COMPUTATIONAL CHEMISTRY AT THE SUPEK SUPERCOMPUTER

Emir Imamagić [a], Kristijan Mrkalj [a], Tomica Hrenar [b], Darko Babić [c] and Borislav Kovačević [c]

[a] University Computing Centre – SRCE, University of Zagreb, Josipa Marohnića 5, Zagreb, Croatia

[b] Department of Chemistry, Faculty of Science, University of Zagreb, Horvatovac 102a, Zagreb, Croatia

[c] Department of Physical Chemistry, Ruđer Bošković Institute, Bijenička 54, Zagreb, Croatia

INTRO

- Scalability of computational chemistry applications on Supek supercomputer
- How to get the most out of the application and the cluster itself?
- Sensitive nature of licence agreements
  - No „explicit” comparisons
- Interested in seeing how your calculations will perform?
  - **computing@srce.hr**

SUPEK SUPERCOMPUTER

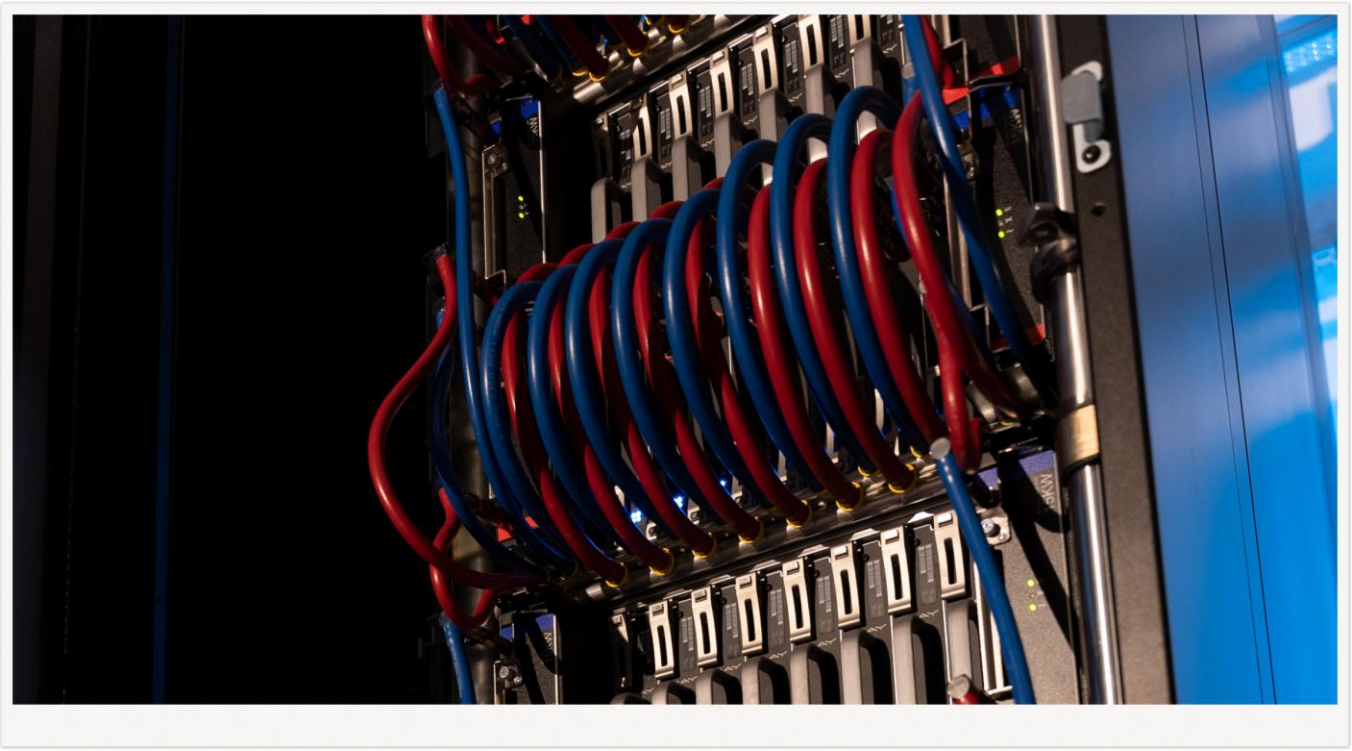

- HPE Cray EX2500
  - single rack, water cooling
- **1 250** TFLOPS R_{max}, **1 830** TFLOPS R_{peak}
- HPE ClusterStor E1000
  - **531** TiB (~**580** TB)
- HPE Slingshot Interconnect
  - **200 Gbit/s**

Purpose	Number	CPU	GPU	RAM [GB]
CPU access node	1	2 × AMD EPYC 7763	-	256
GPU access nodes	1	1 × AMD EPYC 7763	1 × NVIDIA A100 (PCI)	128
CPU work nodes	52	2 × AMD EPYC 7763	-	256
GPU work nodes	20	1 × AMD EPYC 7763	4 × NVIDIA A100 (SXM)	512
Big-memory work nodes	2	2 × AMD EPYC 7763	-	4096

TESTING ENVIRONMENT

Single computational node / CPU work node

- AMD EPYC 7763 (×2)
  - # of CPU cores: 64
  - # of threads: 128
  - Base clock: 2.45 GHz
  - L3 cache: 256 MB
- 256 GB RAM
- HPE Cray ClusterStor E1000 storage system

