Linear-Scaling Quantum Chemistry:

Multilayer

Divide-Expand-Consolidate (DEC)

Scheme

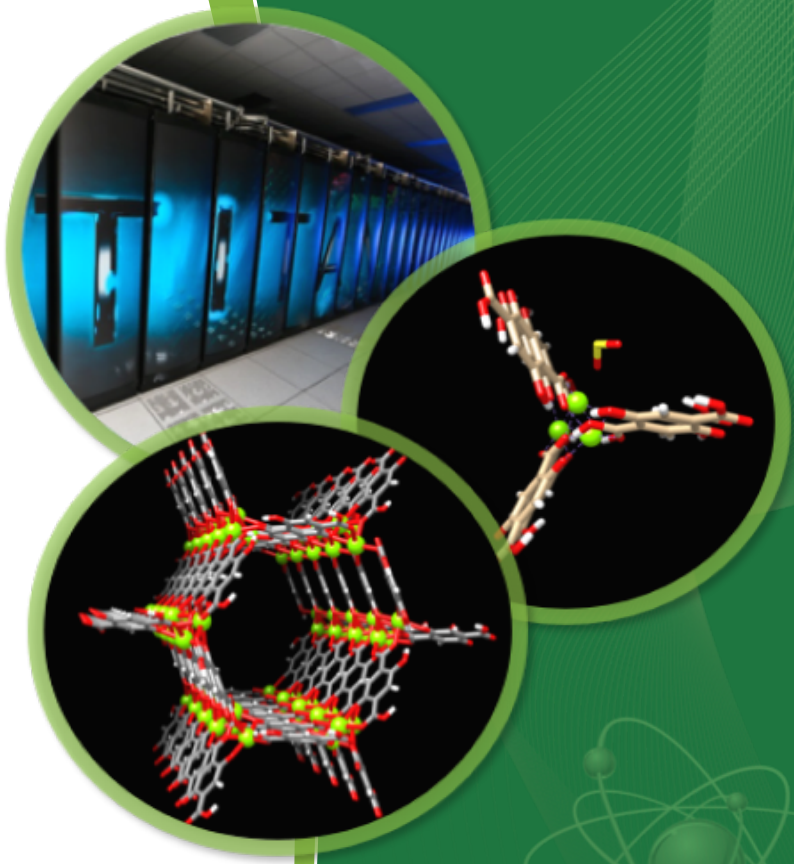

Ashleigh Barnes

OLCF, ORNL

May 16, 2018

The Multilayer DEC Scheme

The Multilayer DEC (ML-DEC) Scheme

Small region of interest within larger system

More time spent on fragments that don't make a strong contribution to property of interest

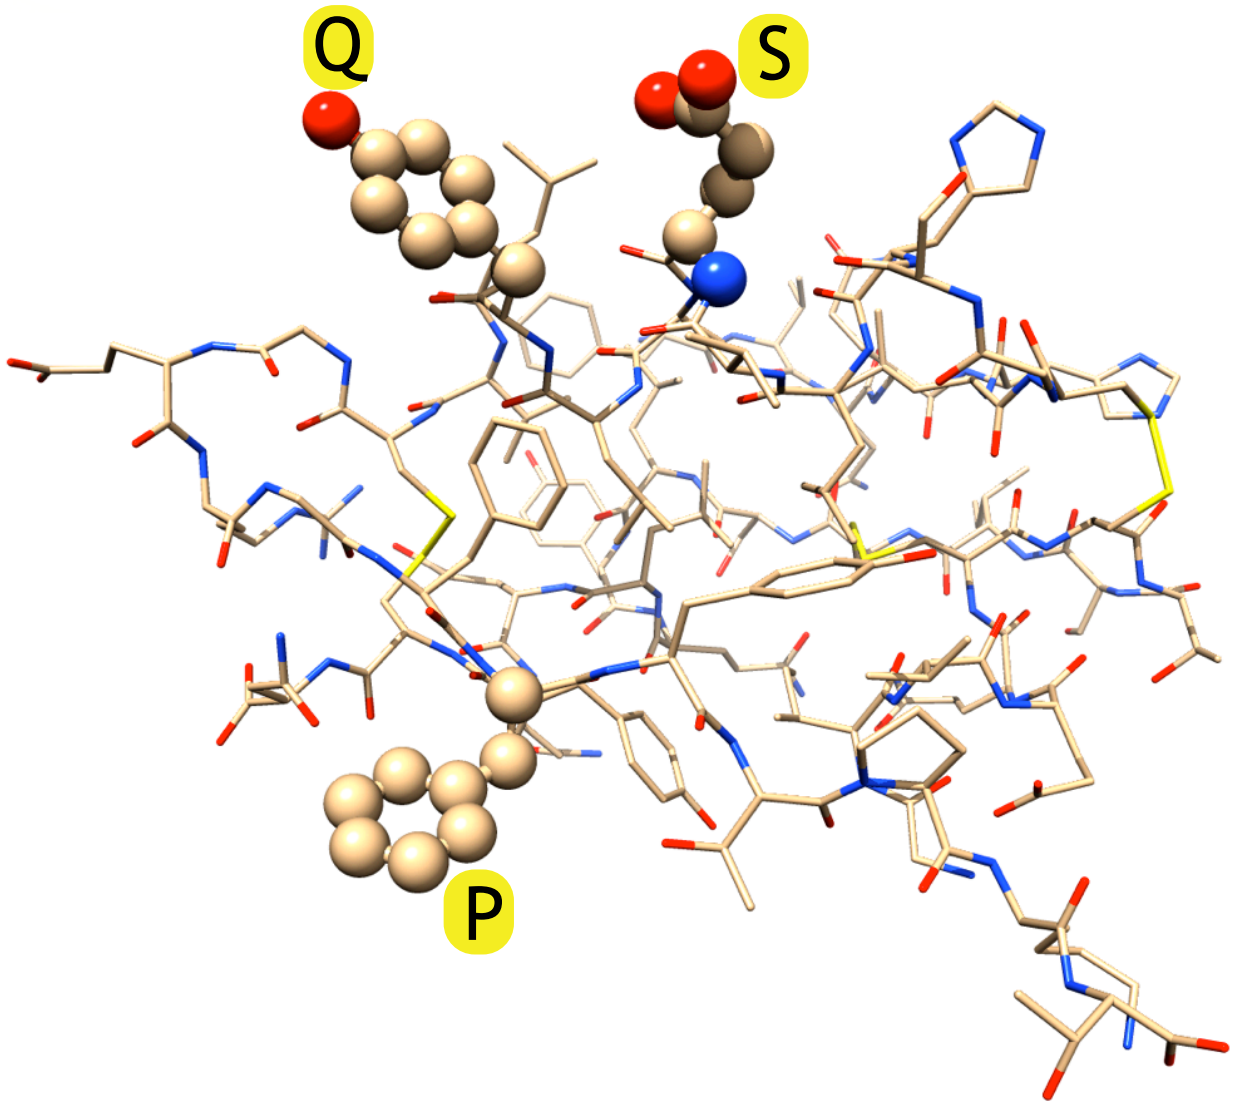

Take advantage of independent nature of fragments

The ML-DEC Scheme

- Example: Region of interest centered on P
  - Q, S evaluated at RI-MP2 level
  - P evaluated at CCSD level
- Intent is to capture important correlation contributions in region of interest – accurate energy differences
- Suited for local effects: adsorption, defect sites, catalytically active sites, etc.

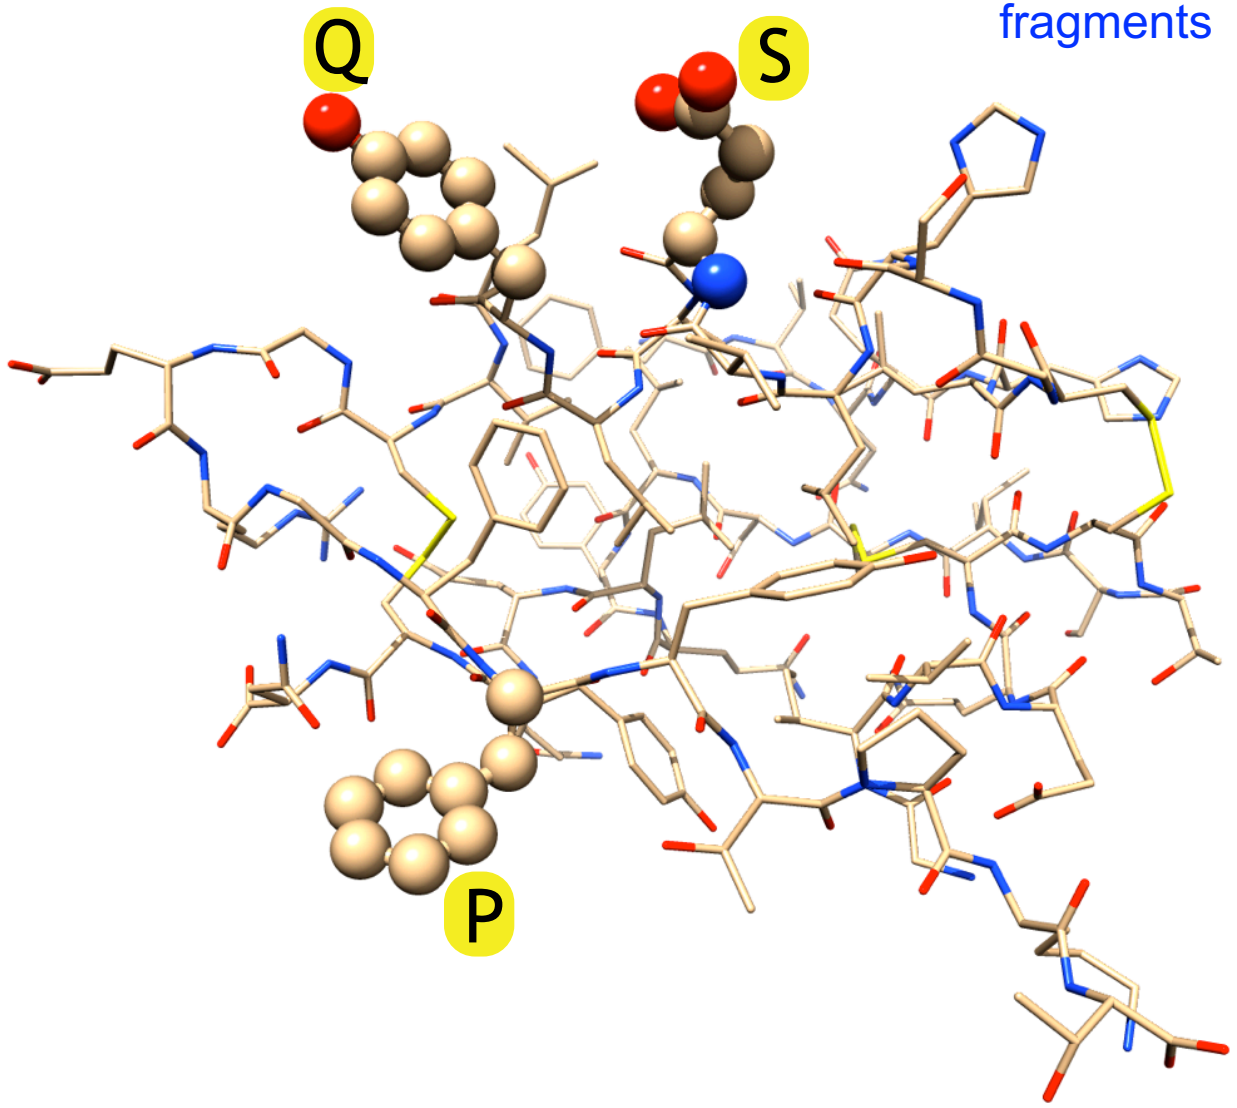

P, Q, S = optimized fragments

The ML-DEC Scheme

High Level

Low Level	RI-MP2	CCSD	CCSD(T)
HF	X		
RI-MP2		X	
CCSD			X

- User defines high and low level layers in input file and assigns atoms to each.
- All fragments optimized at RI-MP2 level of theory

