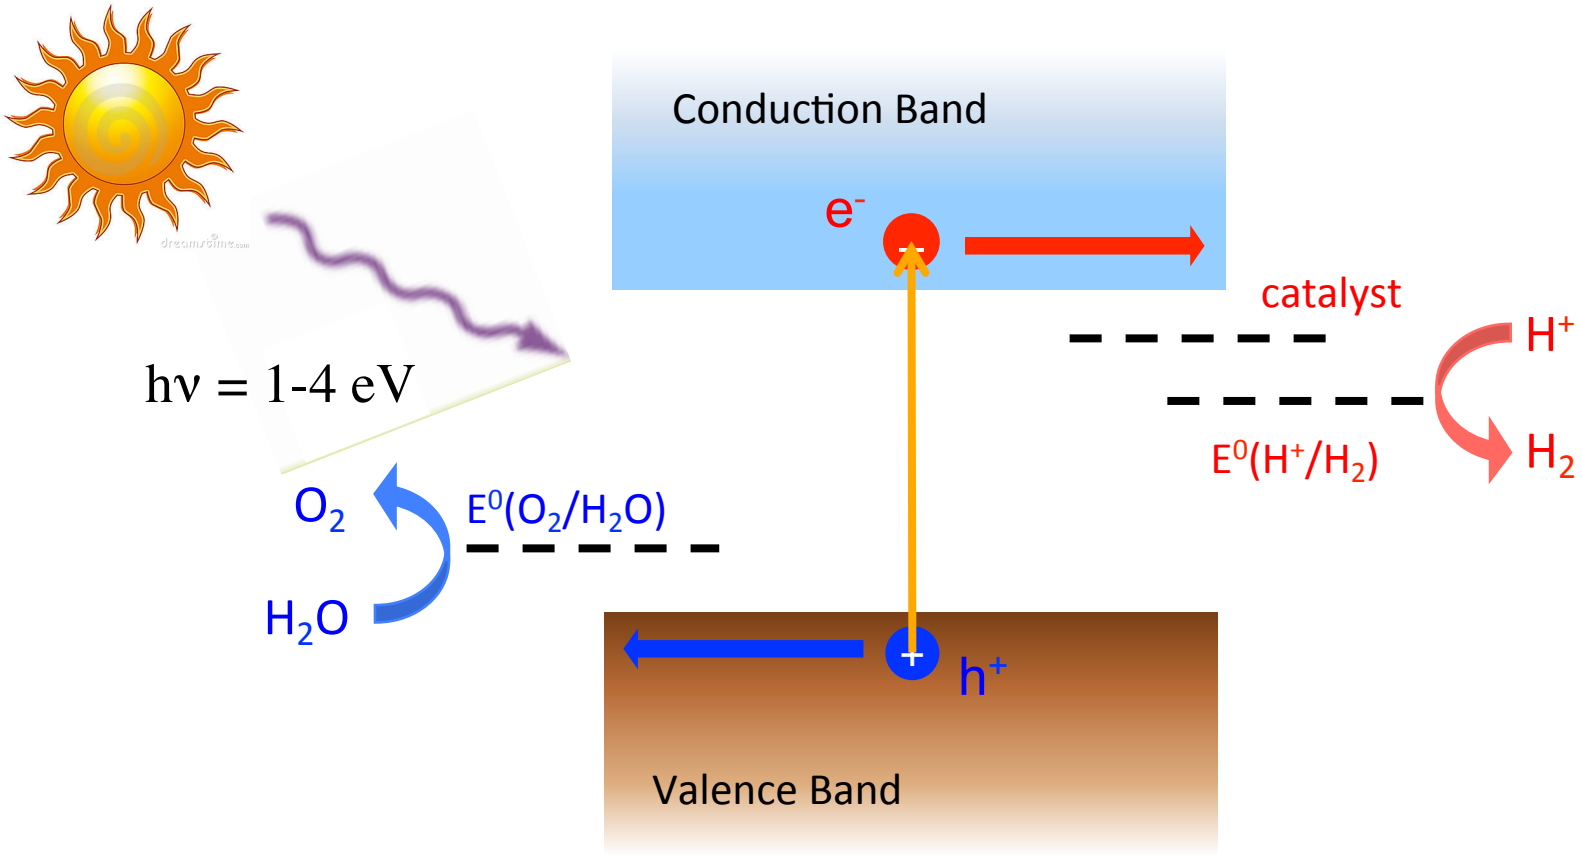

- Importance of band gaps, edges in complex materials
- Need for new methods, algorithms, approximations, intuition