** illustration*

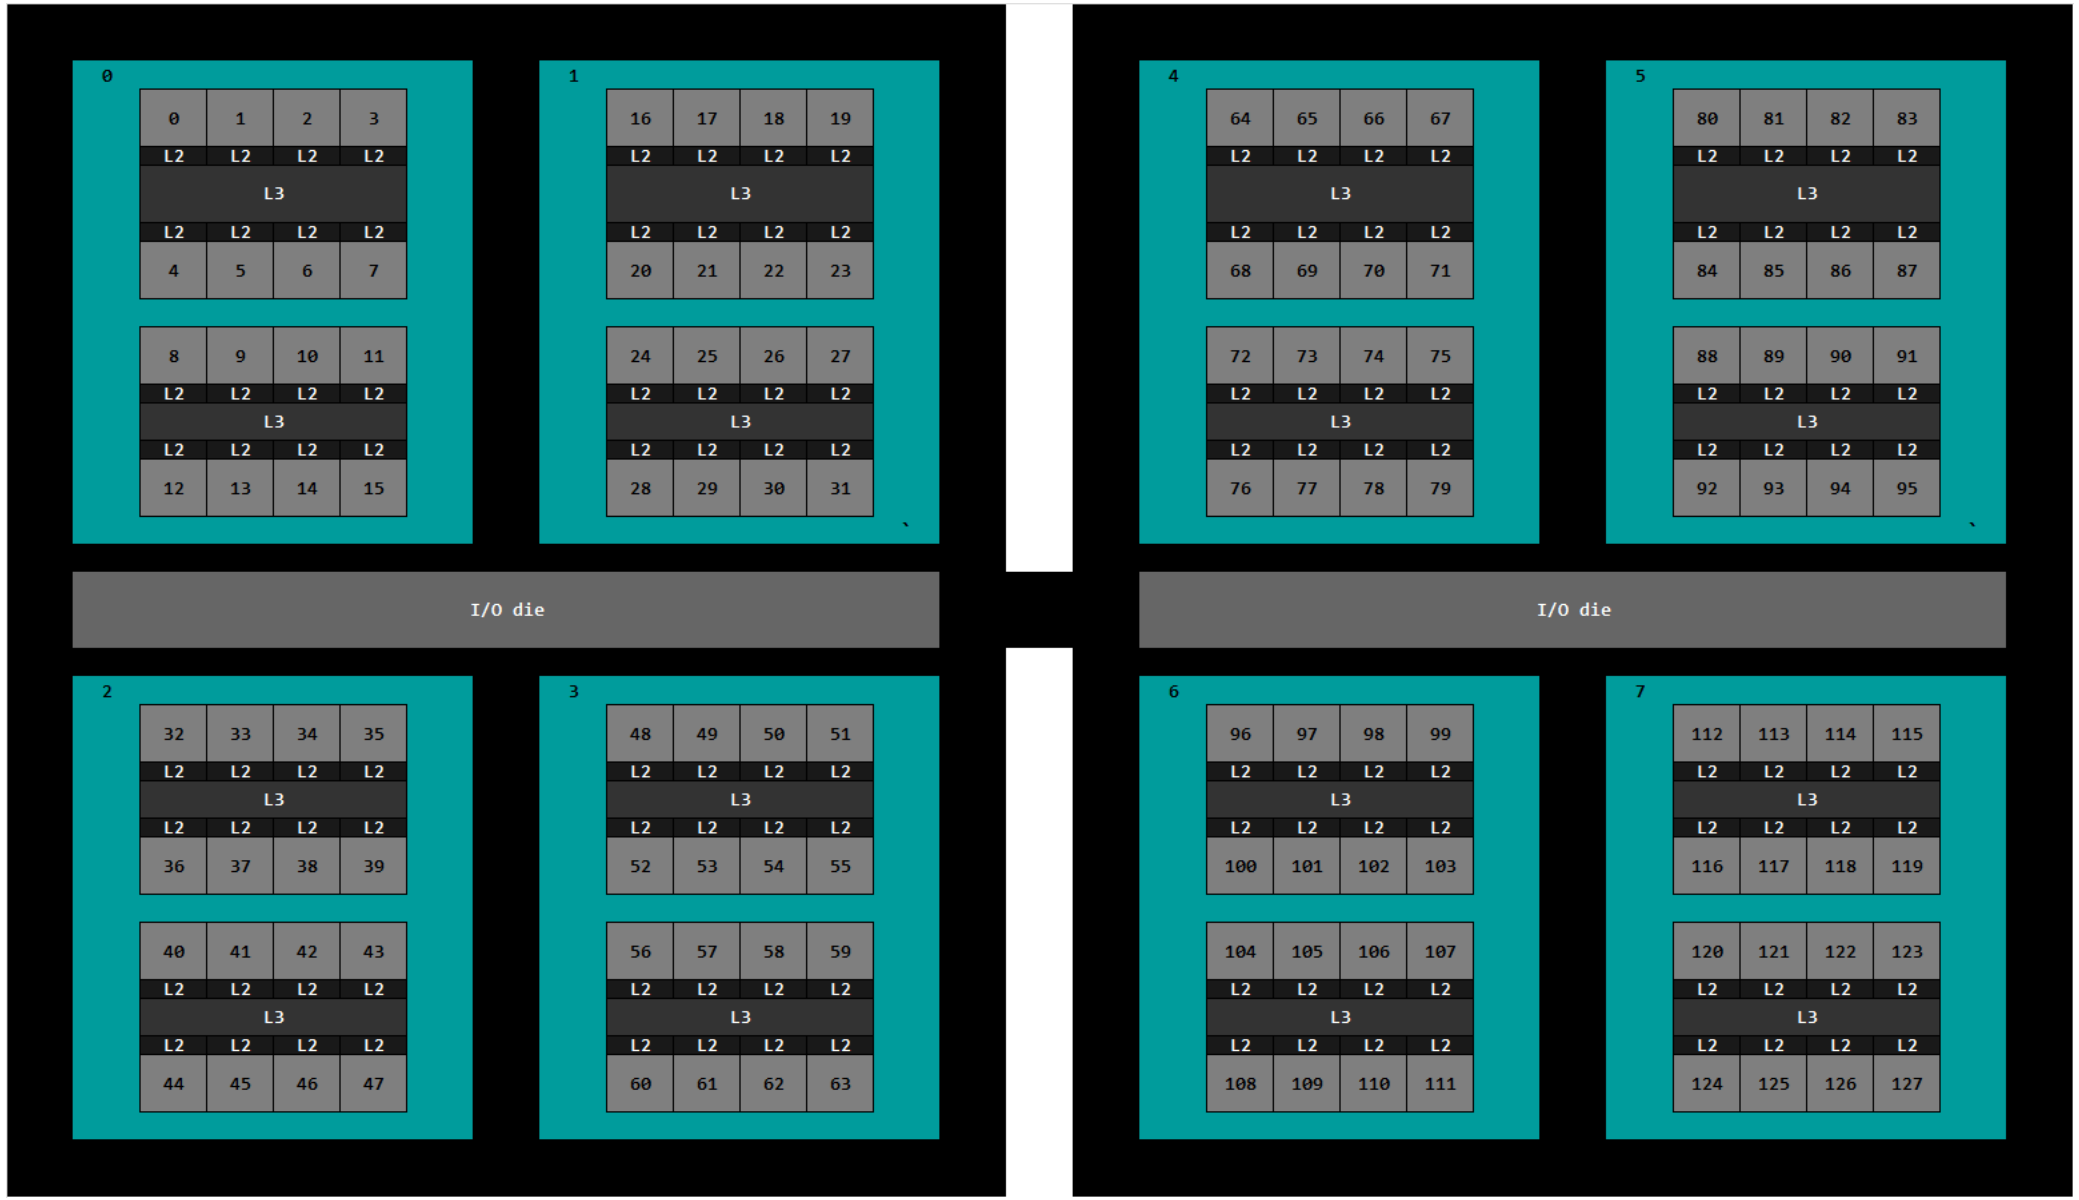

Simplified schematic of two AMD EPYC 7763 CPUs (two sockets)