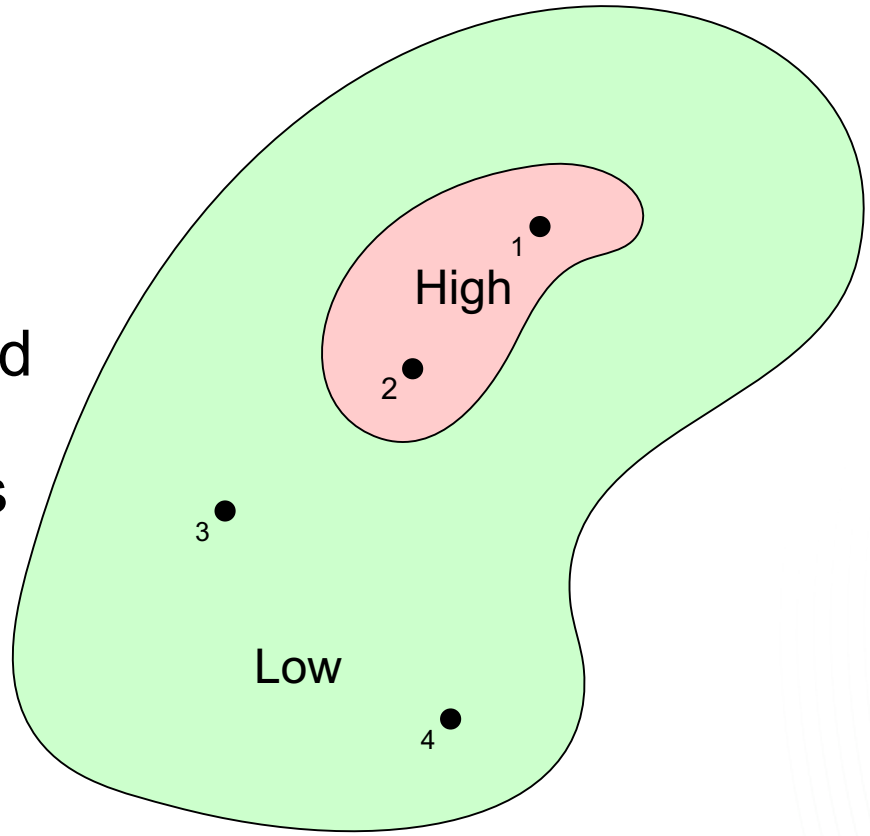

Layers not required to be contiguous
– e.g. multi-site catalysis

- After optimization, atomic fragments assigned to user-specified level of theory.
- Pair fragments are evaluated at the higher level of the two atomic fragments:
  - 1-2: CC_{high}
  - 2-4: CC_{high}
  - 3-4: CC_{low}
- Pair estimates at RI-MP2 level determine if pair level is reduced (not skipped).

The ML-DEC Scheme

Low Level \ High Level	RI-MP2	CCSD	CCSD(T)
HF	X		
RI-MP2		X	
CCSD			X

- Special consideration for HF/RI-MP2:
  - Only optimize low-level atomic fragments that contribute to high-level pairs

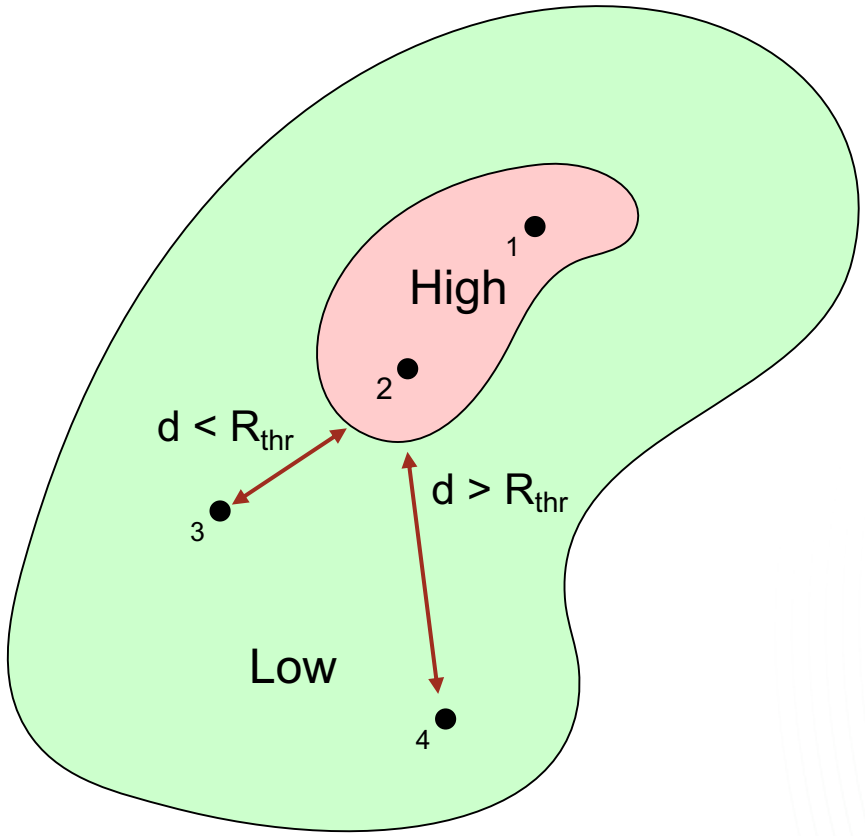

R_{thr}

(recommended: 10 Å)

- Atomic fragments within R_{thr} of any high-level atom are optimized, otherwise skipped
- Unoptimized atomic fragments automatically excluded from pairs
  - 2-3: RI-MP2
  - 2-4: HF
  - 3-4: HF

The ML-DEC Scheme: Summary

$$E_{\mathrm{corr}} = \sum_{P}^{N_{\mathrm{frag}}} \left[E_P + \sum_{Q<P}^{N_{\mathrm{frag}}} \Delta E_{PQ} \right]$$

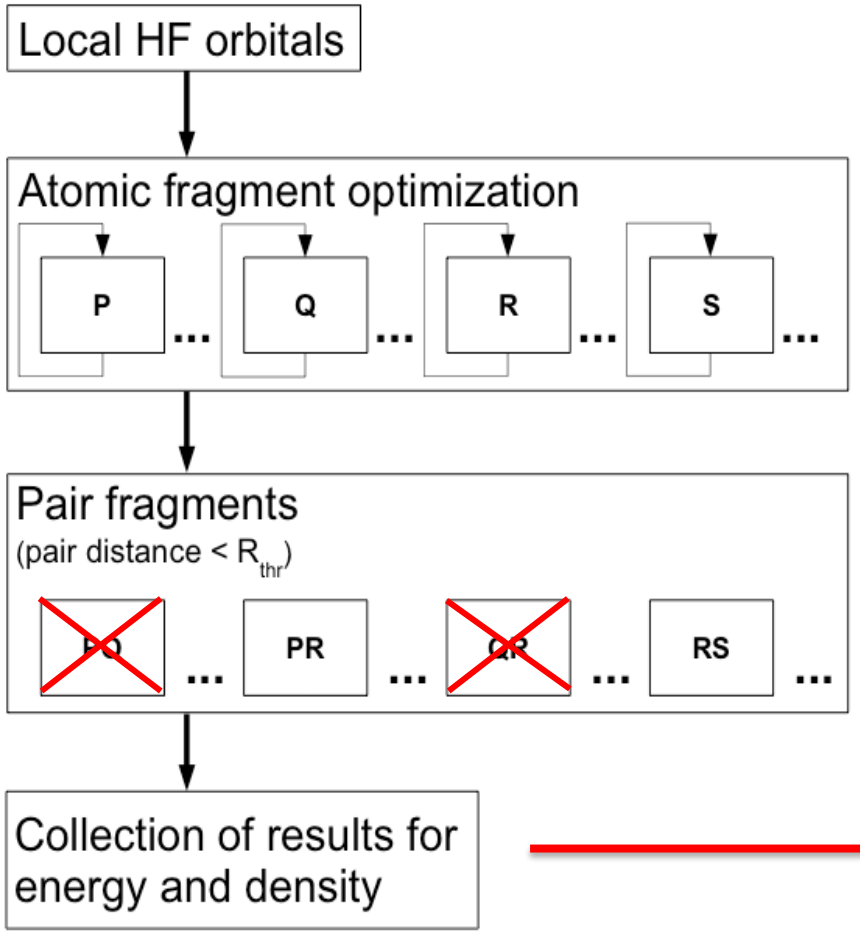

Performed at RI-MP2 level for all fragments

P = Occupied orbital space

$[P]$ = Virtual orbital space

$$E_P = \sum_{\substack{i\in P, j\in P \\ ab\in[P]}} \left(t_{ij}^{ab} + t_i^a t_j^b\right)\left(2g_{iajb} - g_{ibja}\right)$$

Performed at target level

$$\Delta E_{PQ} = \sum_{\substack{i\in P, j\in Q \\ ab\in[P]\cup[Q]}} \left(t_{ij}^{ab} + t_i^a t_j^b\right)\left(2g_{iajb} - g_{ibja}\right) + \quad P \leftrightarrow Q \text{ term}$$

$$E_{\mathrm{corr}} = E_{\mathrm{corr,high}} + E_{\mathrm{corr,low}}$$

M. Ziółkowski, B. Jansík, T. Kjærgaard, and P. Jørgensen, J. Chem. Phys. **133**, (2010).
T. Kjærgaard, P. Baudin, D. Bykov, K. Kristensen, and P. Jørgensen, Wiley Interdiscip. Rev. Comput. Mol. Sci. (2017).