Our Theory & Simulation Efforts are Strongly Connected to other DOE Programs

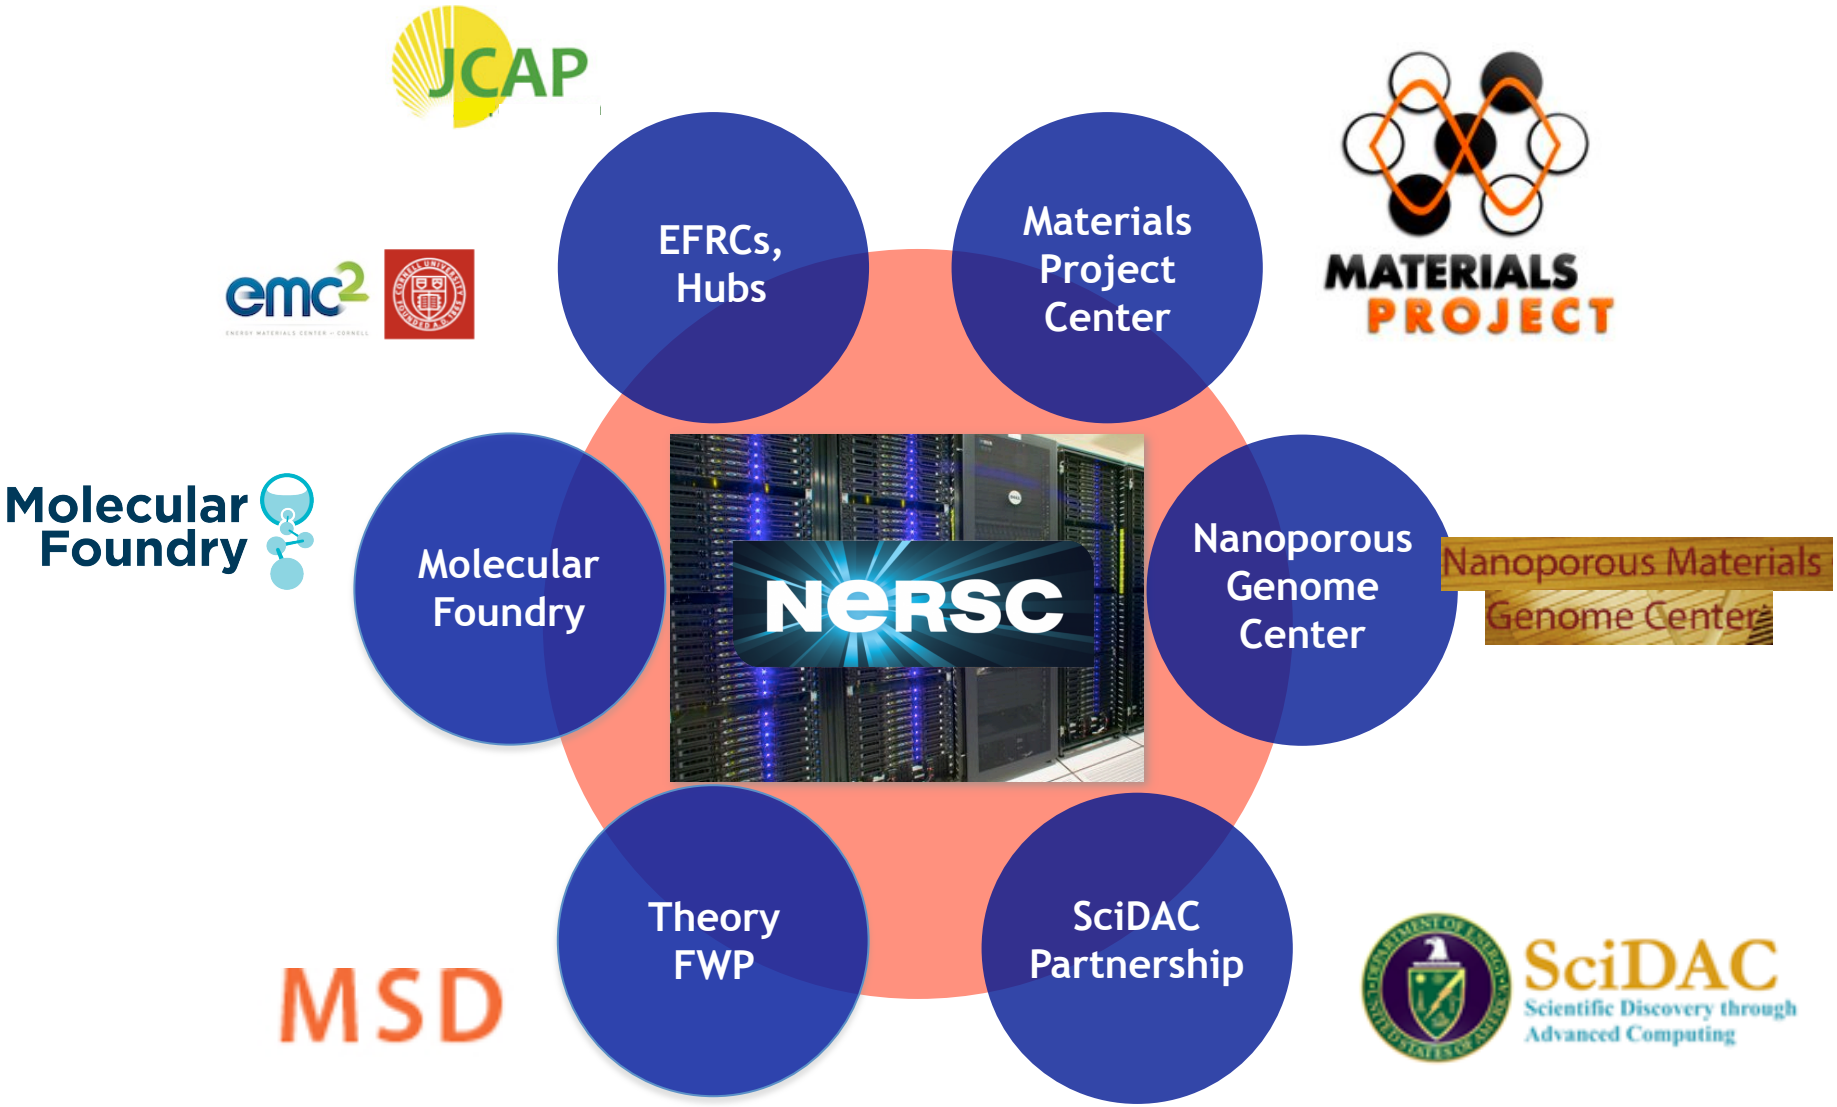

SciDAC Partnership:
Scalable Computational Tools for Discovery and Design of Excited State Phenomena in Energy Materials

PIs: J.R. Chelikowsky and A.A. Demkov, UT-Austin; Y. Saad, University of Minnesota
S.G. Louie, J. B. Neaton, A. Canning, and C. Yang, LBNL
J. Deslippe (National Energy Research Scientific Computing Center)

The objective of our proposed work is to develop and implement new methods and theories to predict electronic excited state phenomena in energy related materials, e.g., materials for photovoltaics, photocatalysis, and electrical energy storage.

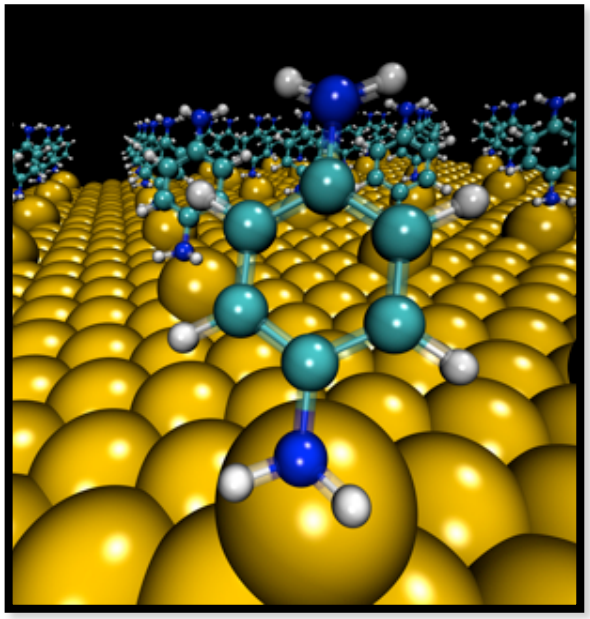

Representative Science Applications

- Chromophore or dye molecules in gas-phase, in solution, or at an oxide surface
- Solvents, e.g. water or electrolyte, in contact with an oxide or metal electrode
- Molecular junctions and self-assembled monolayers
- Multicomponent inorganic semiconductor nanostructures
- Transition metal oxides, with defects, dopants, and with magnetic cations
- Nanophase crystals and clusters appropriate for photovoltaic applications
- Organic molecular crystals and assemblies, and donor-acceptor molecular interfaces