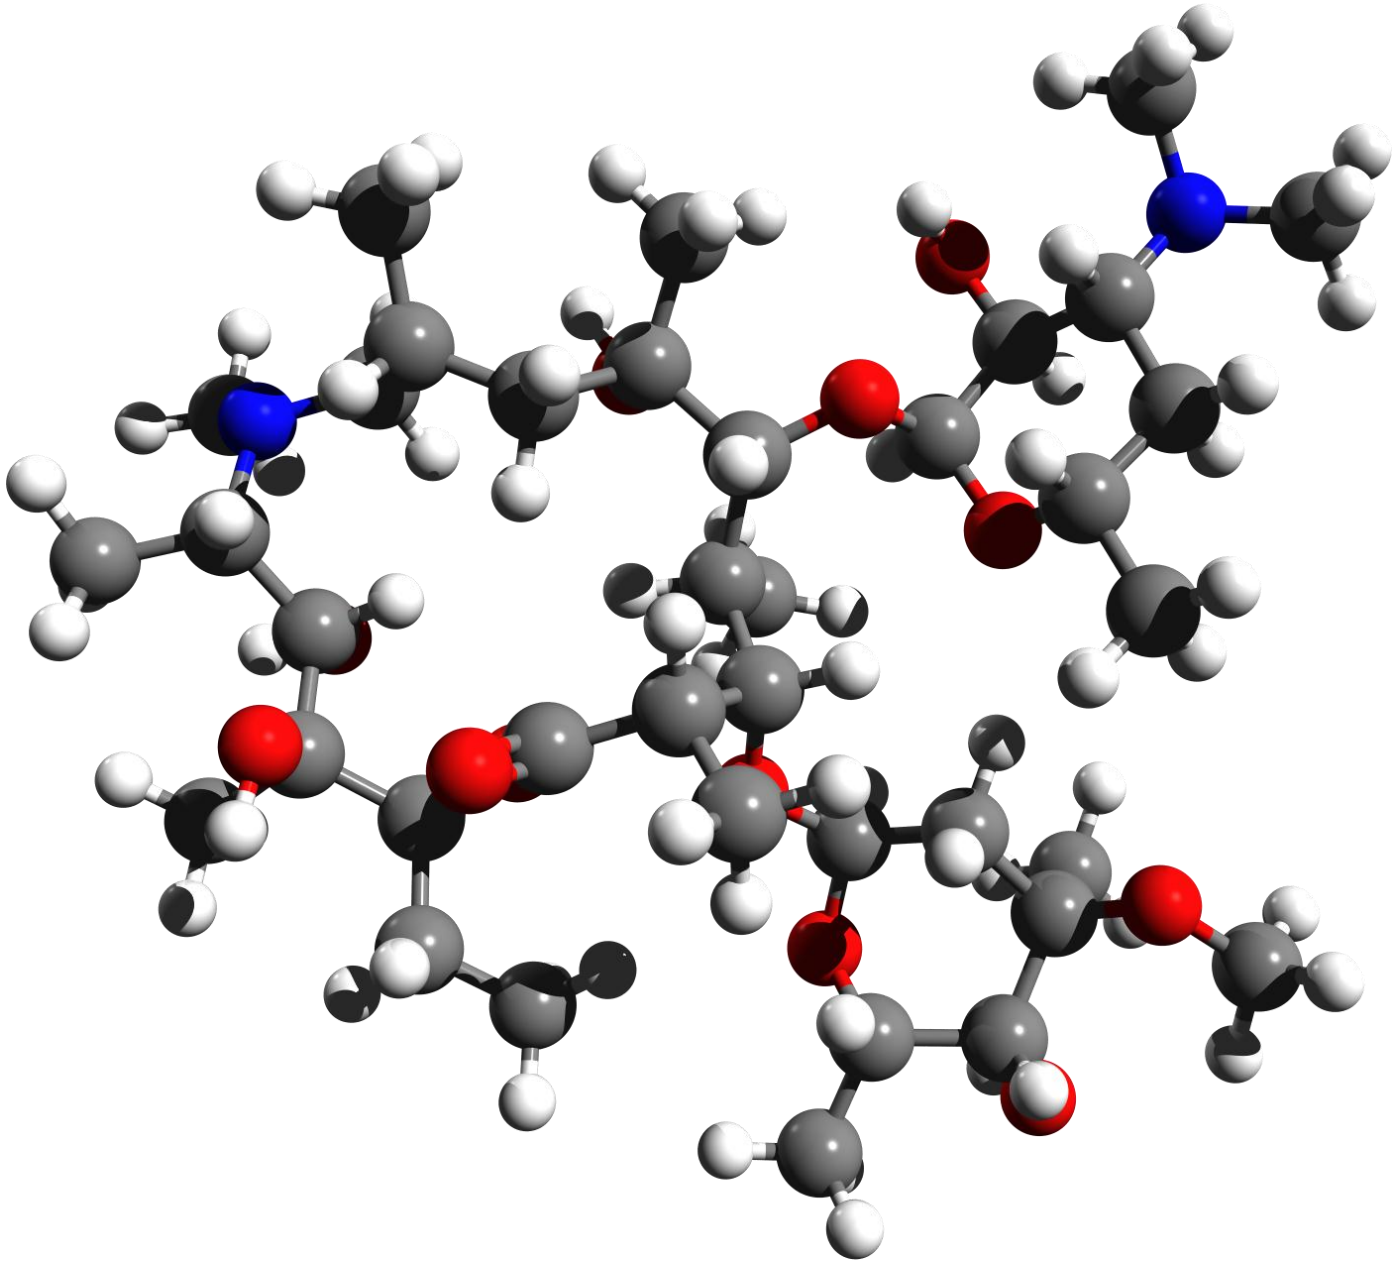

MOLECULE OF CHOICE

- Azythromicin (*Sumamed*)
  - 124 atoms, $C_{38}H_{72}N_2O_{12}$

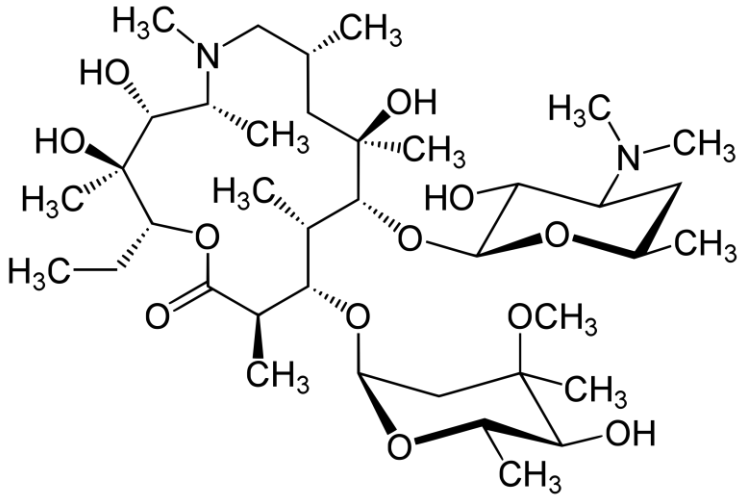

CHOSEN CALCULATIONS

ENERGY AND GRADIENT

- DFT functional: B3LYP
- Basis set: 6-31G(d,p)
- Correction: D3BJ

HARMONIC FREQUENCIES

- DFT functional: B3LYP
- Basis set: 6-31G(d,p)
- Correction: D3BJ

NMR SHIELDINGS (CHEMICAL SHIFTS)

- DFT functional: B3LYP
- Basis set: 6-311+G(d,p)

TD-DFT FOR THE FIRST 8 EXCITED SINGLET STATES

- DFT functional: B3LYP
- Basis set: 6-31+G(d,p)

ENERGY AND GRADIENT

B3LYP, 6-31G(d,p), D3BJ

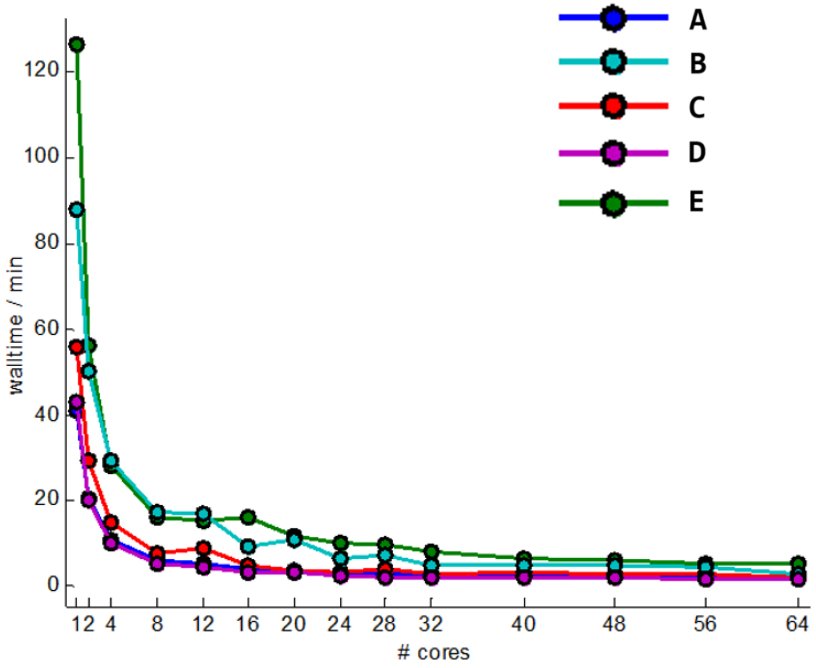

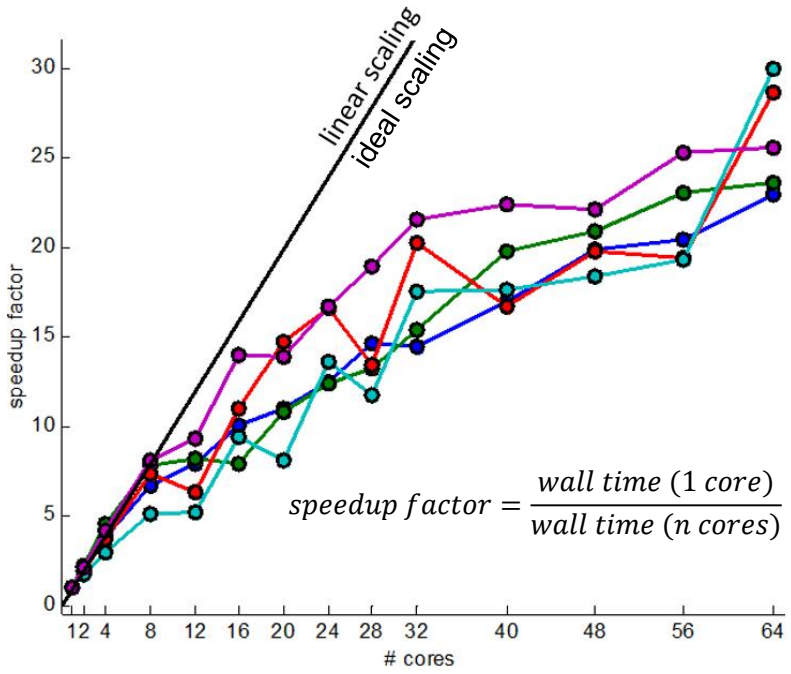

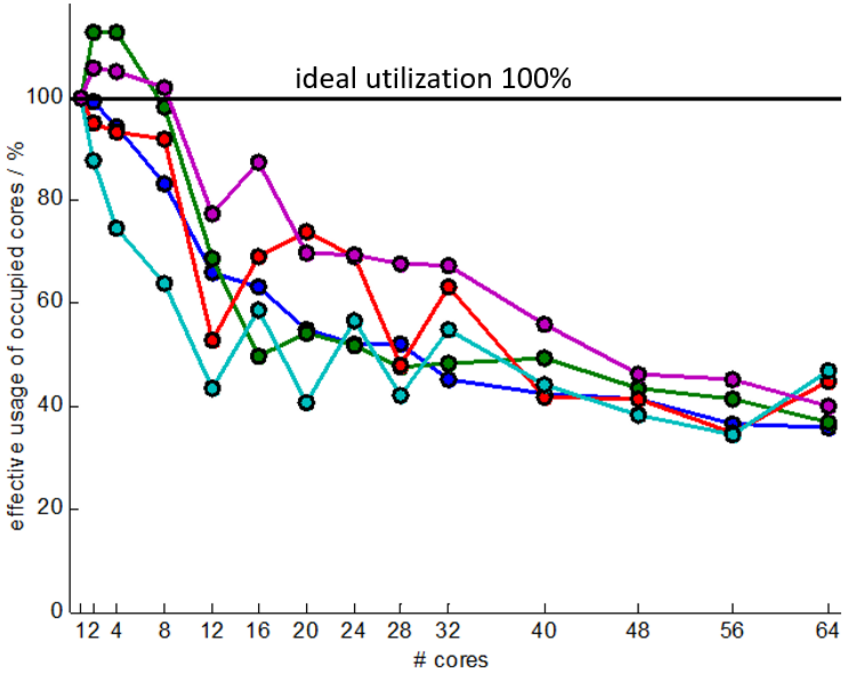

Total SCF energy calculated:

-2503.89260524 a.u.