Test Case: Palmitic Acid Dimer

Palmitic acid dimer
cc-pVTZ, FOT=10^{-4}
1976 basis functions

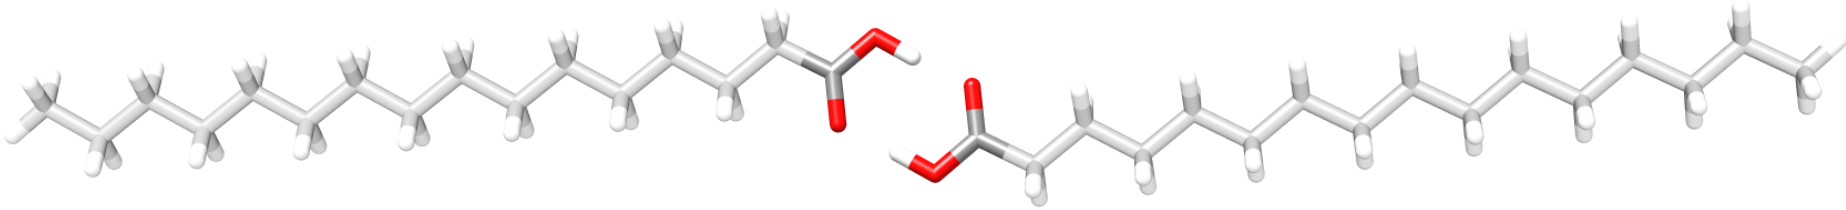

HF/RIMP2, 10.0 Å cutoff

Model	Nproc	RI-MP2 frags	Wall time (min)	E_{corr}
DEC monomer	7	17*, 51	17.51	-3.241(3)
ML-DEC monomer	7	9*, 6	6.71	-1.534(2)
DEC dimer	33	36*, 130	10.19	-6.493(7)
ML-DEC dimer	33	20*, 32	5.51	-3.084(4)

*Atomic fragment optimization

$$\Delta E_{\mathrm{HF/RI\text{-}MP2}} = -0.033$$

$$\Delta E_{\mathrm{RI\text{-}MP2}} = -0.027$$

Test Case: Palmitic Acid Dimer

Palmitic acid dimer
cc-pVTZ, FOT=10^{-4}
1976 basis functions

HF/RI-MP2

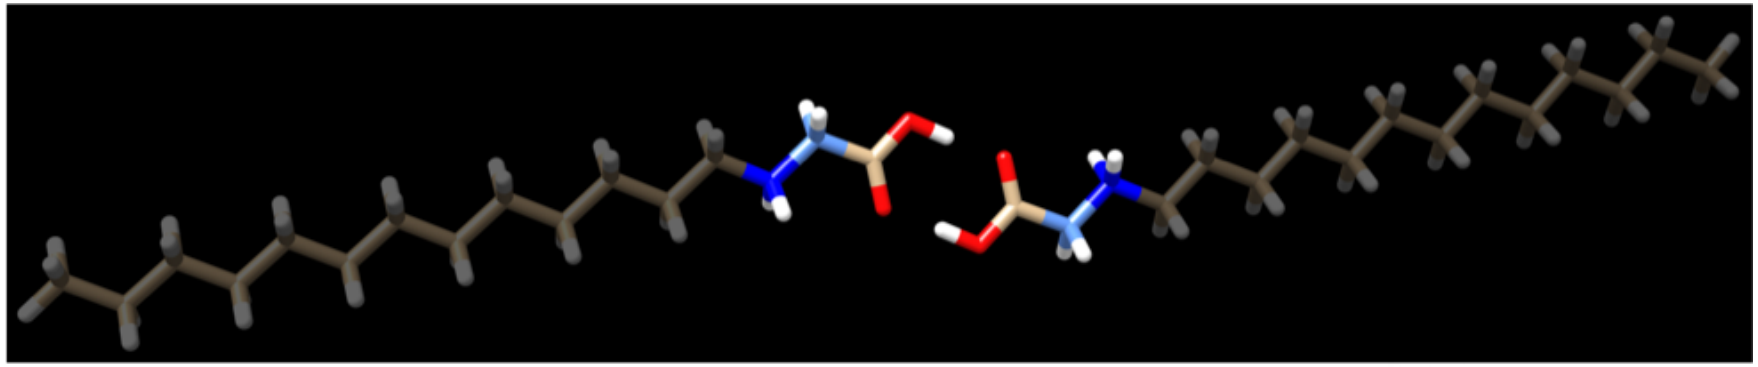

Monomer, 988 basis functions		
Model	Fragment Optimizations	Pair Energy Calculations
ML-DEC	9	6
ML-DEC (ext_1)	10	9
ML-DEC (ext_2)	11	12
DEC-RI-MP2	17	51

Dimer, 1976 basis functions		
Model	Fragment Optimizations	Pair Energy Calculations
ML-DEC	20	32
ML-DEC (ext_1)	22	42
ML-DEC (ext_2)	24	50
DEC-RI-MP2	36	130

Test Case: Palmitic Acid Dimer

Palmitic acid dimer
cc-pVTZ, FOT=10^{-4}
1976 basis functions

HF/RI-MP2

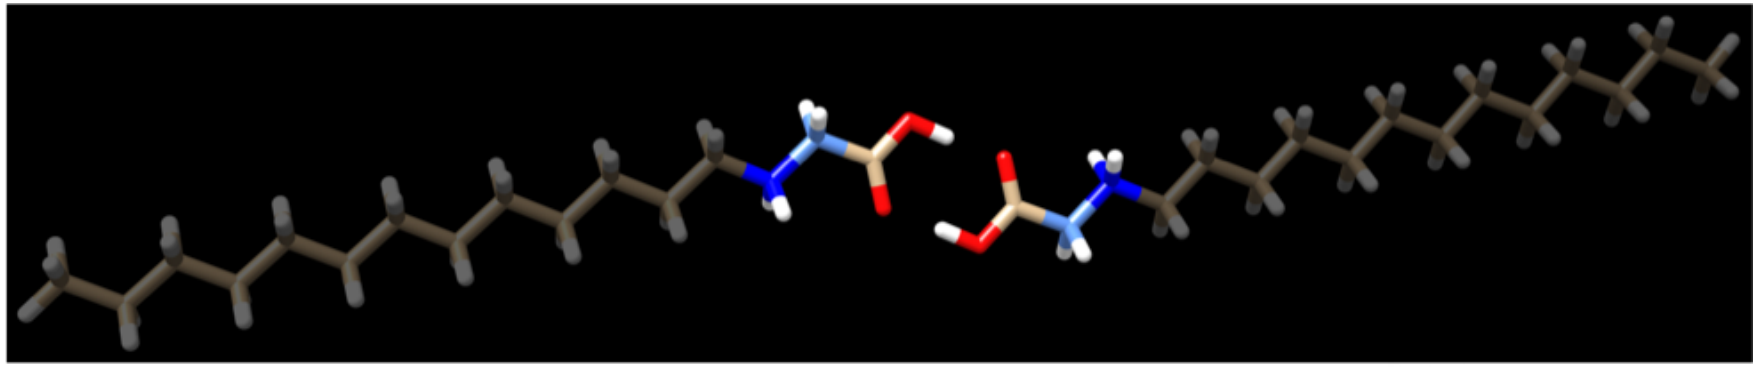

- Even with smallest high level layer, interaction energies agree within FOT.
- Agreement converges to standard DEC RI-MP2 energy as high level is increased.

Model	TTS_m (min)	TTS_d (min)	ΔE (E_h)
ML-DEC	6.71	5.51	-0.033(5)
ML-DEC (ext_1)	7.98	6.21	-0.026(5)
ML-DEC (ext_2)	9.27	6.84	-0.027(6)
DEC-RI-MP2	17.51	10.19	-0.027(8)

Test Case: Palmitic Acid Dimer

Palmitic acid dimer
cc-pVDZ, FOT=10^{-4}
824 basis functions

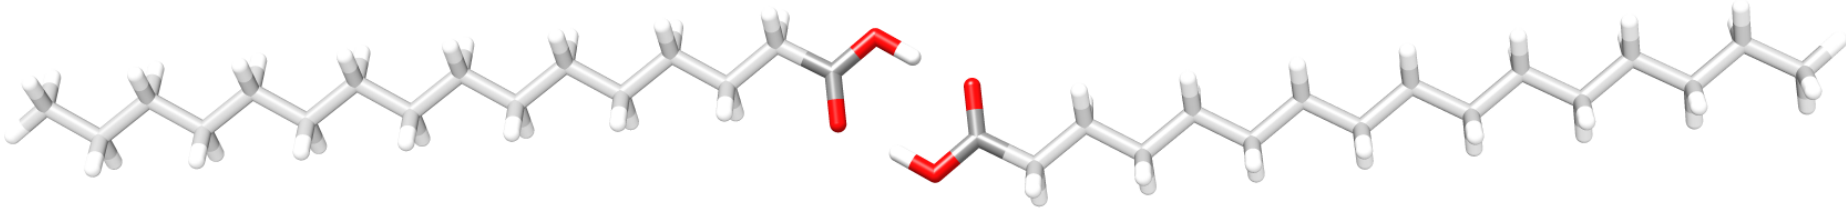

RI-MP2/CCSD

Model	Nproc	CCSD frags	Wall time (h)	E_{corr}
DEC monomer	129	76	3.21	-2.853(3)
ML-DEC monomer	129	14	0.24	-2.655(3)
DEC dimer	129	164	8.4	-5.712(7)
ML-DEC dimer	129	38	1.85	-5.316(7)

$$\Delta E_{\text{RI-MP2/CCSD}} = -0.027$$

$$\Delta E_{\text{CCSD}} = -0.026$$

Application: Acid gas adsorption in Mg-MOF-74

MOFs for CO_2 Capture

High uptake in many MOFs with OMSs
(e.g. Mg-MOF-74, HKUST-1)