BerkeleyGW

Excited States & Transport in Complex Energy Materials

Transition metal oxides

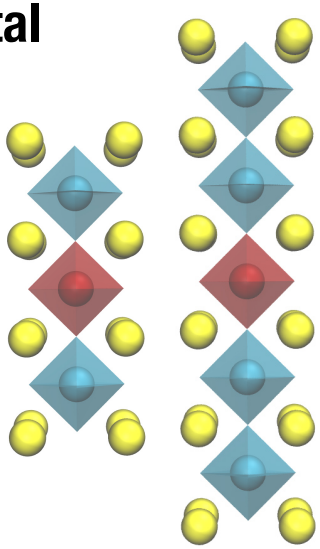

Single-Molecule Junctions & Organic Semiconductors

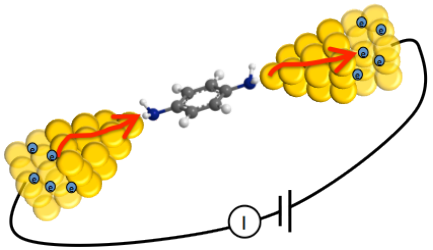

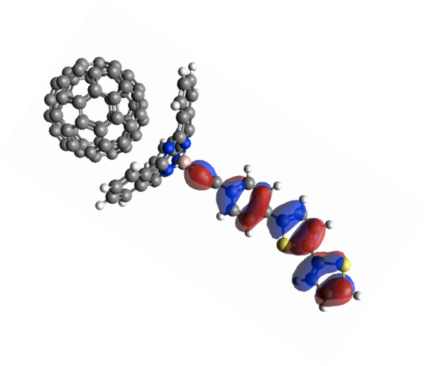

Metal-Organic Frameworks

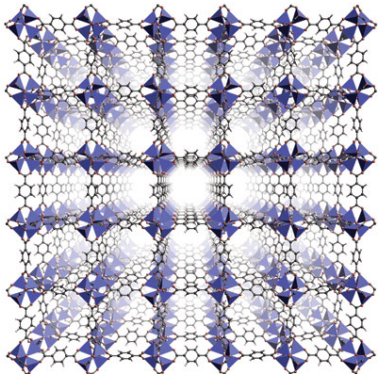

Surfaces and hybrid interfaces

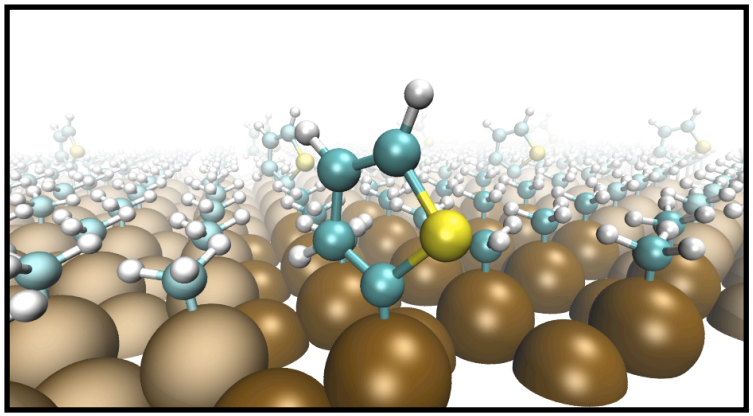

Molecular Foundry

Density Functional Theory and Beyond

- Density Functional Theory (DFT)
  - Structural energetics
  - Chemical specificity, reactivity
  - van der Waals interactions
  - Band gaps, excited states

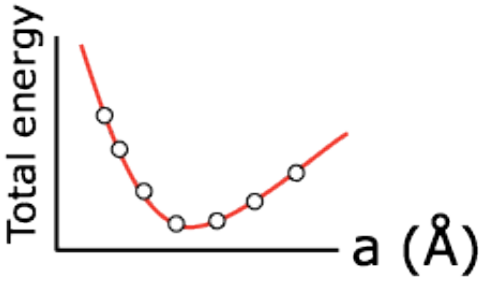

- DFT + Many-Body Perturbation Theory
  - Charged excitations w/ GW
    - GW approximation
    - Electron addition/removal, photoemission, STM
  - Neutral excitations w/ the BSE
    - Bethe-Salpeter equation (BSE)
    - Optical absorption, with e-h interactions

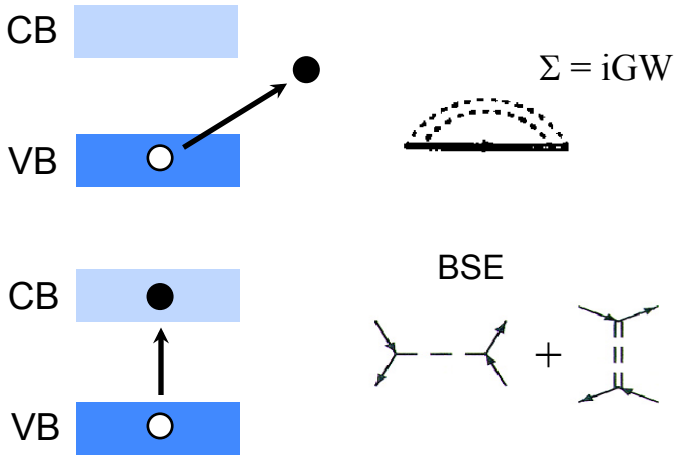

See also Time-dependent DFT, Quantum Monte Carlo, Dynamical Mean Field Theory, wavefunction-based methods

GW: Gap Renormalization & Nonlocal Polarization

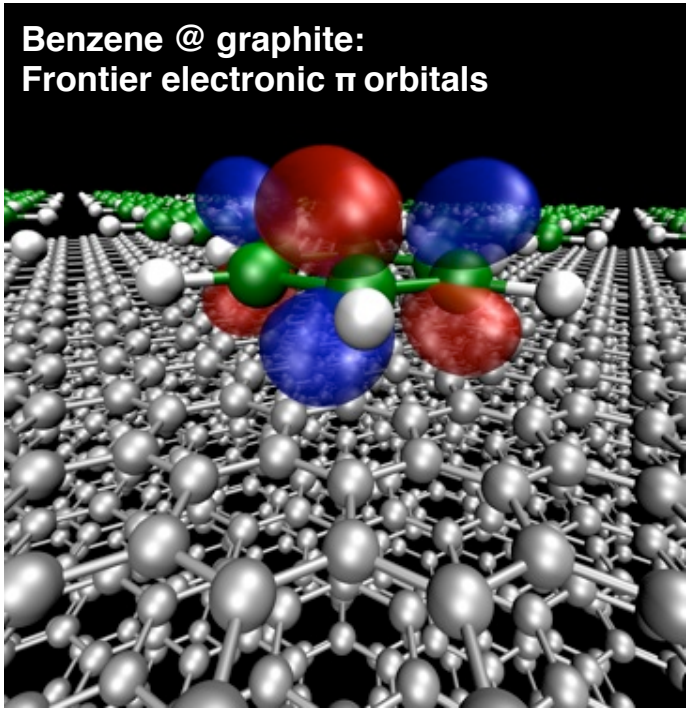

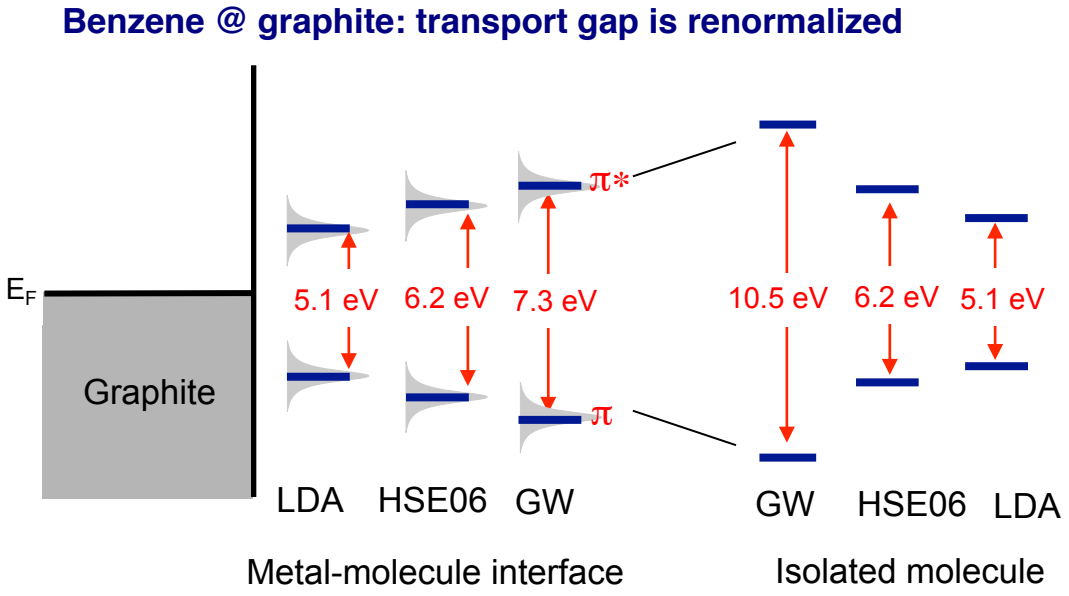

Neaton, Hybertsen, and Louie,
Phys. Rev. Lett. 97, 216405 (2006)