-2503.892602743818 a.u. (without any "RI" approximation)

-2503.895658476004 a.u. (with "RIJCOSX" approximation)

-2503.8930727146 a.u.

-2503.8927872816 a.u.

HARMONIC FREQUENCIES

B3LYP, 6-31G(d,p), D3BJ

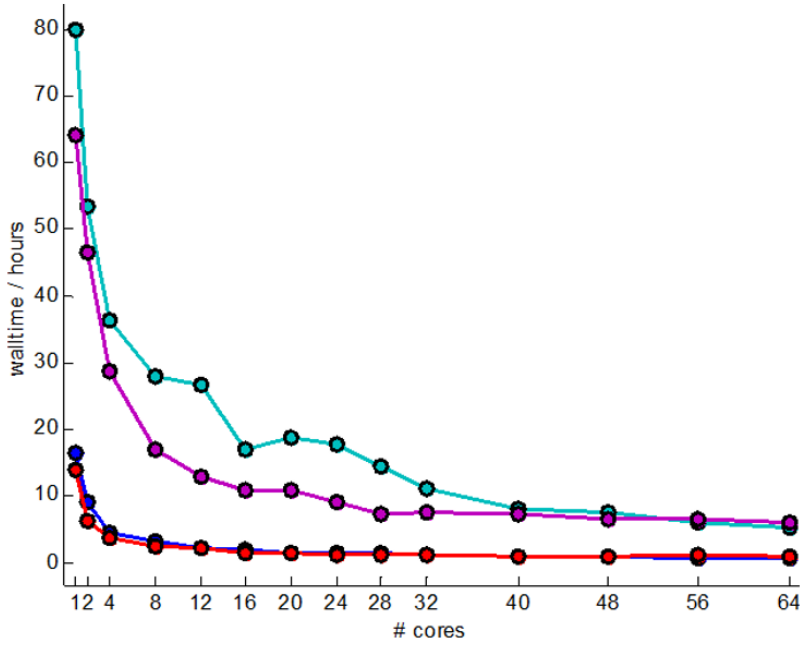

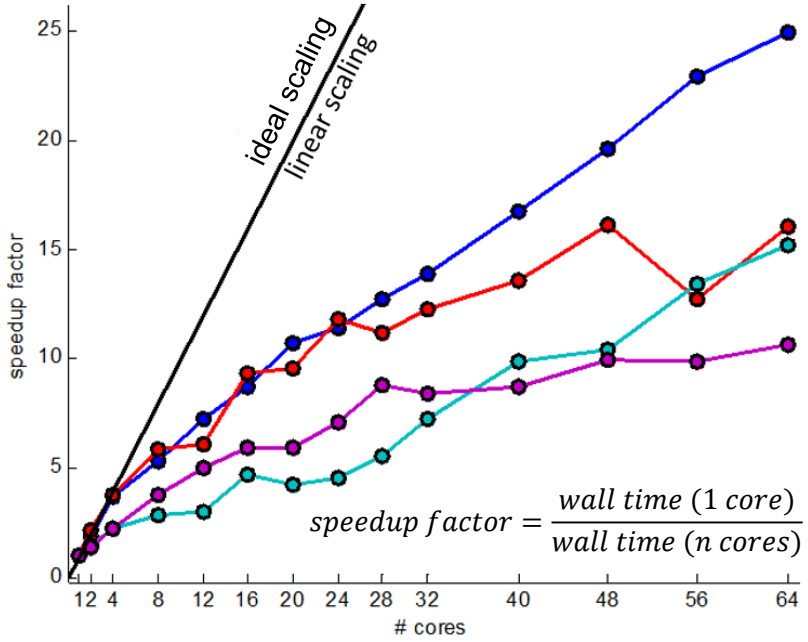

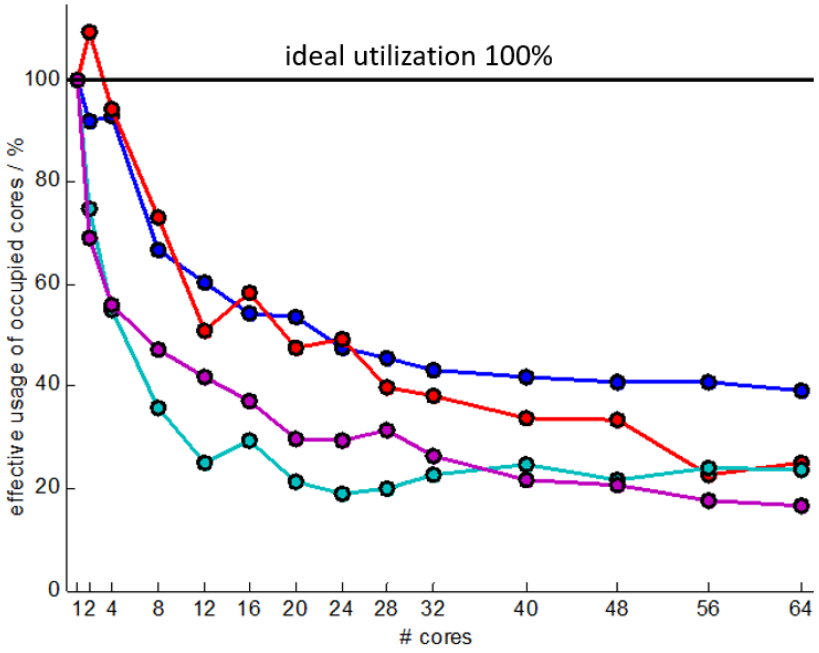