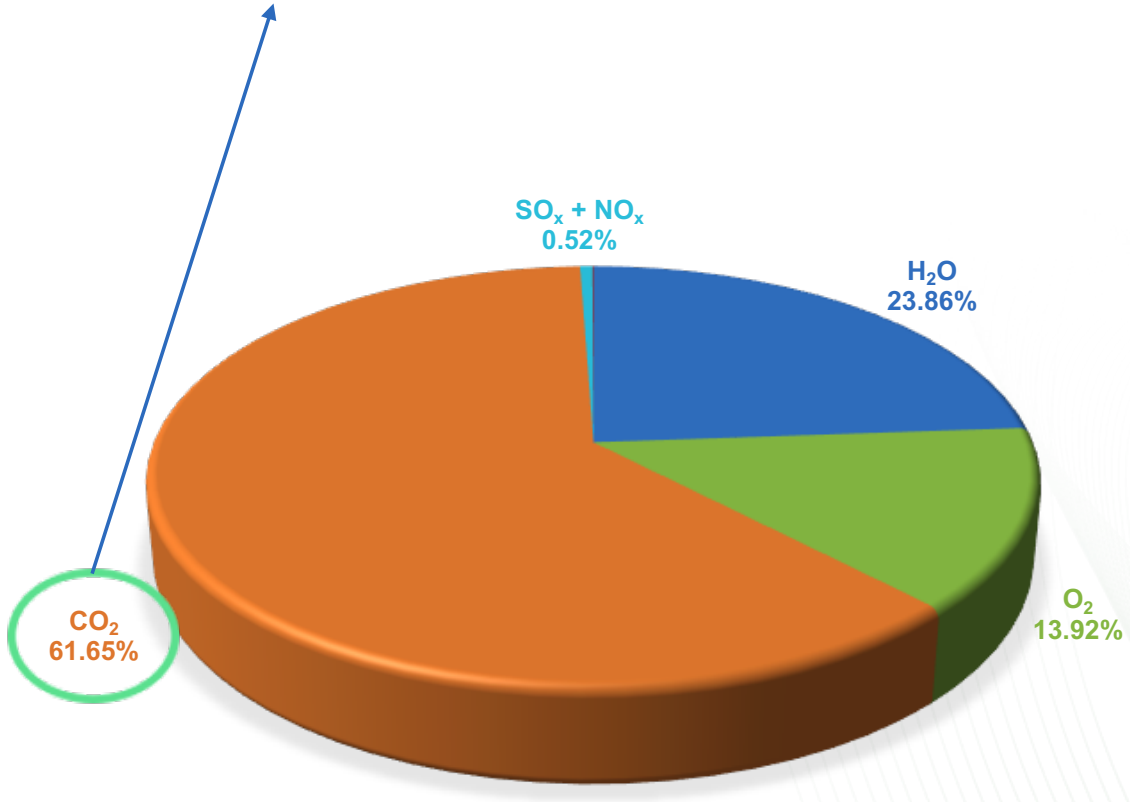

Non-N_2 flue gas components

MOFs for CO_2 Capture

Competitors,
Strong binding to MOF OMSs
→ Lead to degradation of MOFs

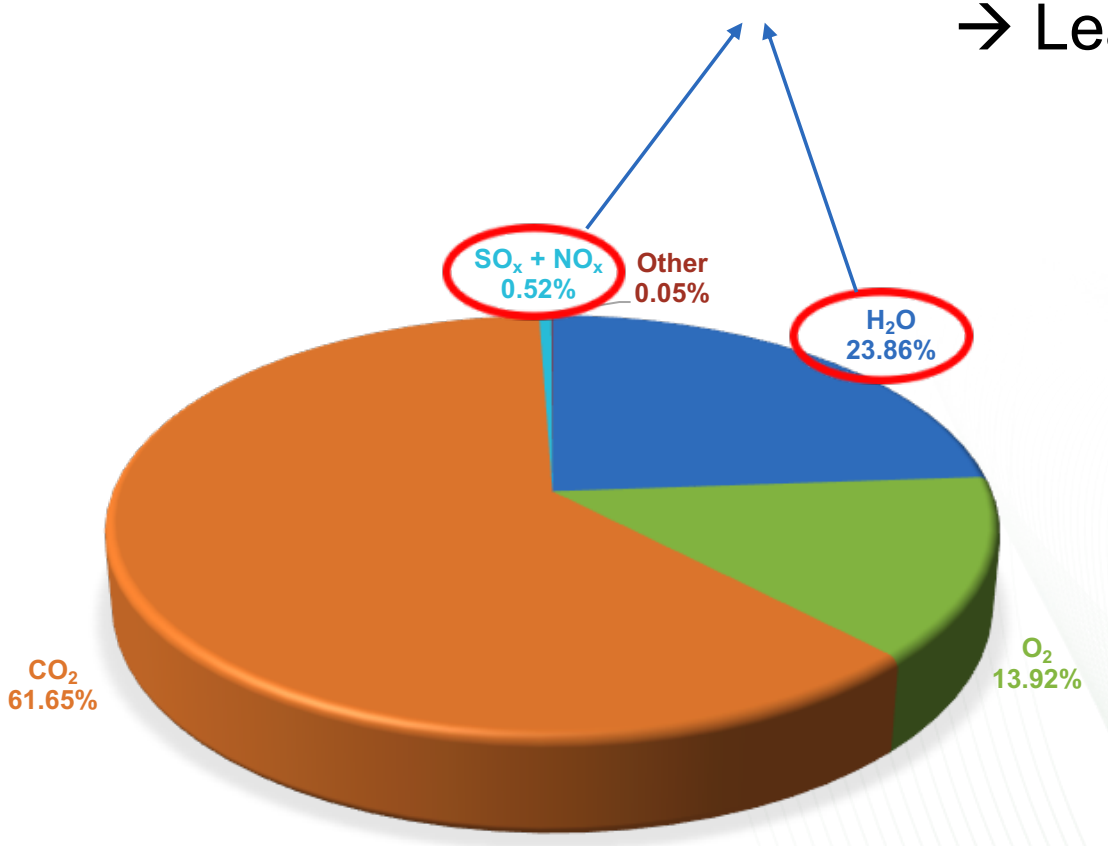

Non-N_2 flue gas components

Modeling Acid Gas adsorption in Mg-MOF-74

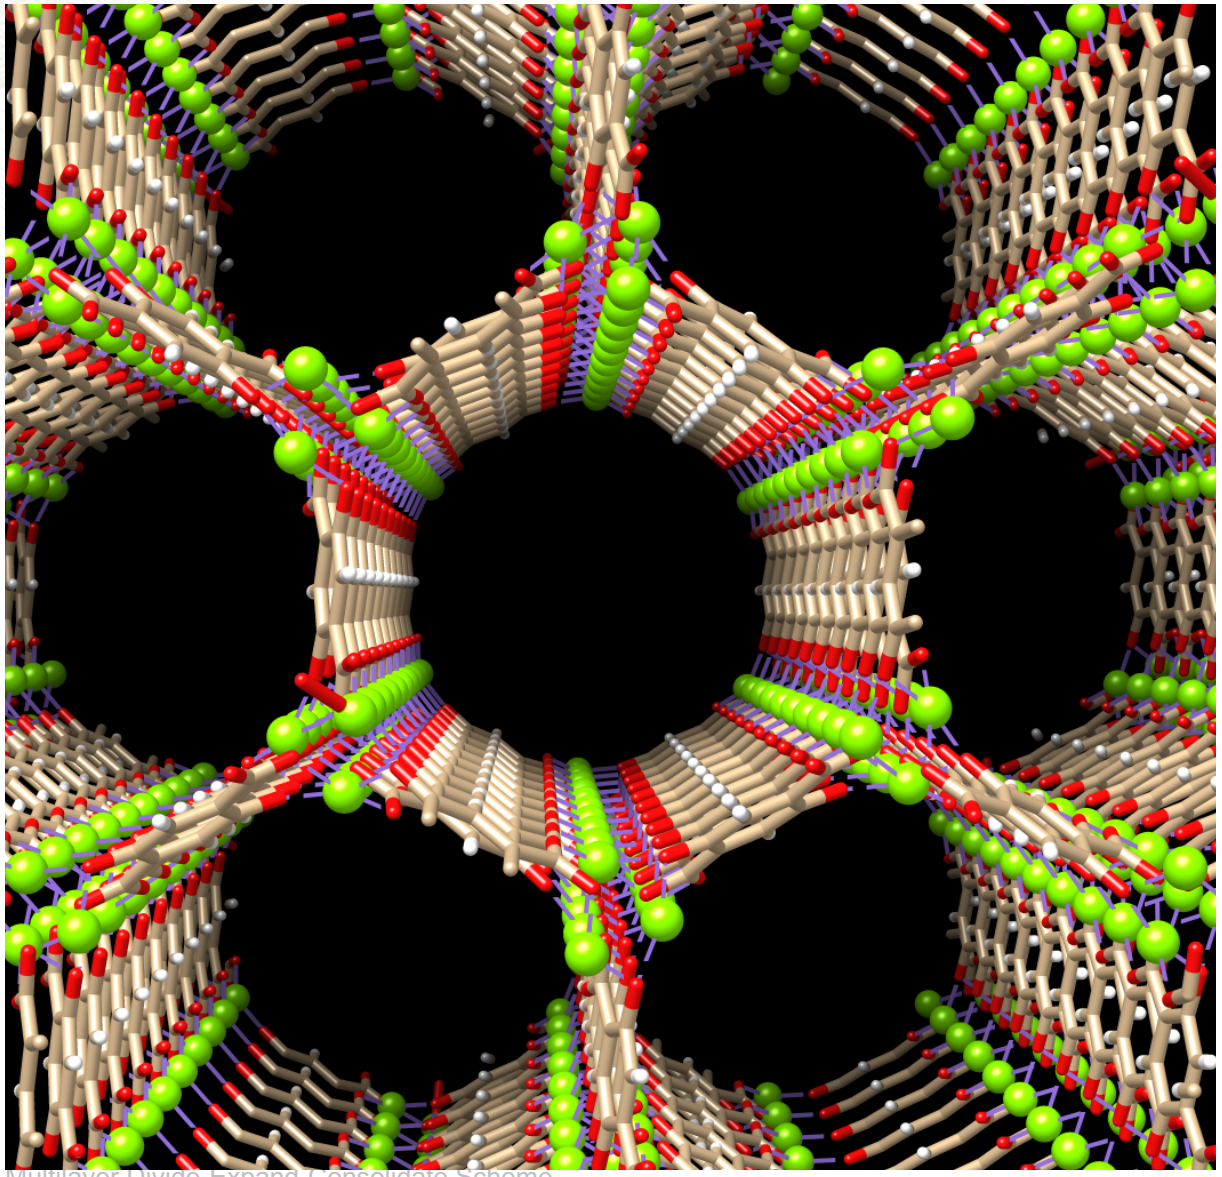

- Plane-wave DFT
  - Results depend on choice of functional, Hubbard U corrections, dispersion corrections etc.
- Small cluster models
  - Suffer from finite size effects. Clusters large enough to replicate environment within the bulk MOF are too expensive for CC.
- Desire to model large finite MOF clusters at CC level of theory while capturing bulk MOF environment.
  - Local interaction – good candidate for ML-DEC scheme.