Biller, Tamblyn, Neaton, and Kronik,
J. Chem. Phys. 135, 164706 (2011)

Refaely-Abramson et al,
Phys. Rev. B 88, 081204 (2013)

GW corrections

- large nonlocal polarization effects
- not captured by standard functionals

Recent work: GW for Physisorbed and Covalent Interfaces

Benzene diamine on Au(111)

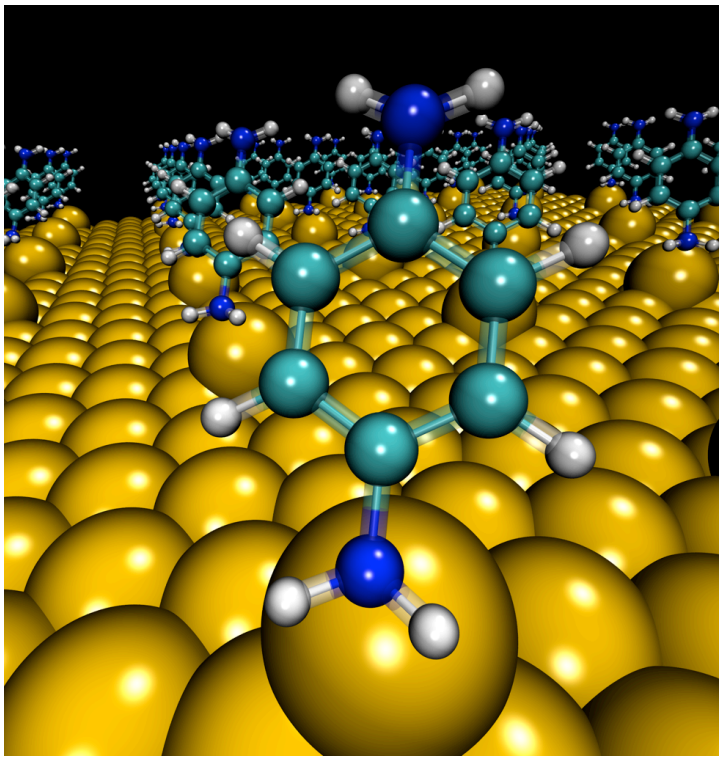

Tamblyn, Darancet, Quek, Bonev, and Neaton,
Phys. Rev. B **84**, 201402(R) (2011)

Thiophene on Functionalized Si(111)

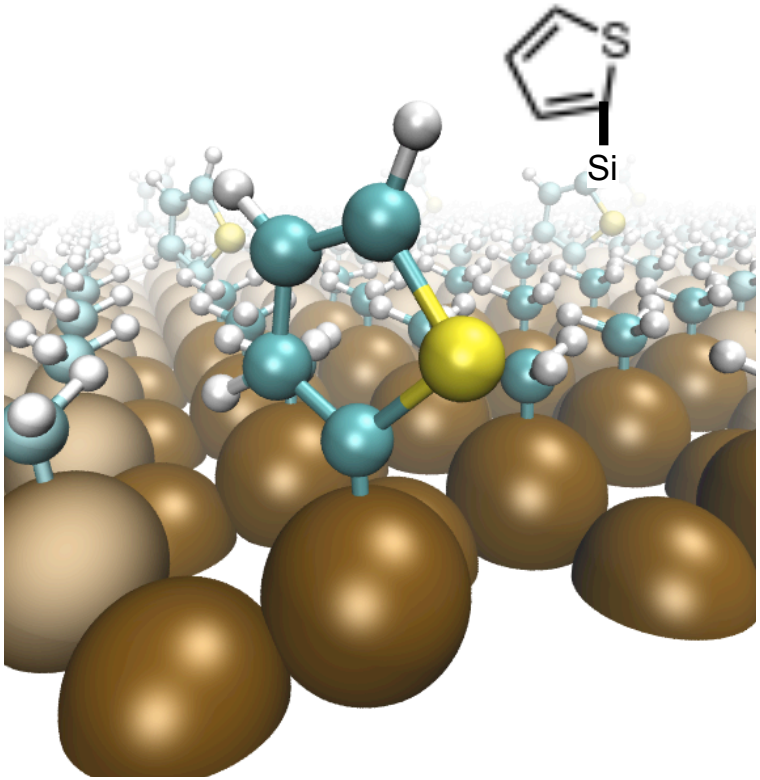

Yu, Doak, Tamblyn,and Neaton,
J. Phys. Chem. Lett. **4**, 1701 (2013)