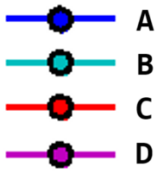

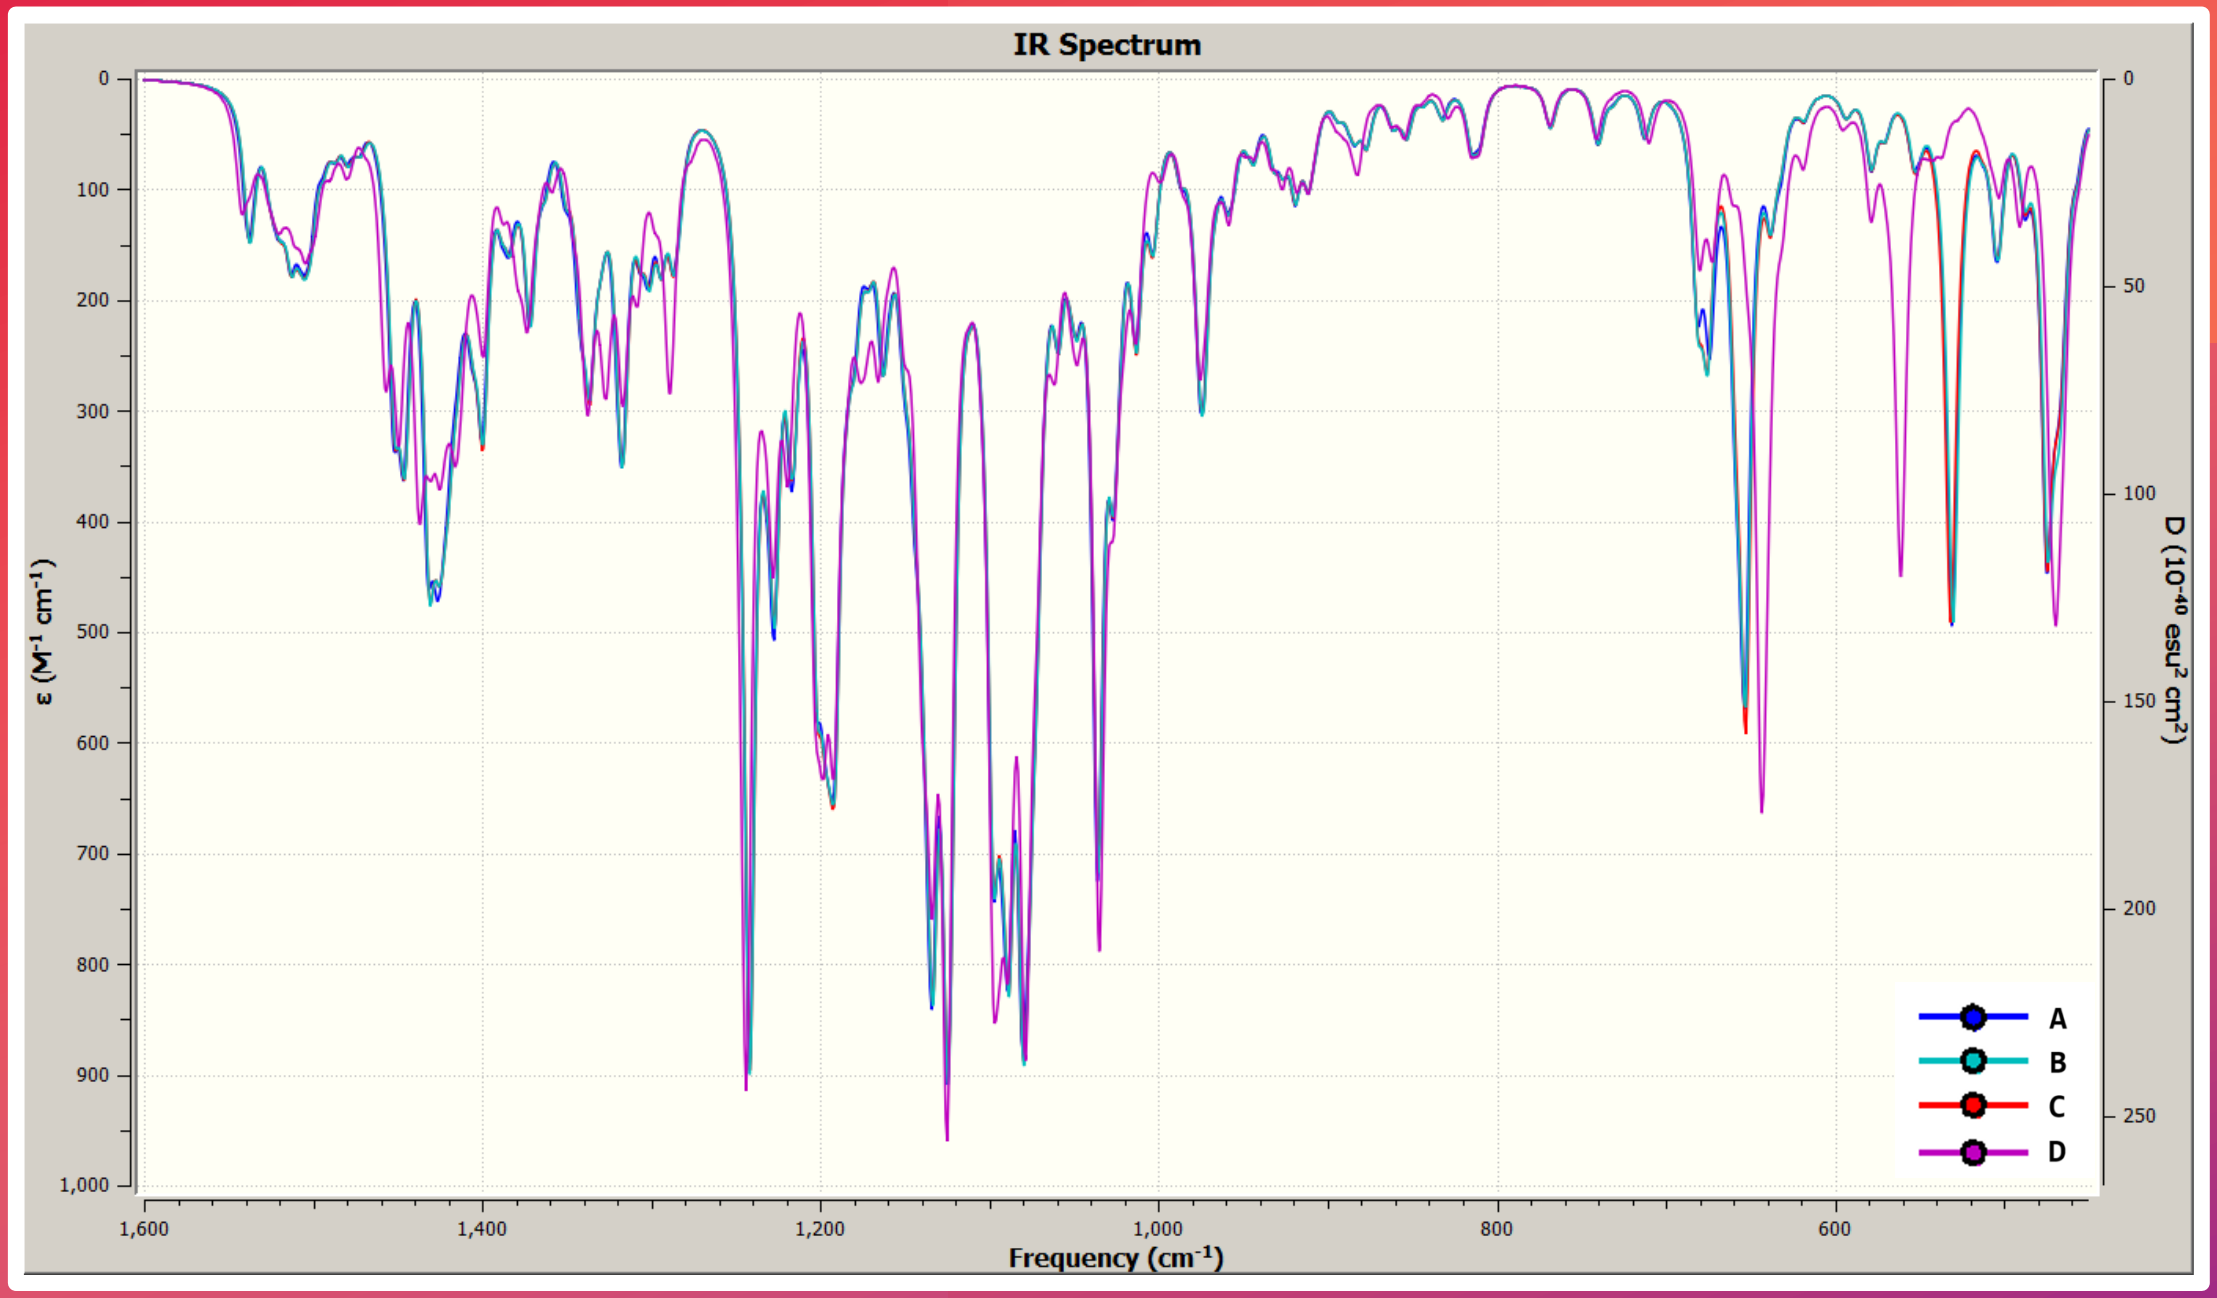
IR Spectrum
ε (M-1 cm-1)
D (10-40 esu2 cm2)
Frequency (cm-1)
0
100
200
300
400
500
600
700
800
900
1,000
50
150
250
1,600
1,400
1,200
1,000
800
600
A
B
C
D

NMR SHIELDINGS (CHEMICAL SHIFTS)

B3LYP, 6-311+G(d,p)