Modeling Acid Gas adsorption in Mg-MOF-74

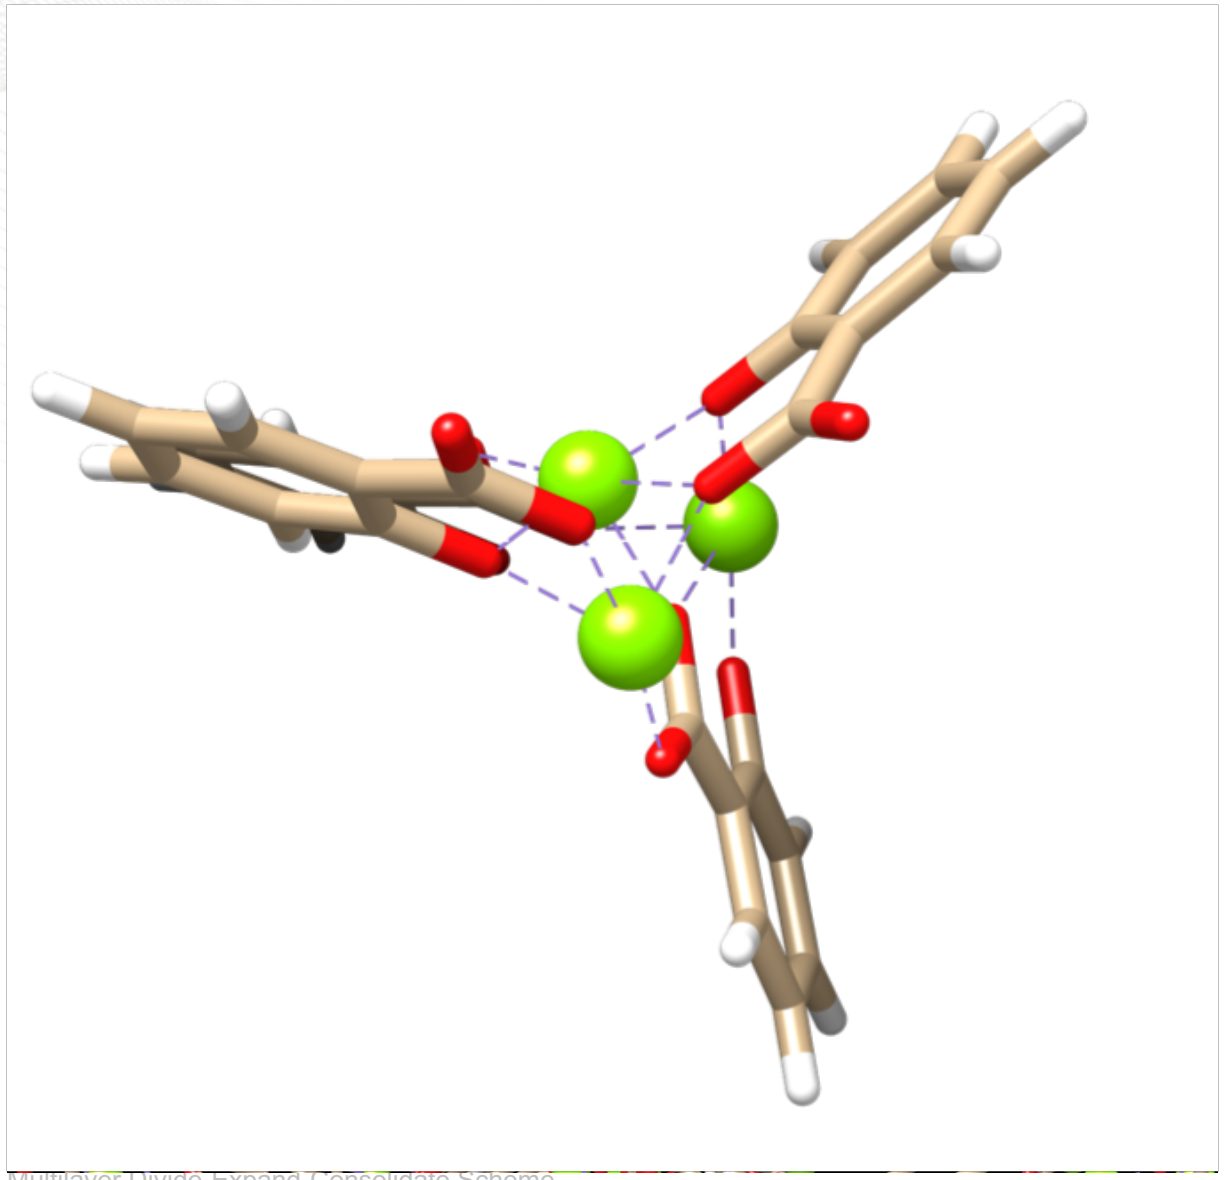

- Plane-wave DFT
  - Results depend on choice of functional, Hubbard U corrections, dispersion corrections etc.
- Small cluster models
  - Suffer from finite size effects. Clusters large enough to replicate environment within the bulk MOF are too expensive for CC.
- Desire to model large finite MOF clusters at CC level of theory while capturing bulk MOF environment.
  - Local interaction – good candidate for ML-DEC scheme.

Modeling Acid Gas adsorption in Mg-MOF-74

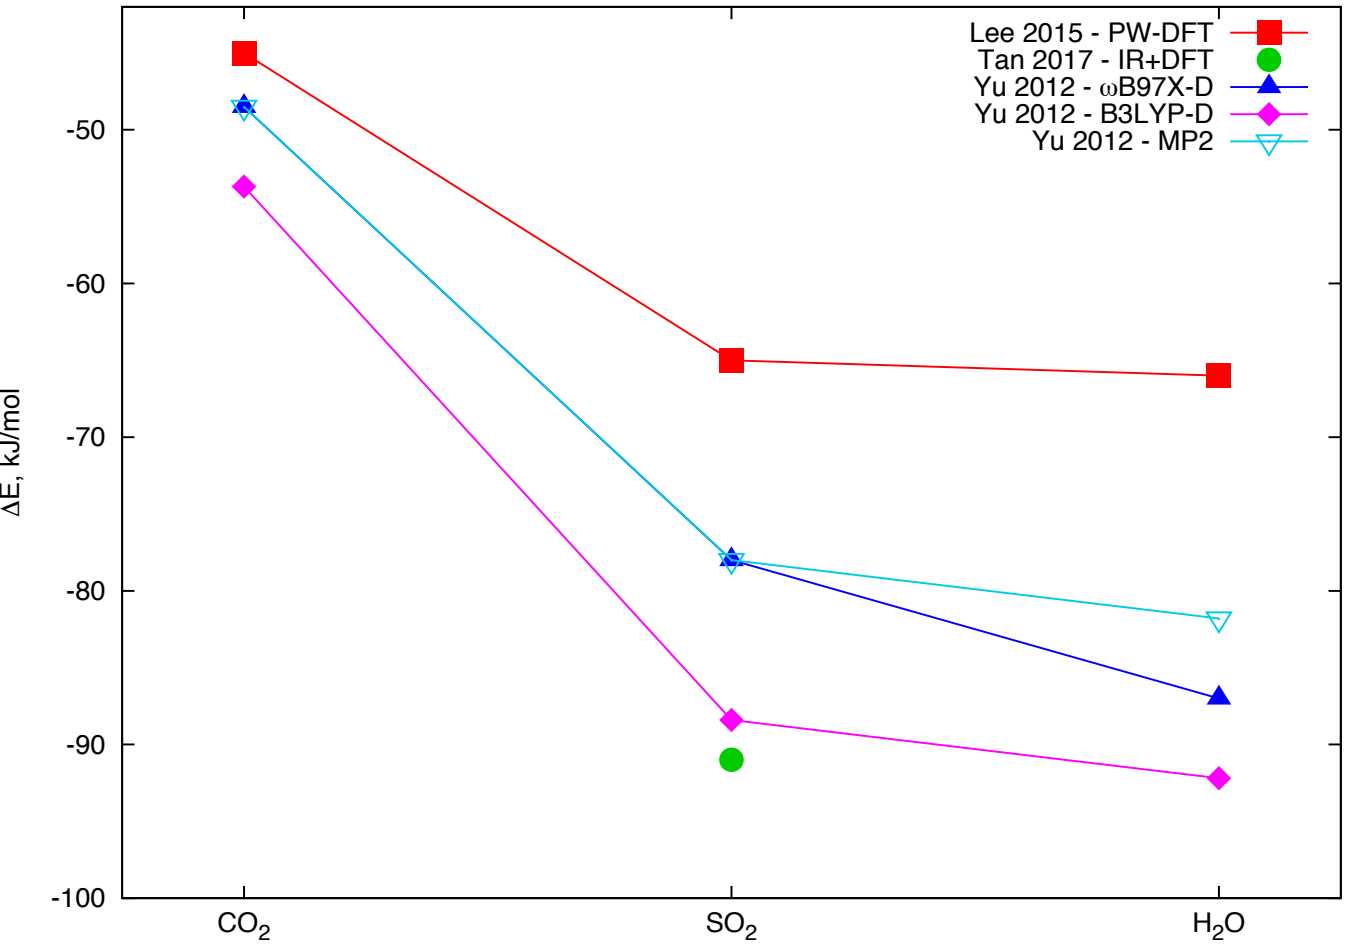

- Plane-wave DFT
  - Results depend on choice of functional, Hubbard U corrections, dispersion corrections etc.
- Small cluster models
  - Suffer from finite size effects. Clusters large enough to replicate environment within the bulk MOF are too expensive for CC.
- Desire to model large finite MOF clusters at CC level of theory while capturing bulk MOF environment.
  - Local interaction – good candidate for ML-DEC scheme.

Modeling Acid Gas adsorption in Mg-MOF-74

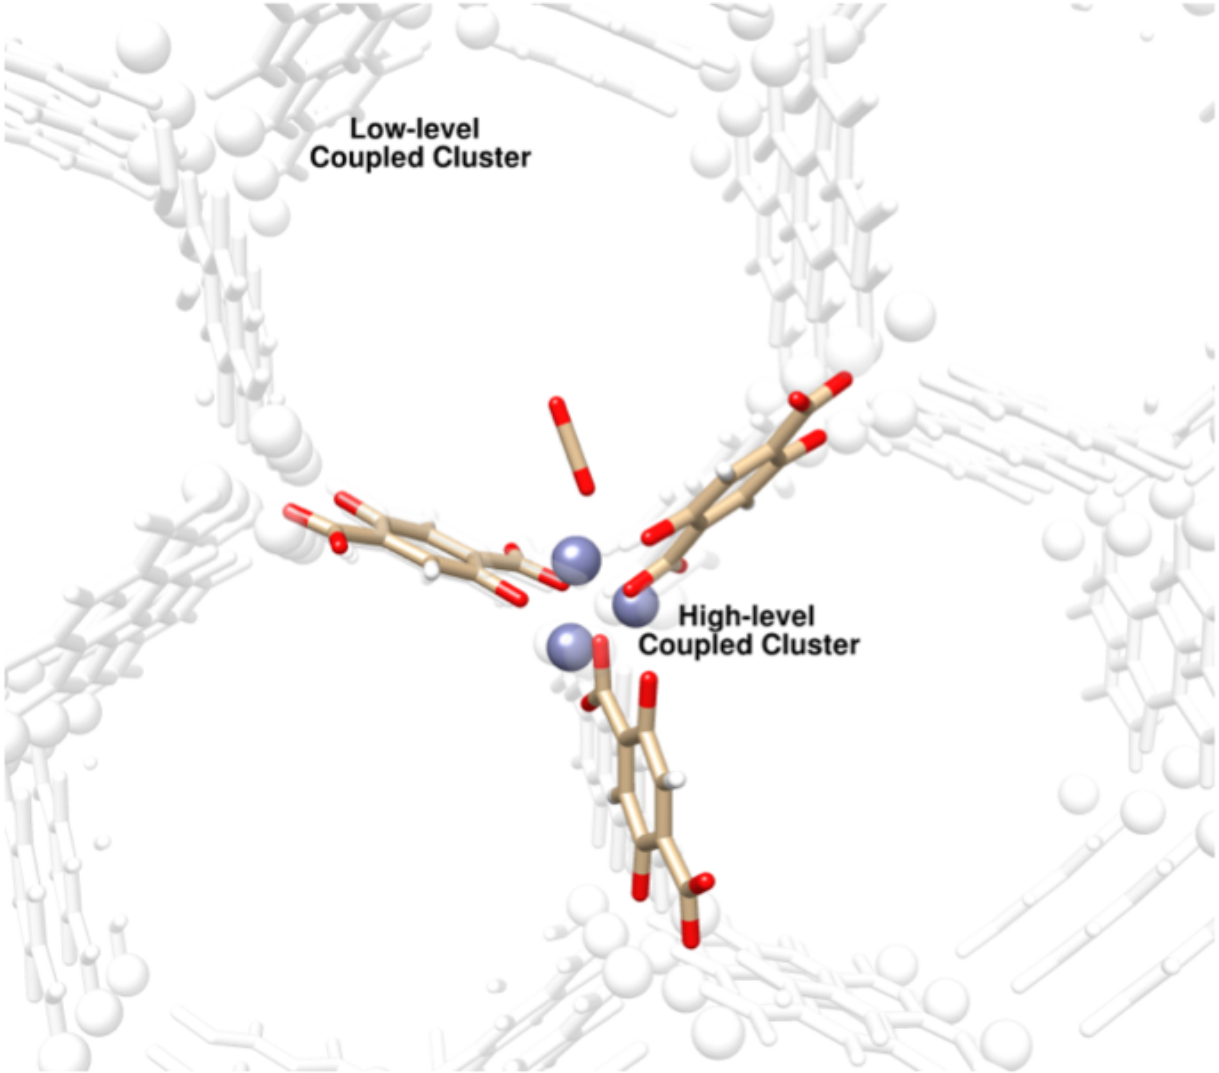

- Plane-wave DFT
  - Results depend on choice of functional, Hubbard U corrections, dispersion corrections etc.
- Small cluster models
  - Suffer from finite size effects. Clusters large enough to replicate environment within the bulk MOF are too expensive for CC.
- Desire to model large finite MOF clusters at CC level of theory while capturing bulk MOF environment.
  - Local interaction – good candidate for ML-DEC scheme.