Longer-Term: Hybrid Systems of 1000-5000 atoms

Chromophore graphene junctions

Gratzel cells

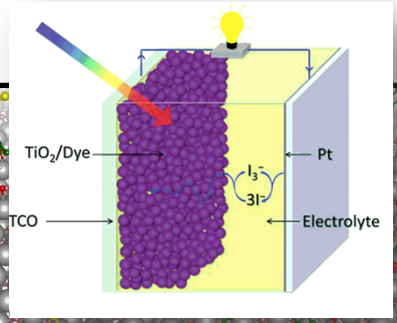

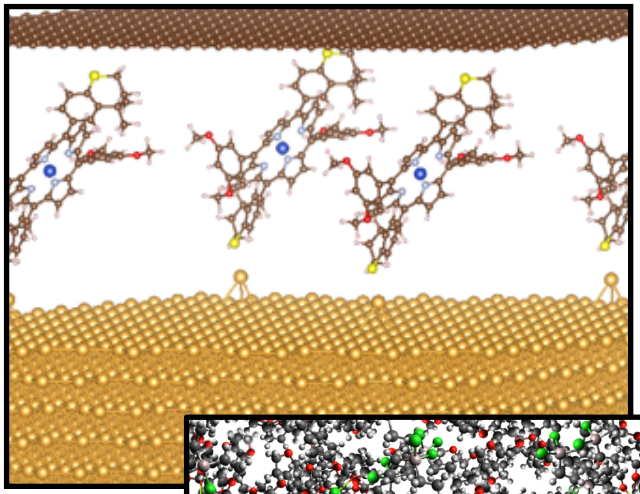

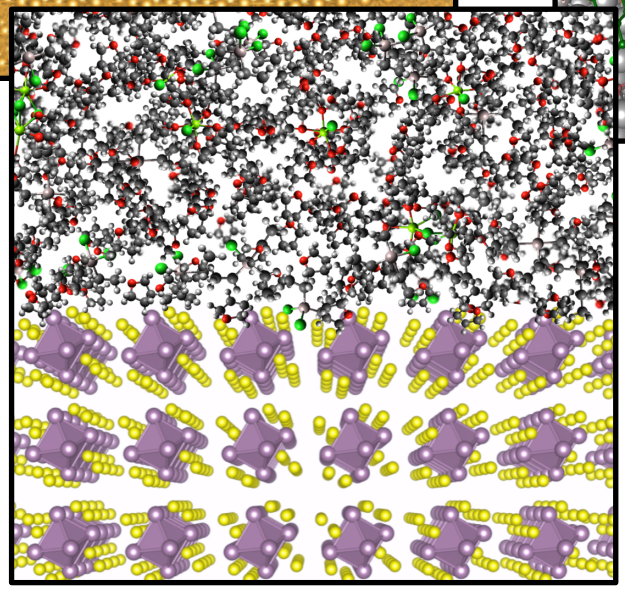

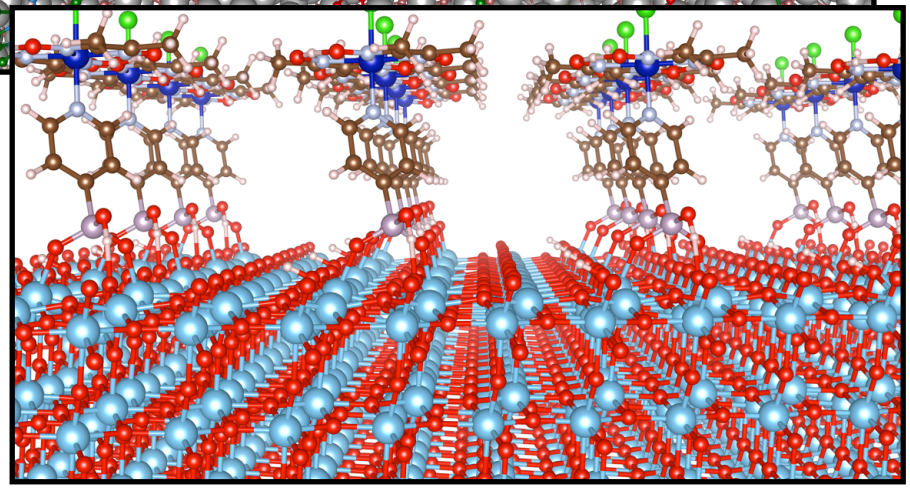

Electrochemical interfaces

Chromophore-Oxide Interfaces

BerkeleyGW Implementation & Scaling

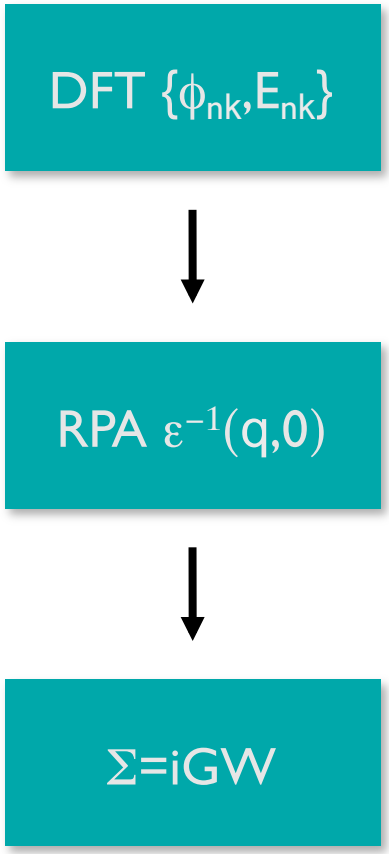

- SCF for ground-state (PW, NC-PPs)
- NSCF for 100-1000's of empty states
- **Bottleneck**: Iterative diagonalization for unoccupied subspace N^3

- Compute $\chi_0(q)$ for G, G'
- Inverts for $\varepsilon^{-1}(q)$ for each G, G'
- Plasmon-pole or fully frequency-dependent
- **Bottleneck**: N^3 computing (v,c) matrix elements and N^4 step for construction of $\chi_0(q)$, plus inversion

- Forms matrix elements to compute Σ, including plasmon pole model
- Outputs $\Sigma - V_{xc}$ for specified nk
- **Bottleneck**: N^3 for vector products & sum over unocc subspace

Beyond DFT Approaches & the Materials Project

- Allow Materials Project Users to reliably calibrate trends DFT with a higher level method

- Use for the development of empirical descriptors for screening band gaps, band edges, and absorption spectra

GW is expensive but simple descriptors are not always broadly applicable

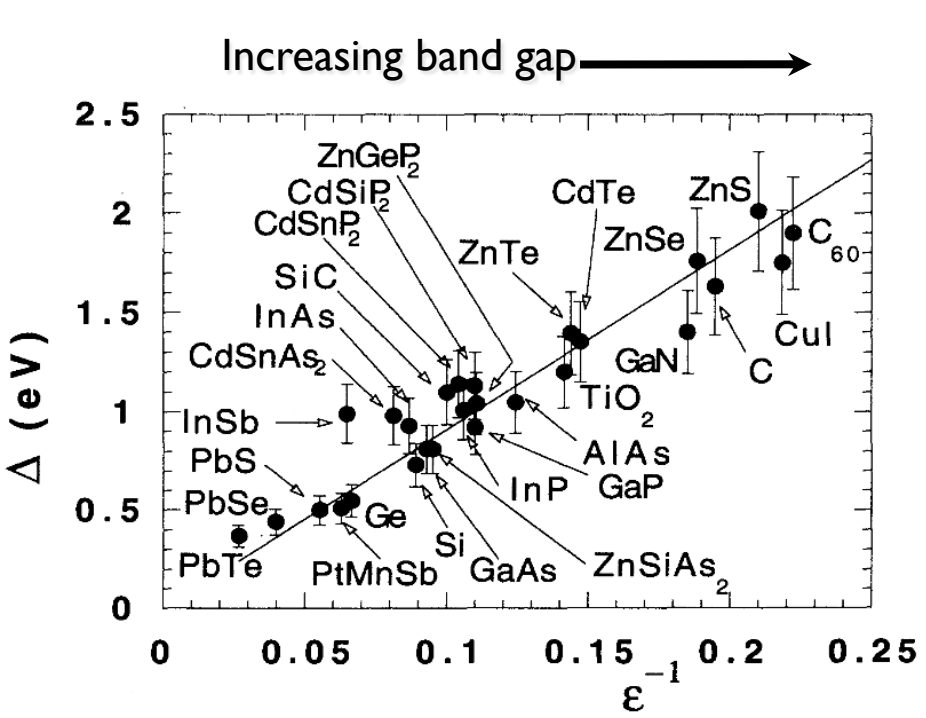

Fiorentini and Baldereschi (1995)