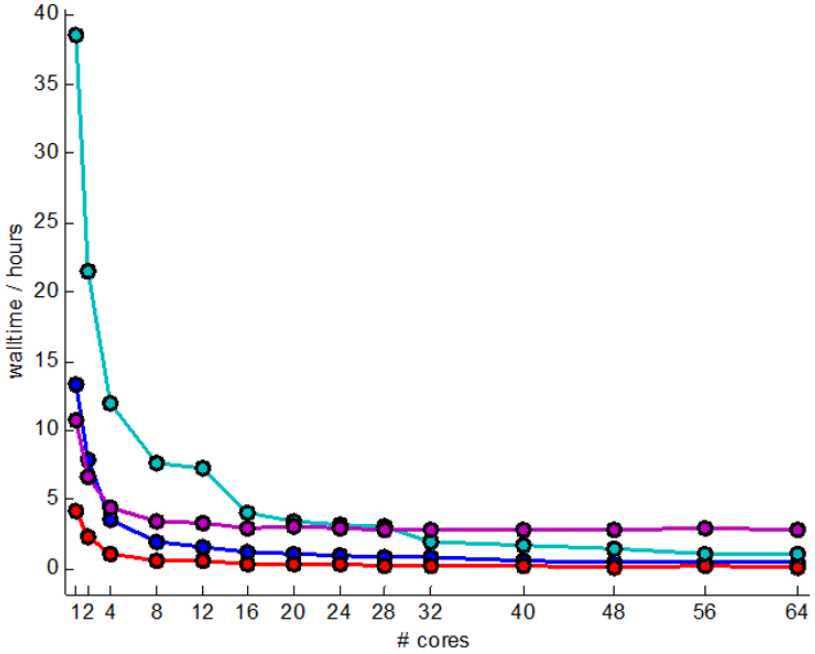

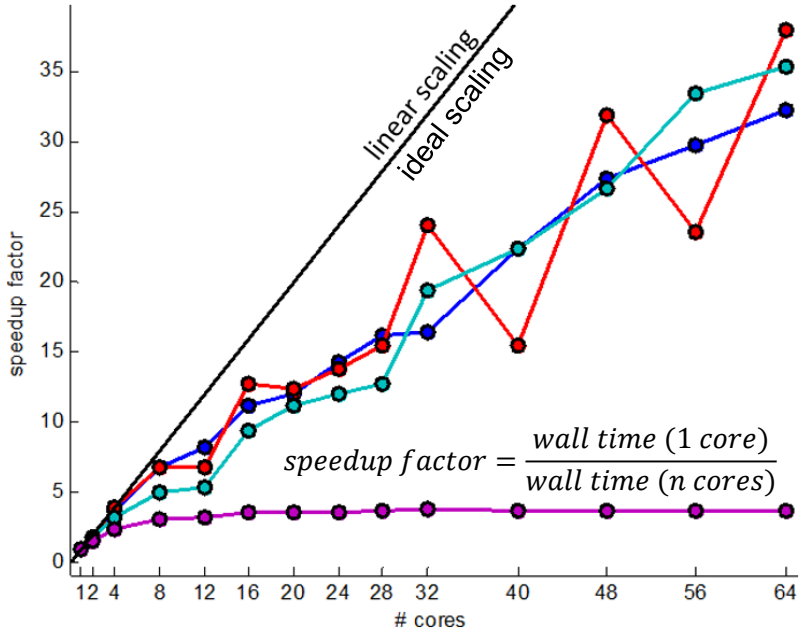

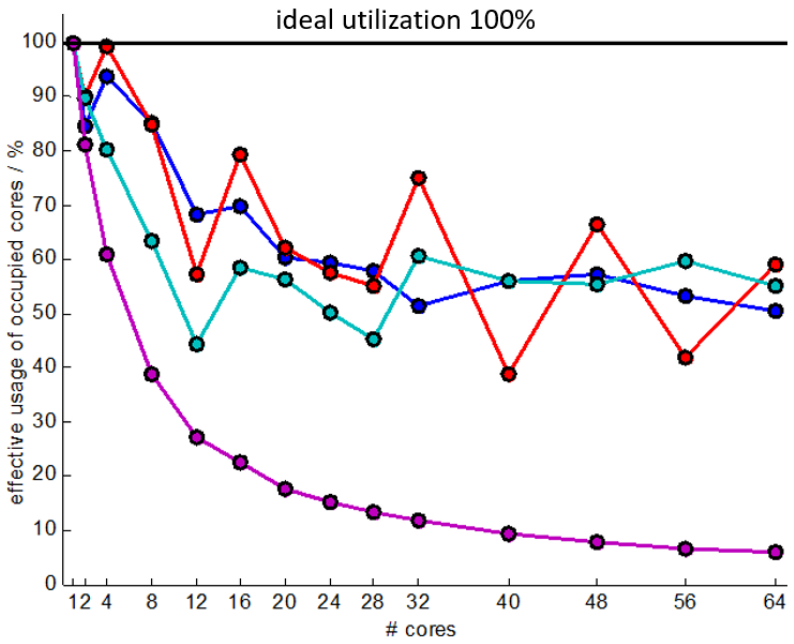

A
B
C
D

calculated NMR shieldings for C atoms