Mg-MOF-74 + CO_2

HF/RI-MP2

Model	NProc	RI-MP2 frags	Wall time (h)	E_{corr}
DEC MOF	65	76*, 633	3.06	-16.131(15)
ML-DEC MOF	65	76*, 135	2.26	-12.876(15)
DEC MOF+CO_2	65	78*, 689	3.78	-16.745(16)
ML-DEC MOF+CO_2	65	78*, 184	2.84	-13.502(16)

*Atomic fragment optimization

$$\Delta E_{\mathrm{HF/RI\text{-}MP2}} = -0.033$$

$$\Delta E_{\mathrm{RI\text{-}MP2}} = -0.022$$

3132 basis functions
cc-pVTZ, FOT=10^{-4}

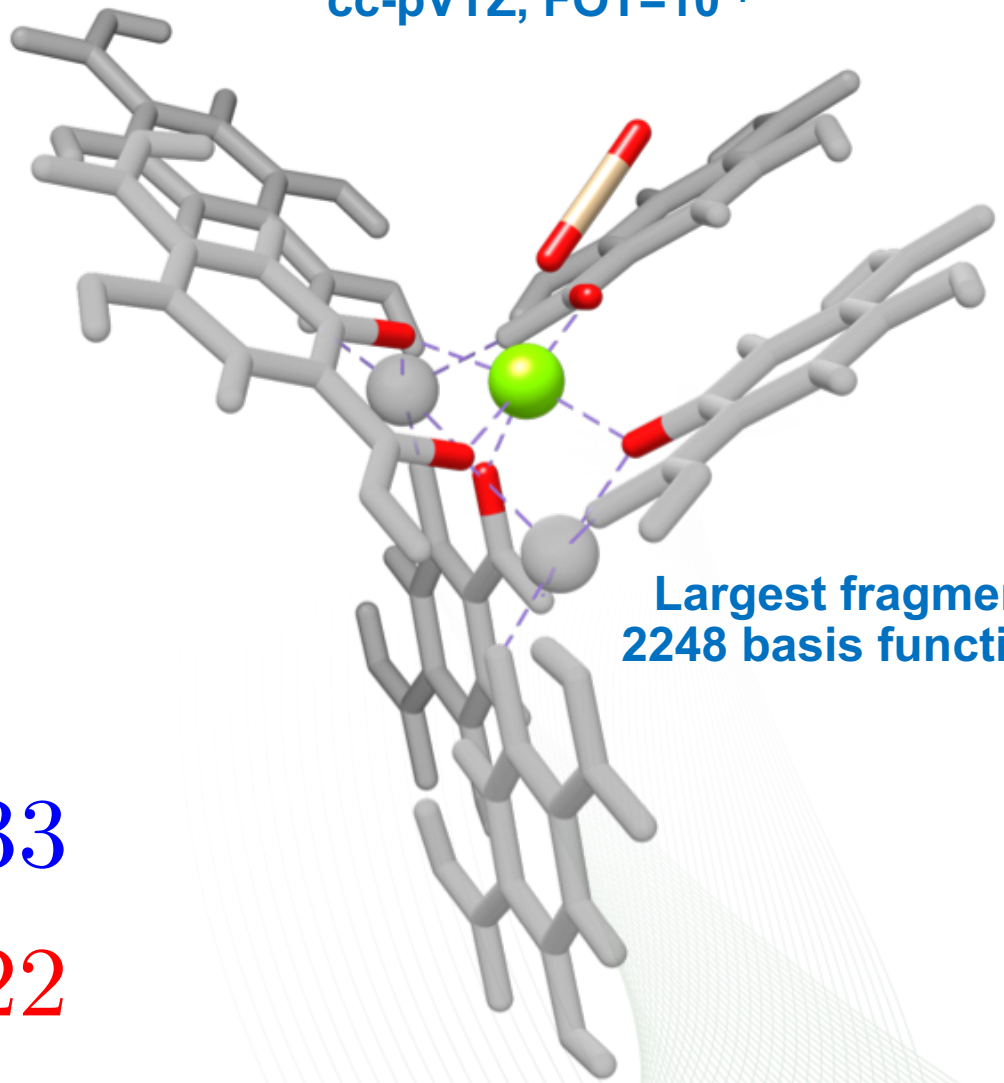

Largest fragment:
2248 basis functions

Mg-MOF-74 + SO_2/H_2O

Interaction energy (E_h)

Gas	HF/RI-MP2	RI-MP2
SO_2	0.028(16)	-0.029(16)
H_2O	-0.099(16)	-0.038(16)
CO_2	-0.033(16)	-0.022(16)

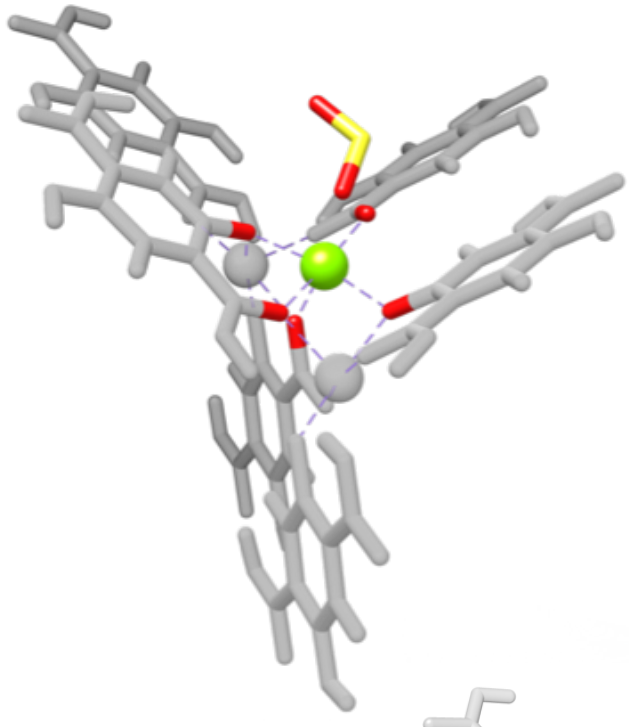

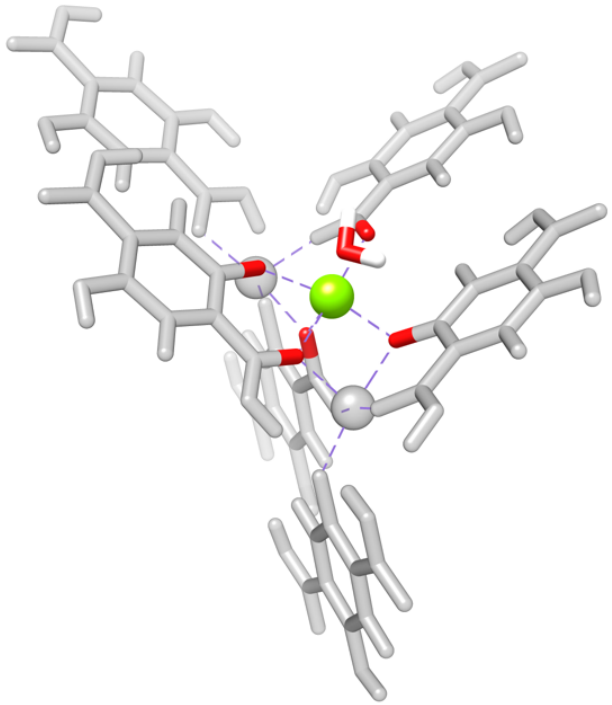

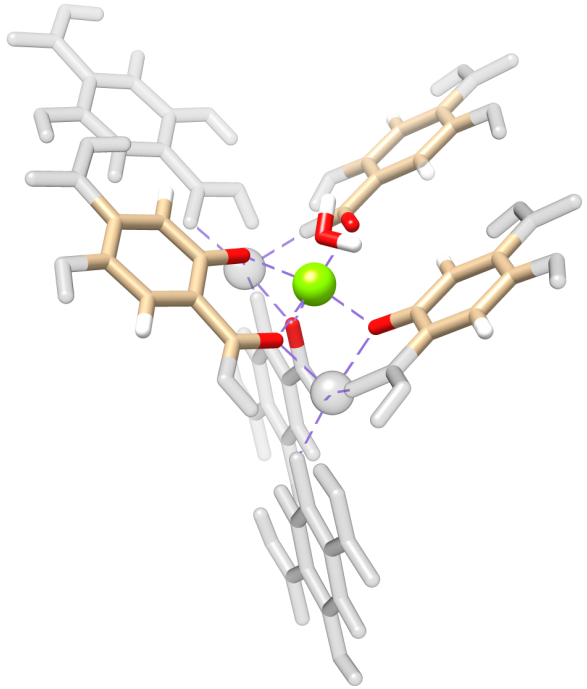

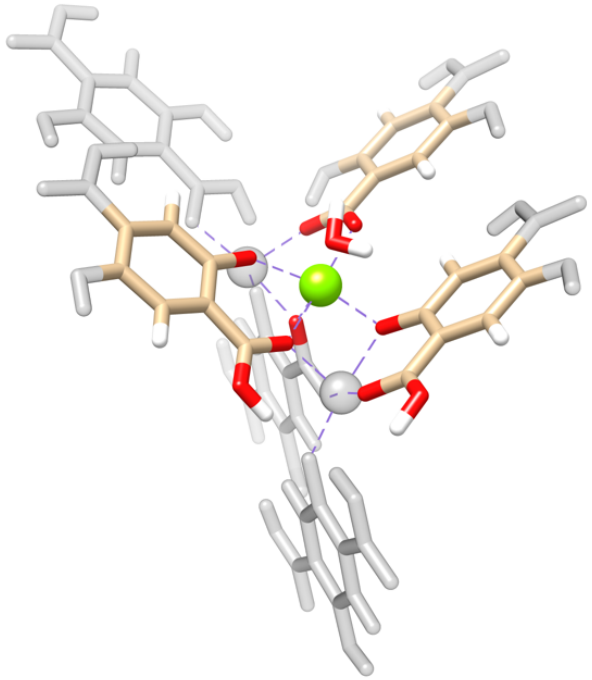

Mg-MOF-74 + SO_2/H_2O

Interaction energy (E_h)

Gas	HF/RI-MP2	RI-MP2
SO_2	**-0.028(16)**	-0.029(16)
H_2O	-0.073(16)	-0.038(16)
CO_2	-0.033(16)	-0.022(16)