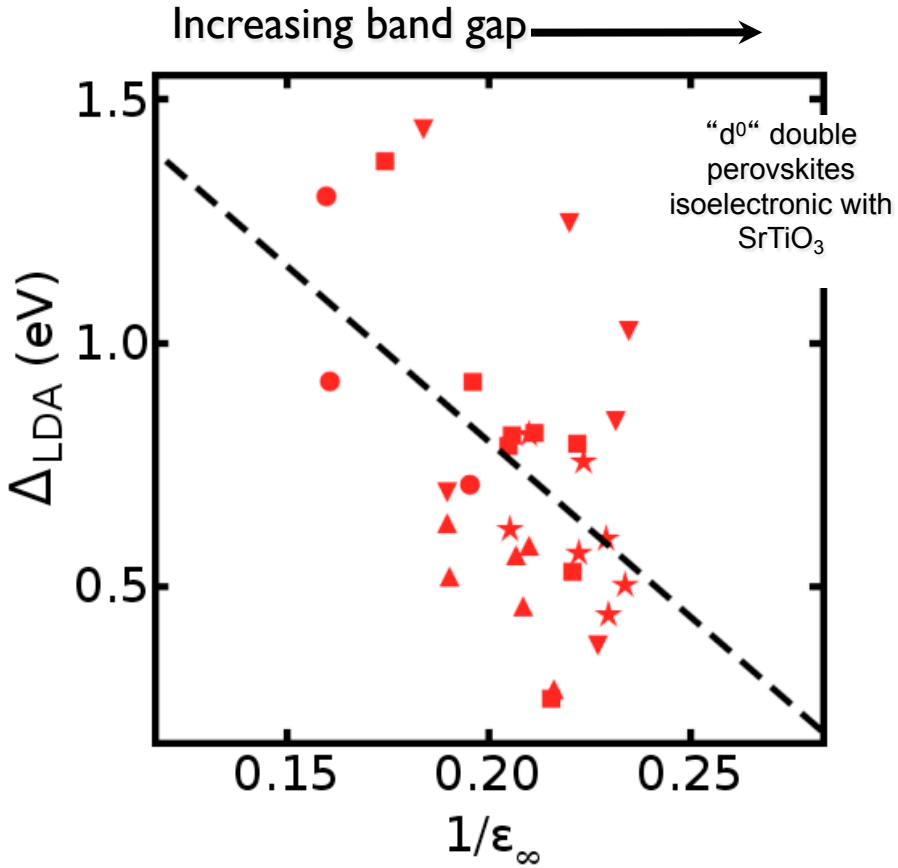

R. F. Berger and J. B. Neaton (2012)

- Gap errors do not follow usual trend in dielectric constant (and gap)
- For oxides, 1/epsilon is not a simple descriptor