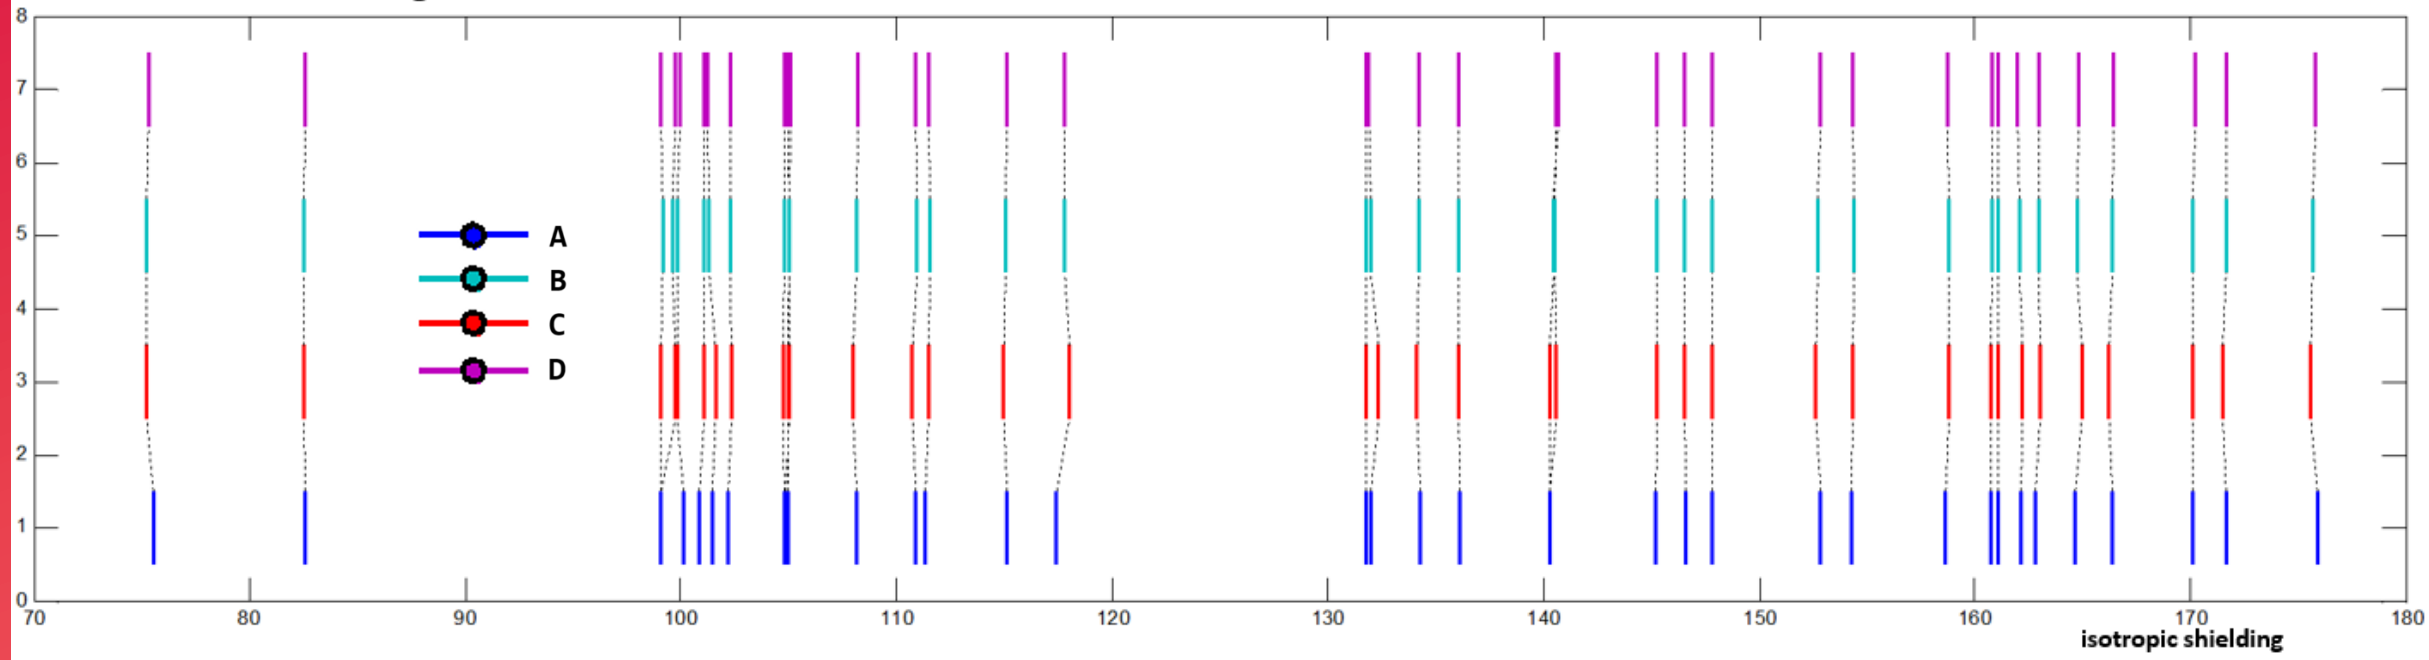

calculated NMR shieldings for H atoms

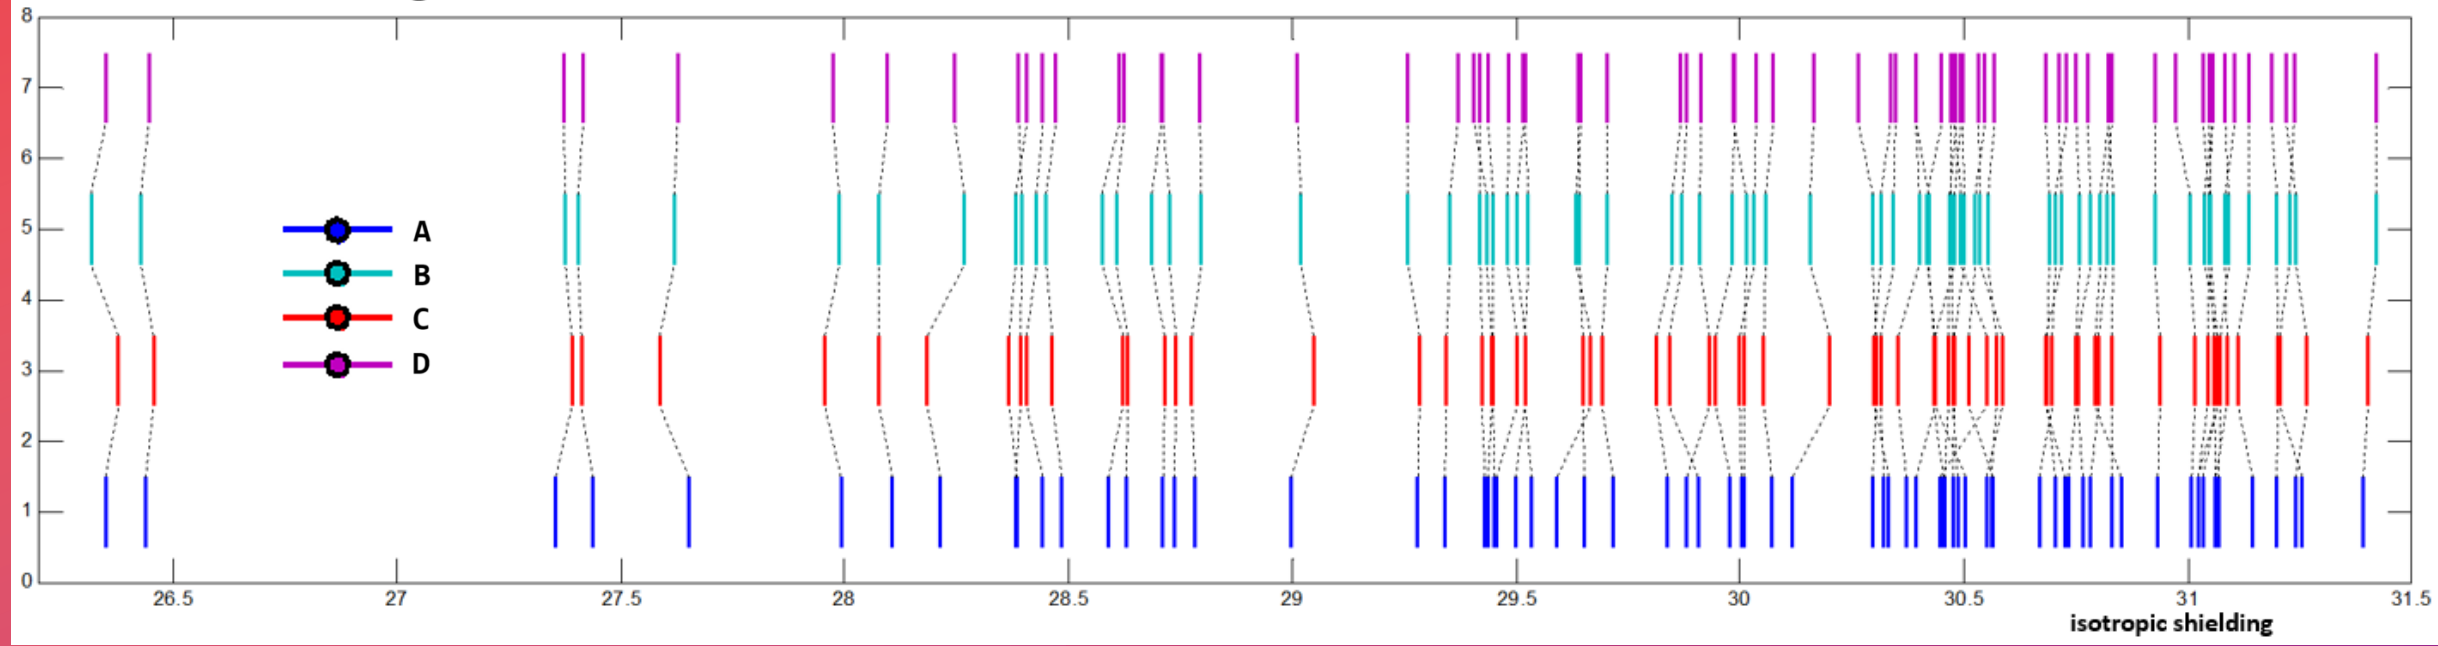

TD-DFT FOR THE FIRST 8 EXCITED SINGLET STATES

B3LYP, 6-31+G(d,p)