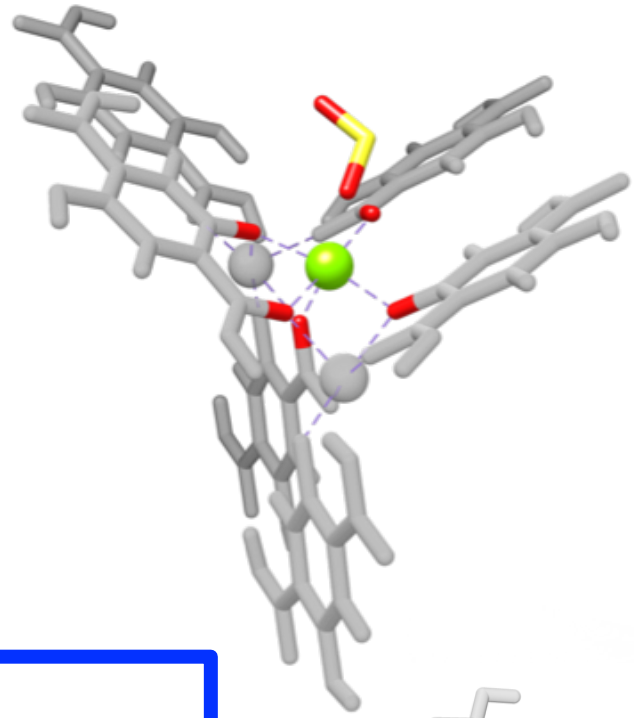

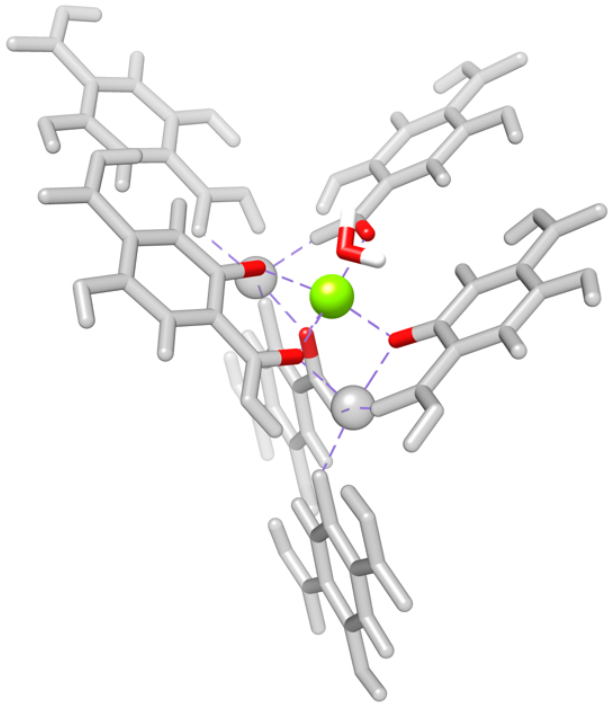

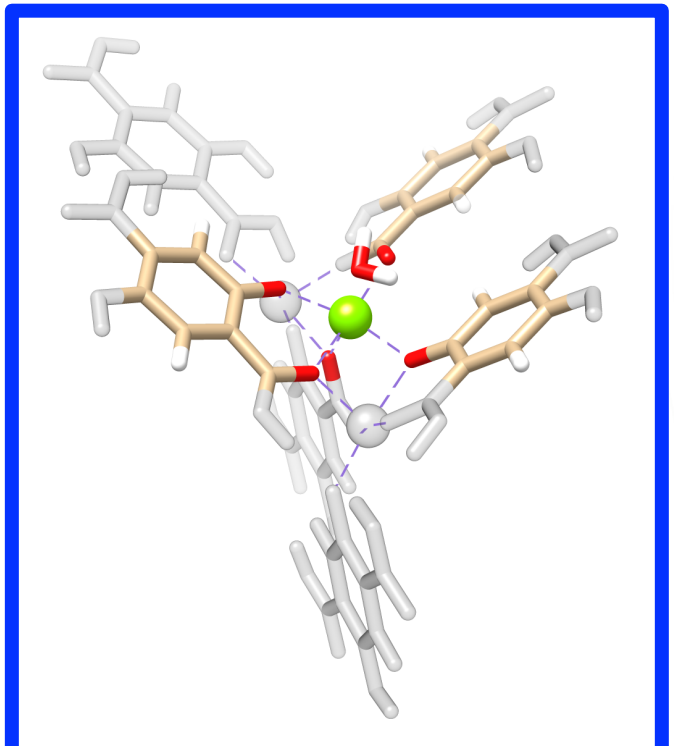

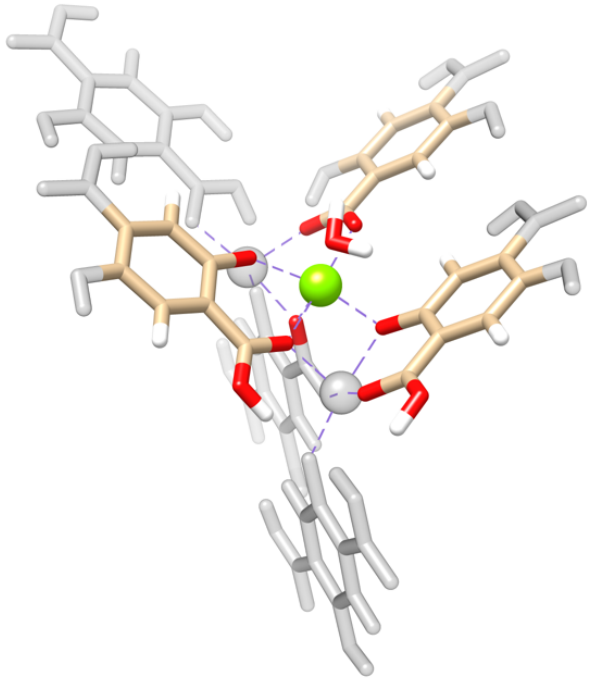

Target system: MOF channel

- HF/RI-MP2 to optimize high level layer.
- RI-MP2/CCSD for final interaction energy calculations.

Model	Fragment Optimizations	Pair Energy Calculations
ML-DEC	202	201
ML-DEC (ext_1)	484	510
DEC-RI-MP2	568	7002
	Most important for HF/RI-MP2	Most important for RI-MP2/CCSD

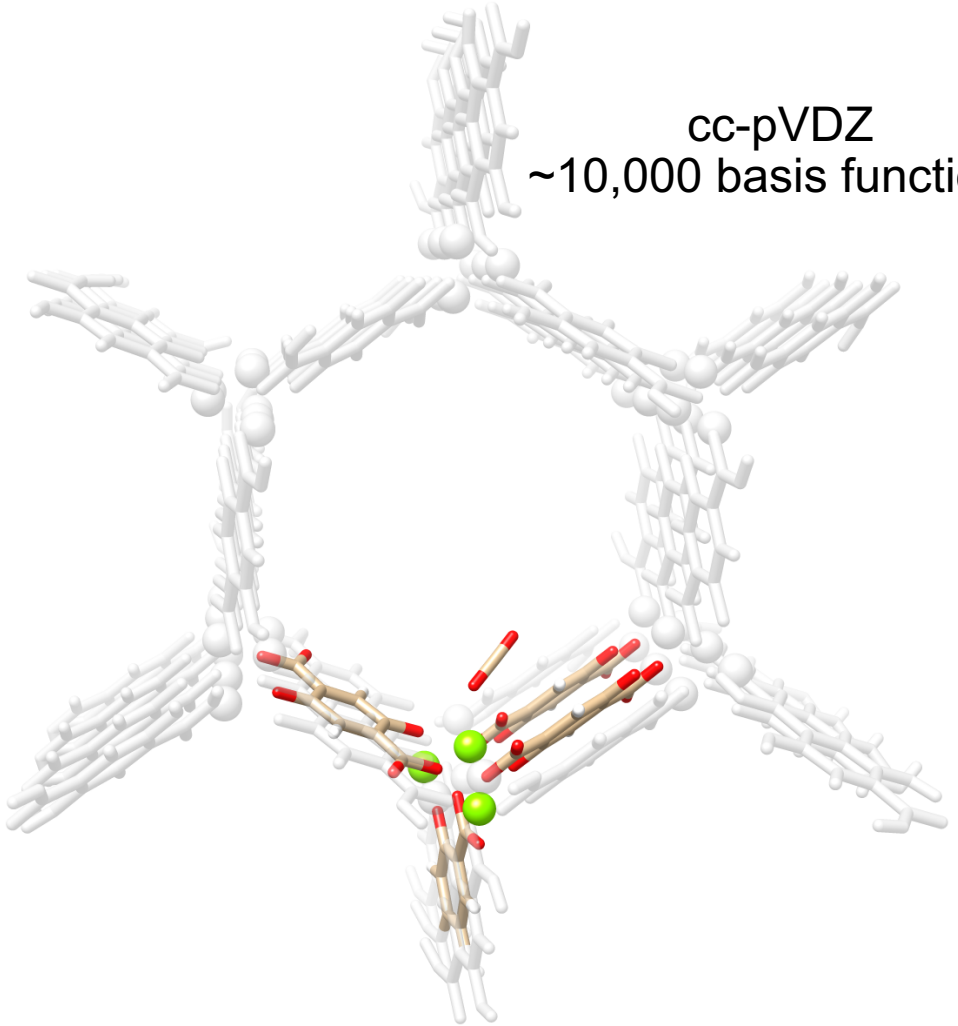

cc-pVDZ
~10,000 basis functions

Targeting Performance on Summit

- LSDalton makes use of MPI, OpenMP, and OpenACC.
  - Linear scaling as long as problem is big enough
- RI-MP2 modules already include OpenACC acceleration
  - jsrun –n 6 –a 1 –g 1 –r 6 –c 7 –brs
  - >2x speedup using 1 GPU/rs compared to 0 GPUs/rs
- Current efforts focus on integration with the TAL-SH tensor algebra library for GPU-acceleration of CC modules (https://github.com/DmitryLyakh/TAL_SH)