New Low-Gap V- and Cr-Based Double Perovskites

$(Sr, Ba)_2(Al, Sc, Ga)VO_6$ $(Sr, Ba)_2(Mg, Ca, Zn)CrO_6$

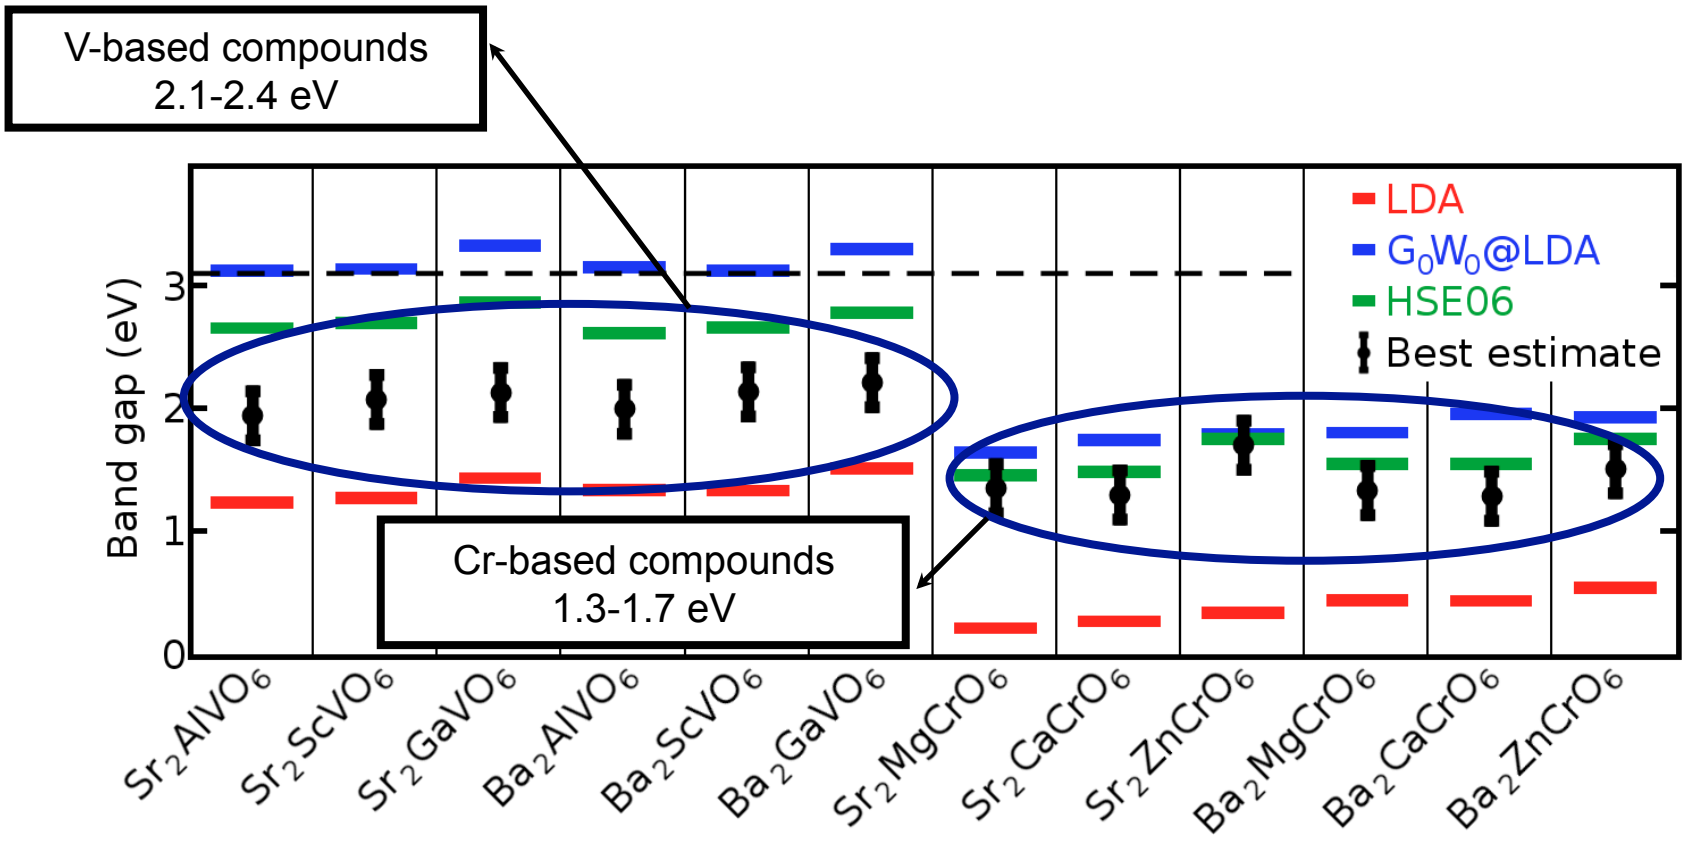

R. F. Berger & J. B. Neaton, Phys. Rev. B 86, 165211 (2012)

Band gaps cover significant parts of the solar spectrum

High-Throughput Data Set

Chemical_formula	Material_ID	Chemical_formula	Material_ID	Chemical_formula	Material_ID	Chemical_formula	Material_ID
Se1	14	SiP	2798	CoSb3	1317	InAs	20305
Te1	19	MoS2	2815	CdSb	1321	InTe	20320
Ge	32	MgSiP2	2961	PdO	1336	InP	20351
C-Diamond	66	CdSiAs2	3078	BaO	1342	PbO	20442
FeS2	226	Cu3AsS4	3345	Mg2Si	1367	PbSe	20667
BeTe	252	ZnSiAs2	3595	MoS2	1434	PbO	20878
PbF2	315	CdSnAs2	3829	BP	1479	OsS2	20905
TlS	322	Cu3PS4	3934	AlP	1550	PbS	21276
Ag2O	353	ZnGeAs2	4008	GaSe	1572	InN	22205
Cu2O	361	ZnSnP2	4175	MoSe2	1634	In2S3	22216
TiO2	390	ZnGeP2	4524	BN	1639	TlInSe2	22232
MnTe	404	CdSiP2	4666	WSe2	1821	InSe	22691
CdTe	406	ZnSnSb2	4756	TlSe	1836	Ca2Sn	22735
Mg2Ge	408	ZnSiP2	4763	SiAs	1863	BiI3	22849
CoAs3	452	ZnSnAs2	5190	SnTe	1883	TlI	22858
CdAs2	471	CdSnP2	5213	CoP3	1944	NaCl	22862
NiP2	486	AgF	7592	Mg3As2	1990	CsCl	22865
HgS	634	Na3Sb	7956	Zn3P2	2071	TlBr	22875
AlN	661	SiC	8062	ZrSe2	2076	PbI2	22893
SnSe2	665	TcS2	9481	ZnO	2133	CsBr	22906
CdS	672	GeAs	9548	AlAs	2172	CuBr	22913
SnSe	691	WS2	9813	ZnTe	2176	CuCl	22914
GeSe	700	Cu3SbSe4	9814	NiS2	2282	AgCl	22922
ZnSb	753	BAs	10044	Mg2Sn	2343	AgI	22925
GaN	805	K3Sb	10159	Cd3P2	2441	TlCl	23167
SnO2	856	Cs3Sb	10378	SrO	2472	TlI	23197
CdP2	913	Cu2SnS3	10519	GaP	2490	AsI3	23218
GeTe	938	GaSe	11342	GaS	2507	AgBr	23231
MnSe	972	Cu2SnSe3	11658	Ca2Si	2517	V2O5	25632
BN	984	K3Sb	14017	GaAs	2534	PbBr2	28077
CdO	1132	Rb3Sb	16319	BeO	2542	Cd7P10	29576
GaSb	1156	Cu2Se	16366	CaO	2605	RbAu	30373
SnS2	1170	Ag8SnSe6	17984	AlSb	2624	ReSe2	541582
ZrS2	1186	Ag8GeSe6	18474	TiO2	2657	Bi2Se3	541837
ZnSe	1190	CoO	19079	CsAu	2667	Cu2S	553942
Sb2Te3	1201	InSb	20012	CdSe	2691	TlGaSe2	567453
MgO	1265	InS	20094	ZnP2	2782	TlInTe2	568517
		MnO	578625	AgI	580941	ReS2	572758

M. K. Y. Chan and G. Ceder, Phys. Rev. Lett. 105, 196403 (2010)

Preliminary G_0W_0@PBE results

25 direct-gap Materials (gaps at Γ)

Cu_2O, CdTe, AlN, CdS, GaN, SnO_2, ZnSe, MgO, ZnO, ZnTe, GaP, BeO, TiO_2, CdSe, $MgSiP_2$, $CdSiAs_2$, $CdSiP_2$, Cu_2SnS_3, InP, NaCl, CsCl, CuBr, CuCl, AgI

G0W0 corrections improves agreement with experiment, modulo certain outliers

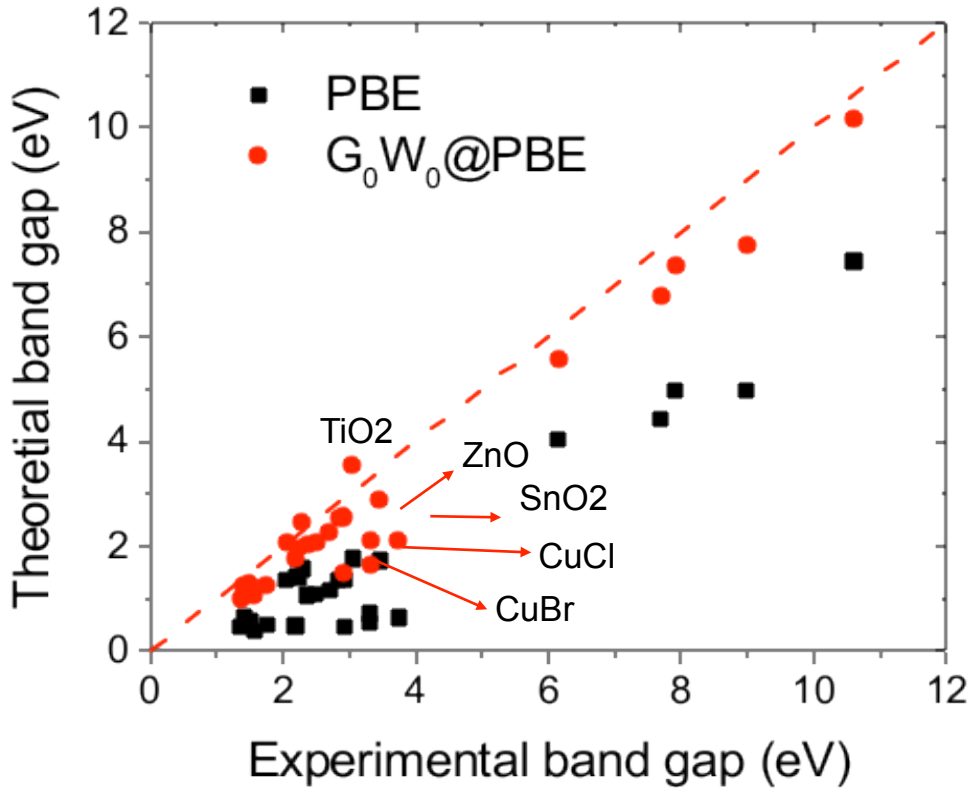

Yan and Neaton, unpublished (2013)

2. Computational Strategies

- The codes we use are: Quantum ESPRESSO, SIESTA/TranSIESTA, SCARLET, VASP, BerkeleyGW

- These codes are characterized by these algorithms:
  - FFTs (FFTW)
  - DFT: Iterative diagonalization (Conjugate Gradient, Davidson, etc)
  - Dense Linear Algebra. Matrix multiplication, Inversion, Diagonaliztion via BLAS/LAPACK. Parallelization via custom routines as well ScaLAPACK
  - Summation Loops (do i=1,n; sum += a(i); enddo)

Our biggest computational challenges are:

- IO Costs ~ 1TB Writes / Reads
- Production of empty states within DFT (typically bound by orthogonalization/subspace-diagonalization)
- N^4 Scaling for GW (Parallel Matrix-Multiplication) will become a bottleneck
- Threaded Library Performance (FFTW, BLAS etc)
- All To All Communication in DFT and GW summation

- Our parallel scaling is limited by:
  - IO
  - ScaLAPACK (For parallel Inversion and diagonalization)
  - All-To-All Communication

- We expect our computational approach and/or codes to change (or not) by 2017 in this way:
  - GW without empty state approaches typically involve solution of linear systems via CG. DFT like problem. Still BLAS like N^4 step.