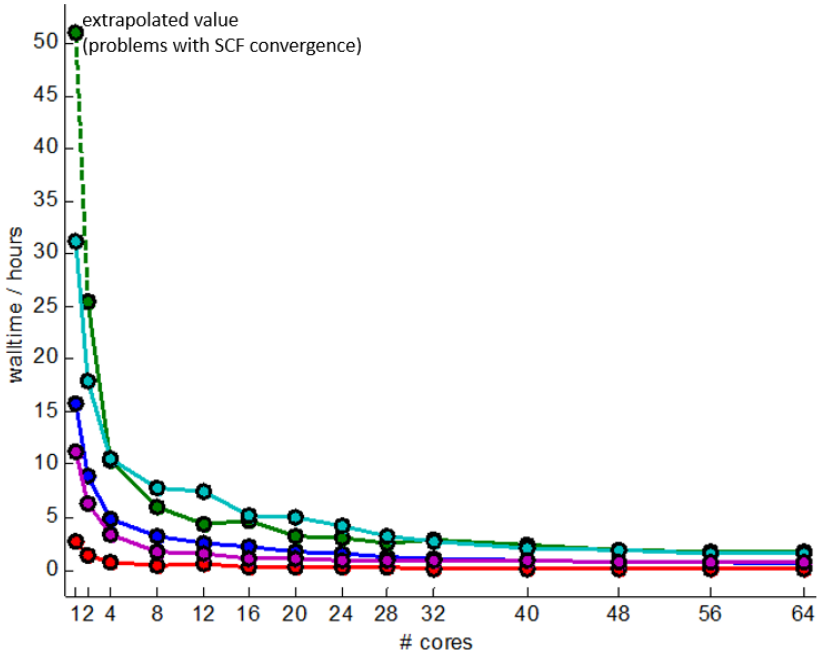

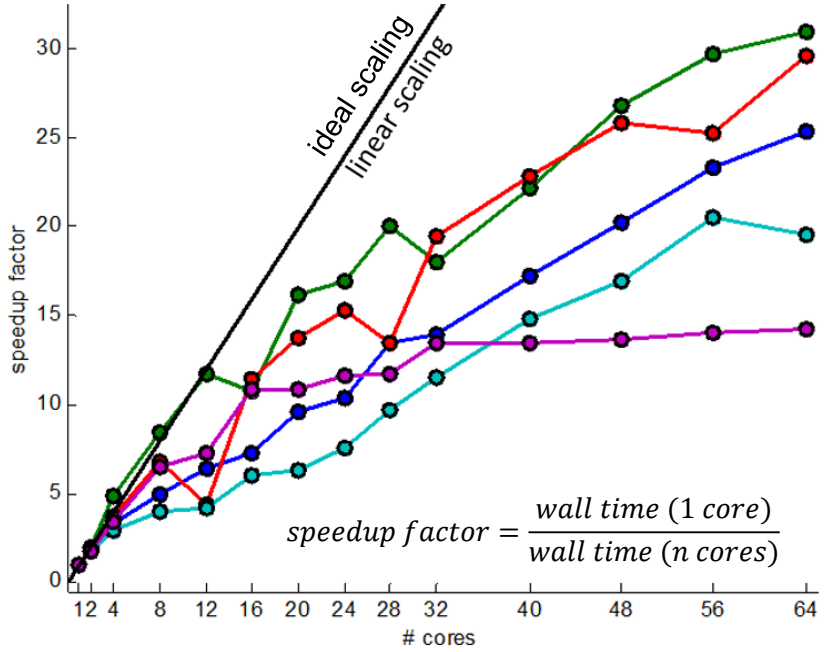

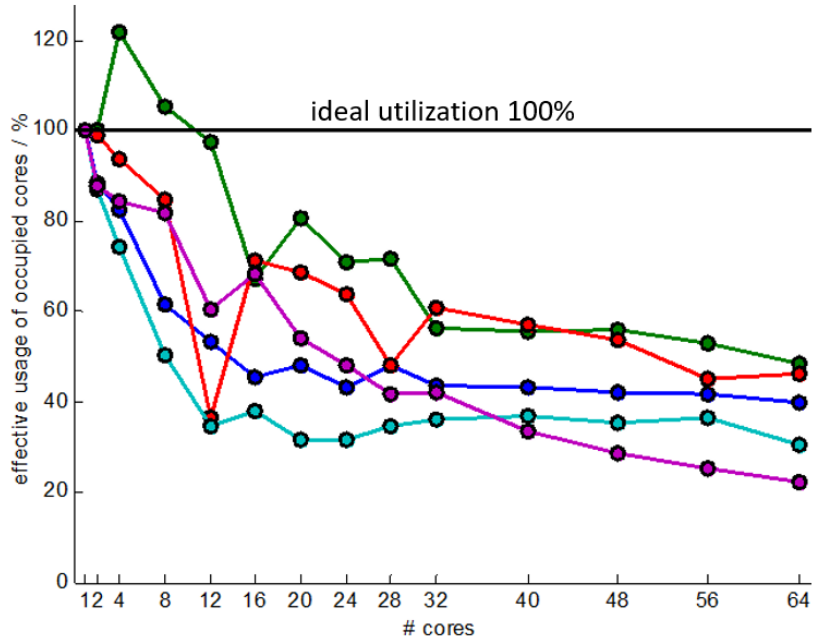

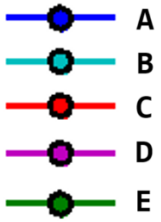

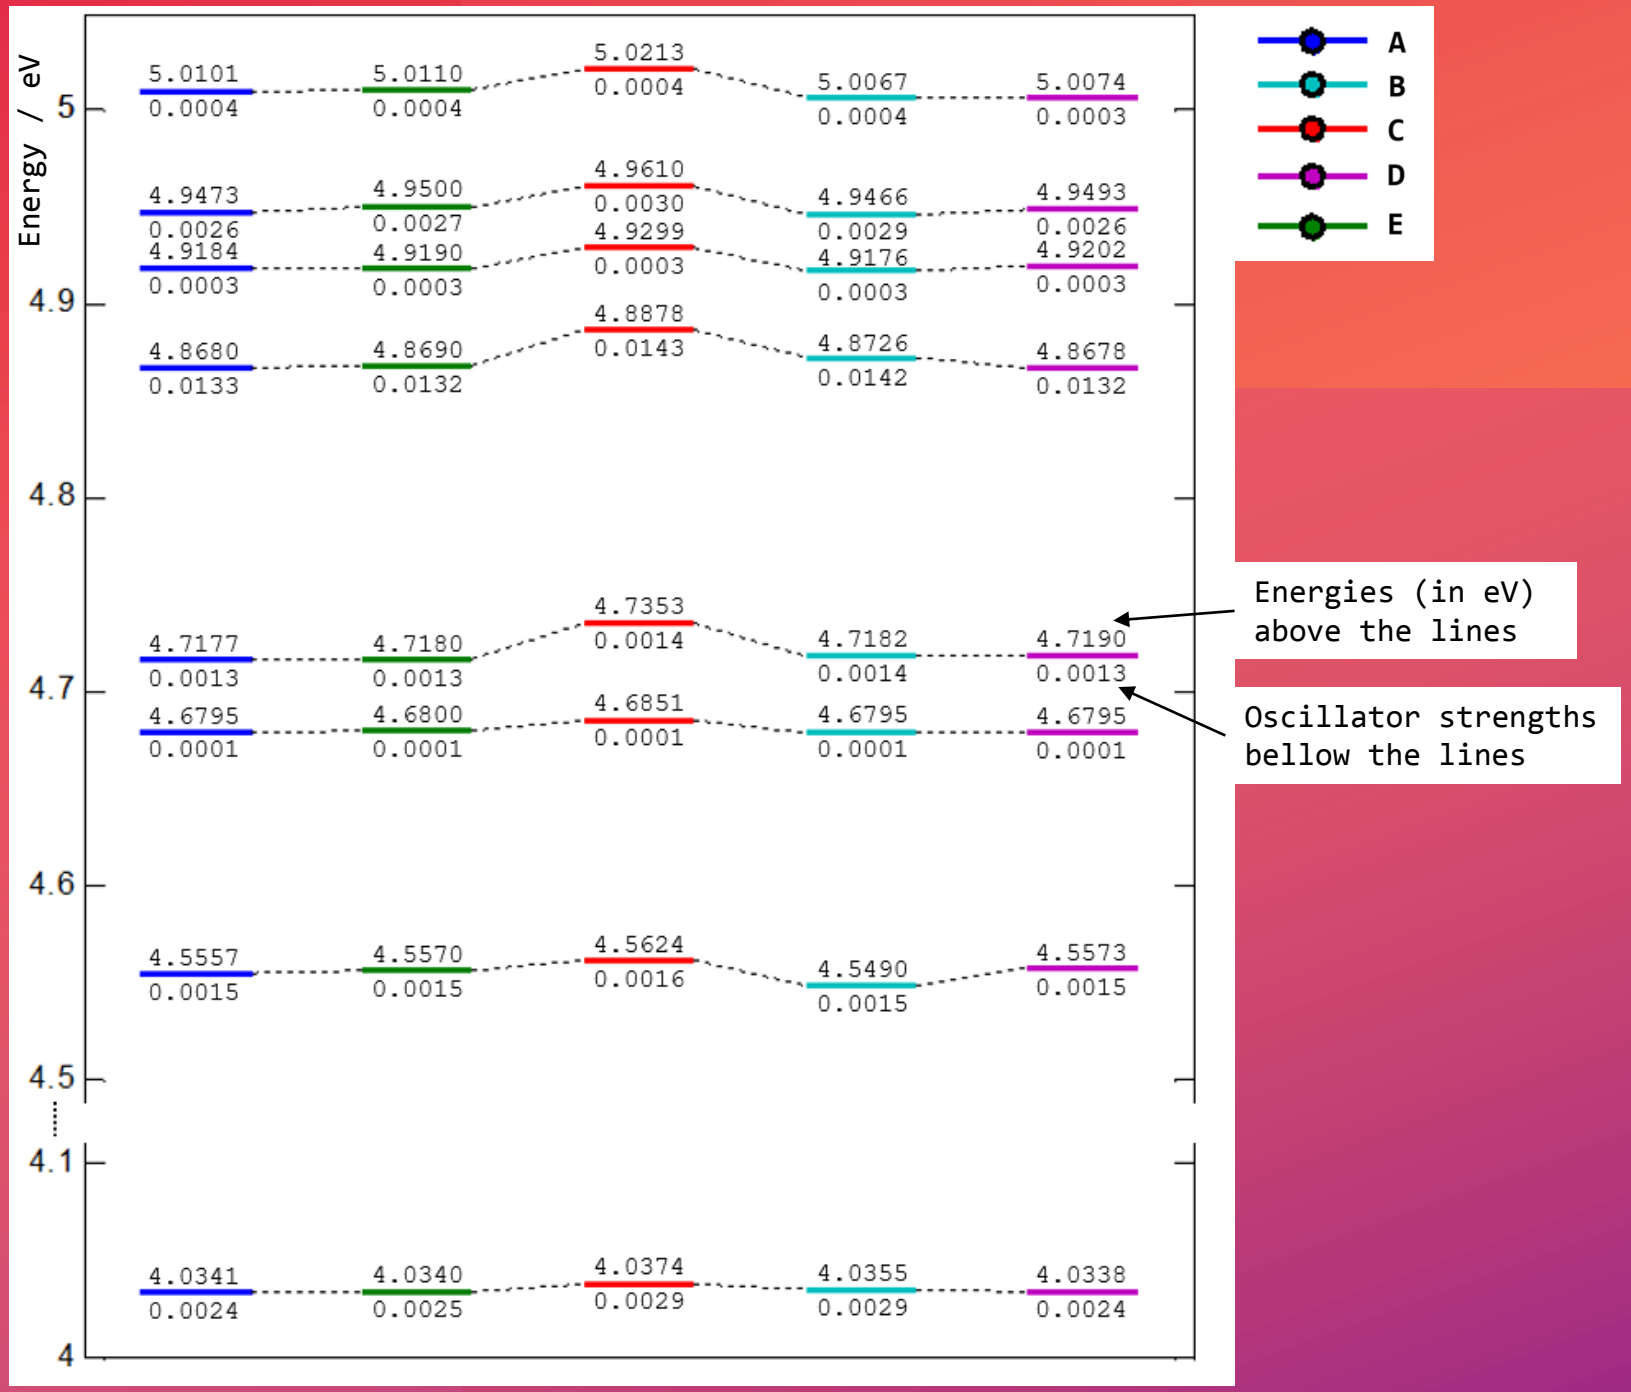

Energies (in eV) above the lines

Oscillator strengths bellow the lines

CONCLUSION

- The scalability drops off after about 16-32 cores for a system of common complexity
  - Further increase of processing cores is giving marginal performance boost
  - Use fewer processing cores and run several jobs instead
- While the absolute values of wall times differ, the “regularity” of performance drops is shared between all of the applications
- No overall winner and no overall loser
- Familiarity with certain software also plays a role when picking the one to work with
  - Sometimes learning a new application is worth the effort