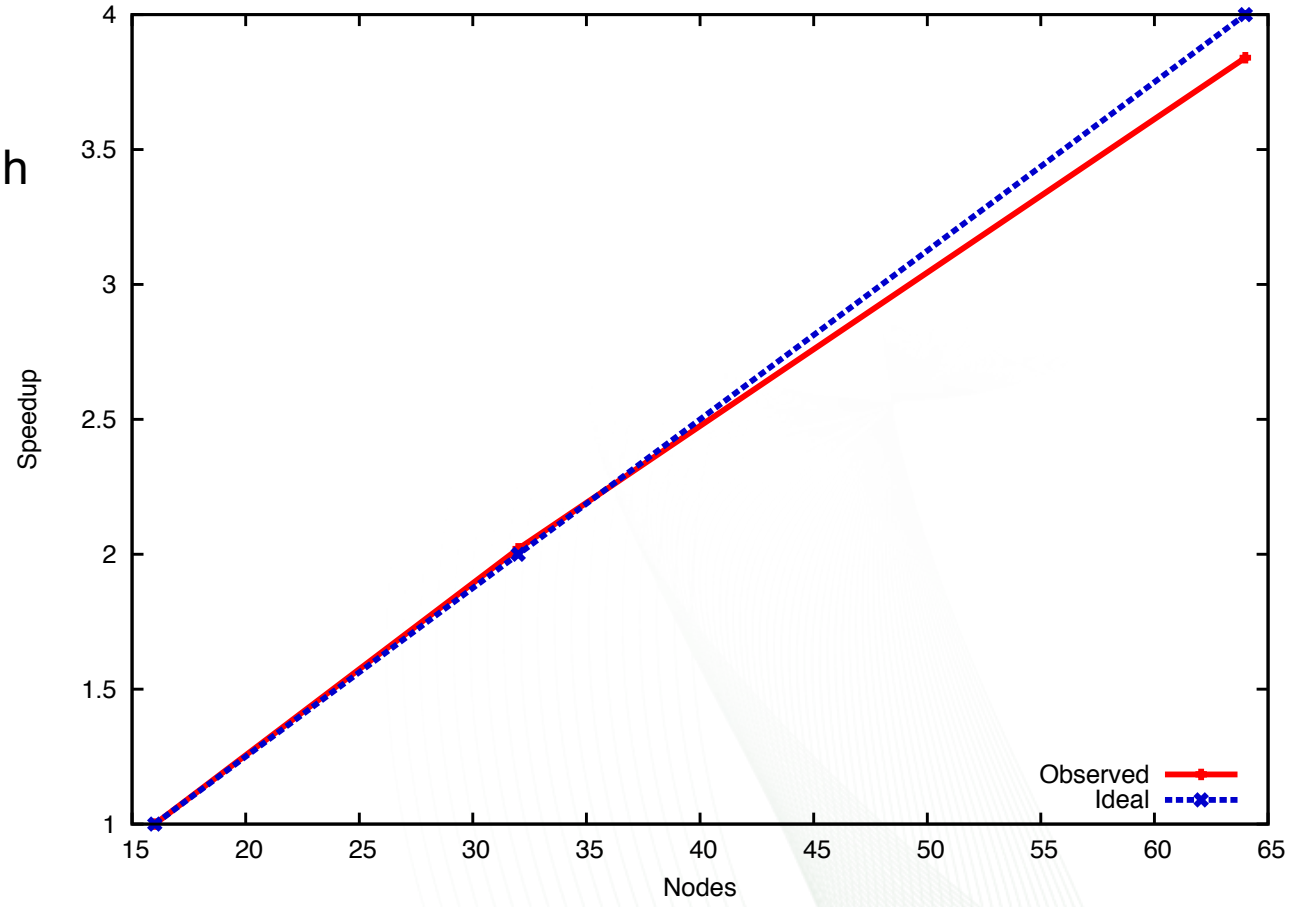

GPU Offloading of Tensor Contractions with TAL-SH

- CC calculations require evaluation of tensor contractions. Efficient handling of these contractions is key to good performance of these modules.
- Currently utilizing ScaTeLib: a Scalable Tensor Library in order to distribute tensor contractions across multiple ranks via tiling.
- E.g.: $$C_{ijkl} = \alpha \sum_{a,b} A_{ajlb} B_{kbai} + \beta C_{ijkl}$$
  - A and B must be sorted appropriately: $A_{ajlb} \Rightarrow A_{jlab}$, $B_{kbai} \Rightarrow B_{abki}$
  - Call dgemm
  - Returns contribution to local C tile as C_{jlki}. Reorder $\Rightarrow C_{ijkl}$ and add to local tile.

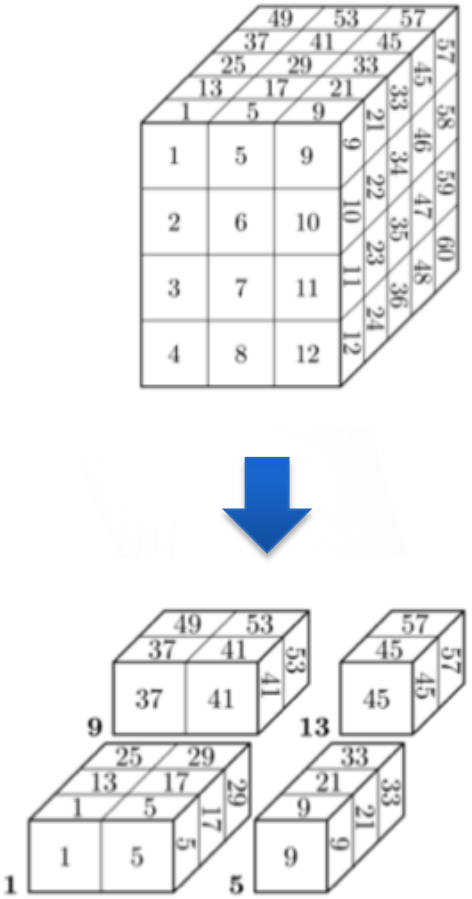

GPU Offloading of Tensor Contractions with TAL-SH

- TAL-SH: Tensor Algebra Library for Shared-memory systems. Integrated as backend for ScaTeLib.

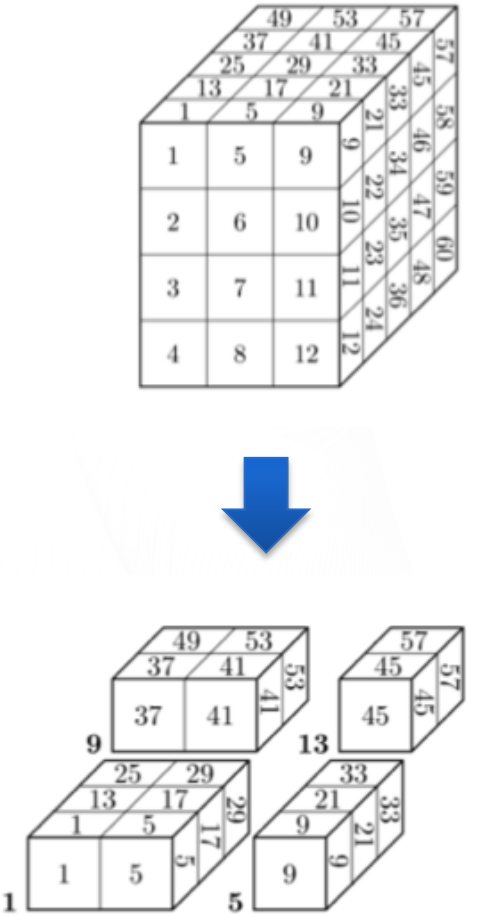

GPU Offloading of Tensor Contractions with TAL-SH

- TAL-SH: Tensor Algebra Library for Shared-memory systems. Integrated as backend for ScaTeLib.

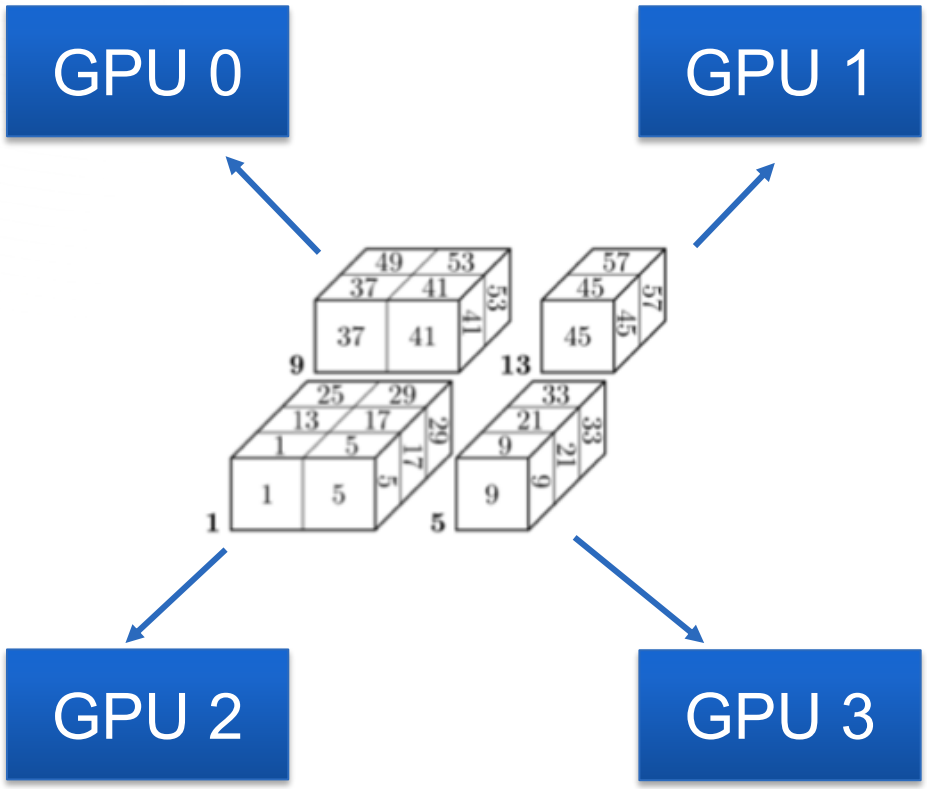

- Handles sorting, no explicit reorder calls necessary in ScaTeLib
- Asynchronous task scheduling
- Tasks are pipelined to overlap computation and data transfer
  - 2 active tasks per GPU at any time

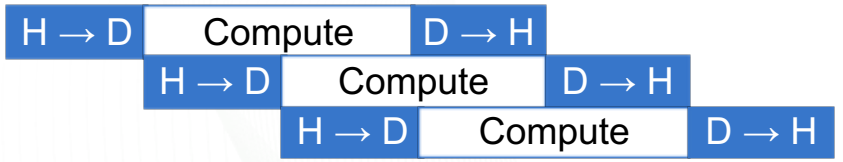

- Basic implementation provides ~2.25x speedup for contraction of two 4th-order tensors with tile size of 80 (Titan).

GPU Offloading of Tensor Contractions with TAL-SH

- TAL-SH: Tensor Algebra Library for Shared-memory systems. Integrated as backend for ScaTeLib.

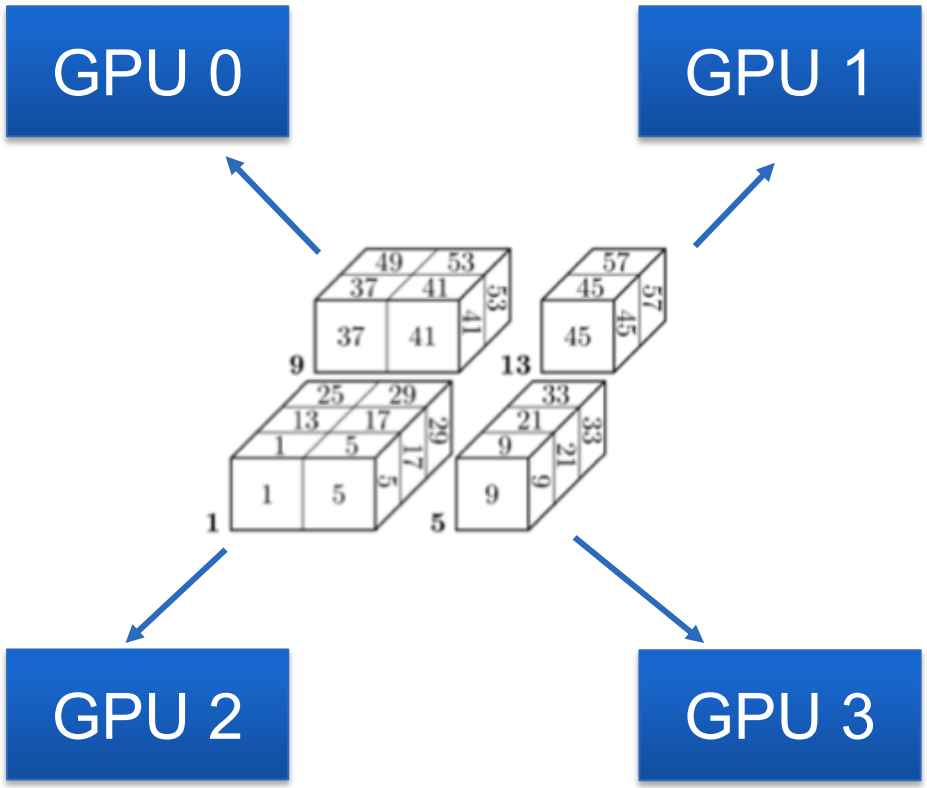

- Targeted performance improvements:
  - OpenMP compatibility
  - Reduction of data transfer
  - Optimization of tile sizes used within LSDalton
    - Ensure time spent in data transfer < time spent in computation

Acknowledgements

Dmytro Bykov: LSDALTON CAAR Liason

Dmitry Lyakh: TAL-SH (https://github.com/DmitryLyakh/TAL_SH)

This research used resources of the Oak Ridge Leadership Computing Facility, which is a DOE Office of Science User Facility supported under Contract DE-AC05-00OR22725