3. Current HPC Usage

- Machines currently using:
  - Hopper/Edison
  - Argonne
  - Local Cluster (Lawrencium)
- Hours used in 2012-2013:
  - 15 million NERSC
- Typical parallel concurrency and run time, number of runs per year:
  - DFT: 100-few thousand cores (~100 materials)
  - GW: 1000-Few tens of thousand cores (~100 materials)
- Data read/written per run
  - DFT: Up to 100's of GB
  - GW: Up to a few TB

3. Current HPC Usage

- Memory used per (node | core | globally)
  - Node: As much as possible. Often run with with fewer than max MPI tasks per node to fit. OpenMP helps here.
  - Typically require 2GB / core in large cases
  - 100s of GB to 10's of TB in large case
- Necessary software, services or infrastructure
  - Optimized math libraries (FFTW, BLAS, LAPACK, ScaLAPACK)
  - Optimized libraries for parallel IO (HDF5)
- Data resources used (/scratch,HPSS, NERSC Global File System, etc.) and amount of data stored
  - Some users are using ~ 20+ TB of /scratch and global scratch space
  - Shared project directory for Molecular Foundry and SciDac projects ~1TB
  - HPSS, currently using ~ 100 TB for long term archiving

4. HPC Requirements for 2017

- Compute hours needed (in units of Hopper hours)

250 Million Hours (~10^3 from now, but anticipate algorithmic improvements, reductions in empty states etc..)

- Changes to parallel concurrency, run time, number of runs per year
  - We will continue increasing concurrency to study larger problems – optimal number of cores scales roughly as N^2
  - We have added hybrid MPI/OpenMP support, that allows scaling to higher core counts
- Changes to data read/written
  - Data written scales as N^2. Expect an increase of approximately 100x to ~ 100TB.
- Changes to memory needed per (core | node | globally)
  - Memory scales as N^2. Expect an increase of approximately 100x to 100 -1000 TB globally.
- Changes to necessary software, services or infrastructure
  - Continued access to high quality math libraries on new architectures
  - Improved IO bandwidth

5. Strategies for New Architectures

- Our strategy for running on new many-core technologies (GPUs or MIC) is:
  - We are testing phase. Basic MPI/OpenMP Xeon-Phi version available, with parts of code optimized; GPU branch (CUDA-Fortran) in active development.
- To date we have prepared for many core by:
  - Ported code to a Hybrid MPI/OpenMP model
  - Profiling key code loops and libraries on GPUs and MICs
  - Optimized key code loops for on-node parallelism and vectorization

5. Strategies for New Architectures

- We are already planning to do:
  - Work with NERSC to ensure threaded or accelerated math libraries perform well for DFT and GW
  - Continue to port remaining BerkeleyGW executables to GPU and MIC
- To be successful on many-core systems we will need help with:
  - Optimizing DFT for many-core. Quantum ESPRESSO. David Prendergast working on this.
  - Optimizing IO in Both Quantum ESPRESSO and BerkeleyGW
  - Identify key kernels in BerkeleyGW and ensuring vectorization

5. Summary

- What new science results might be afforded by improvements in NERSC computing hardware, software and services?

- Recommendations on NERSC architecture, system configuration and the associated service requirements needed for your science
  - Continue high quality consulting and application support services
  - Minimize amount of code changes and allow code portability (we want to run the same code on commodity machines and other HPC centers)

- NERSC generally refreshes systems to provide on average a 2X performance increase every year. What significant scientific progress could you achieve over the next 5 years with access to 32X your current NERSC allocation?

Acknowledgements

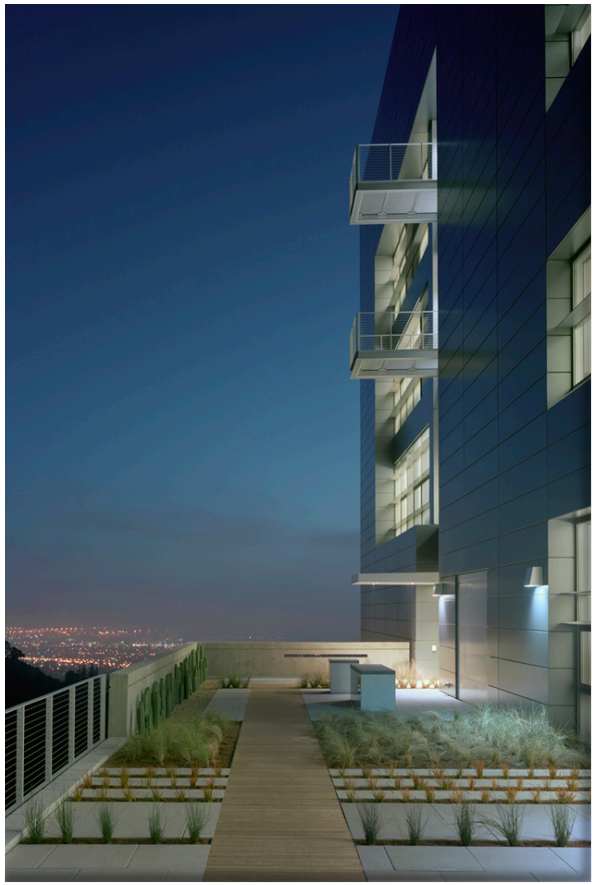

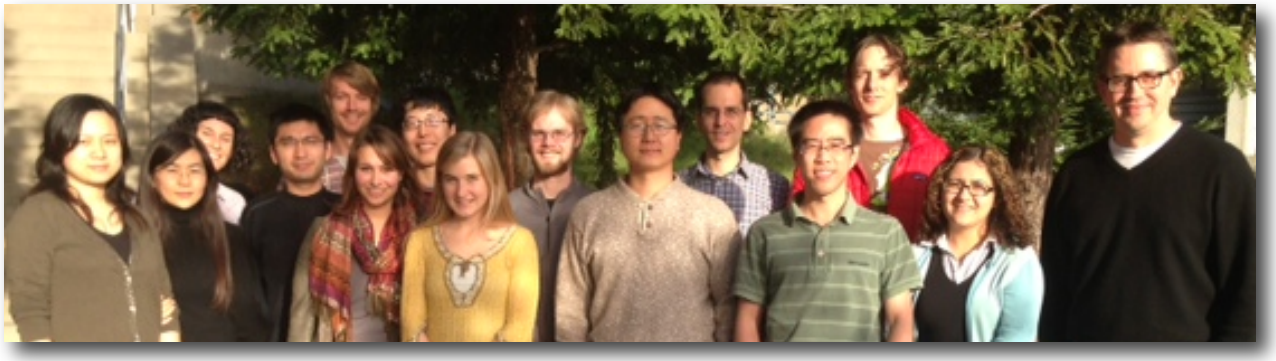

Molecular Foundry

Theory Facility
David Prendergast
Steve Whitelam

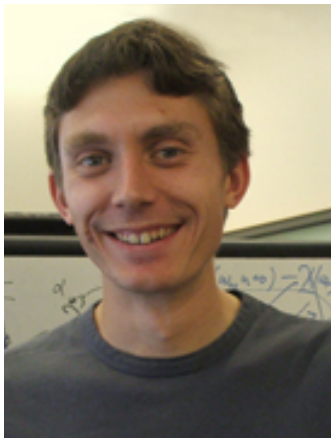

JCAP
JOINT CENTER FOR ARTIFICIAL PHOTOSYNTHESIS

NERSC

SciDAC
Scientific Discovery through Advanced Computing

MATERIALS PROJECT