

QMC USING VERY LARGE CONFIGURATION INTERACTION-TYPE EXPANSIONS

ARXIV:1510.00730V2 [PHYSICS.CHEM-PH]

A. Scemama¹, T. Applencourt, E. Giner, M. Caffarel

18/12/2015

¹Lab. Chimie et Physique Quantiques, IRSAMC, UPS/CNRS, Toulouse, France

QMC=Chem :

- 2013 : Oxygen atom, 100 000 determinants
- 2014 : F₂ dissociation curve, 10 000 determinants¹
- 2014 : Transition metals, 10 000 – 48 000 determinants²

Today

- Water : 84 000 determinants [-76.4354(7)]
- G2 set : ~ 200 000 determinants (w/wo pseudo-potentials)
- Cl atom : 750 000 determinants

¹arXiv:1408.3672 [physics.chem-ph]

²arXiv:1409.3671 [physics.chem-ph]

- *Exact* solution in a given basis set
- Basis : All possible Slater determinants (\mathcal{D}_k) of *spin-orbitals*

$$\Psi_{\text{FCI}} = \sum_k c_k \mathcal{D}_k \quad (1)$$

- Spin-Free formalism :

$$\mathcal{D}_k = D_k^\uparrow \times D_k^\downarrow \quad (2)$$

- Size of the FCI space :

$$\binom{N_{\text{MO}}}{N_{\text{elec}\uparrow}} \times \binom{N_{\text{MO}}}{N_{\text{elec}\downarrow}} \quad (3)$$

CI in cc-pVDZ

- 9 \uparrow and 8 \downarrow electrons in 19 MOs
- 92378 \uparrow determinants and 75582 \downarrow determinants
- $6.98 \cdot 10^9$ basis functions
- The size of the space is huge, but the space is **empty** :
 $\sim 10^6$ coefficients $|c_k| > 10^{-12}$

\implies CIPSI algorithm to select Slater determinants *on the fly*.³

³http://github.com/LCPQ/quantum_package

Ψ can be expressed in a bi-linear form:

$$\Psi(\mathbf{R}) = \sum_k c_k \mathcal{D}_k = \sum_i \sum_j C_{ij} D_{i\uparrow}(\mathbf{R}_\uparrow) D_{j\downarrow}(\mathbf{R}_\downarrow) \quad (4)$$

- $\mathbf{D}_\uparrow(\mathbf{R}_\uparrow)$: vector of $N_{\text{det}\uparrow}$ elements
- $\mathbf{D}_\downarrow(\mathbf{R}_\downarrow)$: vector of $N_{\text{det}\downarrow}$ elements
- \mathbf{C} : $N_{\text{det}\uparrow} \times N_{\text{det}\downarrow}$ matrix. The matrix contains N_{det} elements.

\mathbf{C} is constant in a QMC calculation \implies preprocessing.

1. Encoding of determinants⁴ :

HF : (000011111111,000111111111) = (255,511)

Single \uparrow : (010001111111,000111111111) = (1151,511)

Double $\uparrow\downarrow$: (010001111111,010101111111) = (1151,1407)

⁴arXiv:1311.6244 [physics.comp-ph]

DETERMINATION OF C

1. Encoding of determinants⁴ :

$$\text{HF} \quad : \quad (000011111111,000111111111) = (255,511)$$

$$\text{Single } \uparrow \quad : \quad (010001111111,000111111111) = (1151,511)$$

$$\text{Double } \uparrow\downarrow \quad : \quad (010001111111,010101111111) = (1151,1407)$$

2. Separate D_{\uparrow} and D_{\downarrow} of \mathcal{D}_k : two lists of N_{det}

$$000011111111 = 255 \quad \left| \quad 000111111111 = 511$$

$$010001111111 = 1151 \quad \left| \quad 000111111111 = 511$$

$$010001111111 = 1151 \quad \left| \quad 010101111111 = 1407$$

⁴arXiv:1311.6244 [physics.comp-ph]

DETERMINATION OF C

1. Encoding of determinants⁴ :

$$\text{HF} \quad : \quad (000011111111,000111111111) = (255,511)$$

$$\text{Single } \uparrow \quad : \quad (010001111111,000111111111) = (1151,511)$$

$$\text{Double } \uparrow\downarrow \quad : \quad (010001111111,010101111111) = (1151,1407)$$

2. Separate D_{\uparrow} and D_{\downarrow} of \mathcal{D}_k : two lists of N_{det}

$$000011111111 = 255 \quad | \quad 000111111111 = 511$$

$$010001111111 = 1151 \quad | \quad 000111111111 = 511$$

$$010001111111 = 1151 \quad | \quad 010101111111 = 1407$$

3. Sort both lists, remove consecutive repetitions : One list of $N_{\text{det}\uparrow}$,
one list of $N_{\text{det}\downarrow}$

$$000011111111 = 255 \quad | \quad 000111111111 = 511$$

$$010001111111 = 1151 \quad | \quad 010101111111 = 1407$$

⁴arXiv:1311.6244 [physics.comp-ph]

DETERMINATION OF C

1. Encoding of determinants⁴ :

$$\text{HF} \quad : \quad (000011111111, 000111111111) = (255, 511)$$

$$\text{Single } \uparrow \quad : \quad (010001111111, 000111111111) = (1151, 511)$$

$$\text{Double } \uparrow\downarrow \quad : \quad (010001111111, 010101111111) = (1151, 1407)$$

2. Separate D_{\uparrow} and D_{\downarrow} of \mathcal{D}_k : two lists of N_{det}

$$000011111111 = 255 \quad \left| \quad 000111111111 = 511$$

$$010001111111 = 1151 \quad \left| \quad 000111111111 = 511$$

$$010001111111 = 1151 \quad \left| \quad 010101111111 = 1407$$

3. Sort both lists, remove consecutive repetitions : One list of $N_{\text{det}\uparrow}$, one list of $N_{\text{det}\downarrow}$

$$000011111111 = 255 \quad \left| \quad 000111111111 = 511$$

$$010001111111 = 1151 \quad \left| \quad 010101111111 = 1407$$

4. For each \mathcal{D}_k , find the indices i and j of D_{\uparrow} and D_{\downarrow} (binary search $\mathcal{O}(\log N_{\text{det}\uparrow})$), and set $C_{ij} = c_k$

Takes 3 seconds on a single CPU core for $10^6 \mathcal{D}_k$, $\mathcal{O}(N_{\text{det}})$.

⁴arXiv:1311.6244 [physics.comp-ph]

REMARK ON ORDERING

D_{i-1} and D_i are likely to differ by a **single substitution** :

List index	Decimal	Binary	Determinant
1	15	00001111	1234⟩
2	23	00010111	1235⟩
3	27	00011011	1245⟩
4	29	00011101	1345⟩
5	30	00011110	2345⟩
6	39	00100111	1236⟩
7	43	00101011	1246⟩
8	45	00101101	1346⟩
...

In QMC=Chem : *We do all-electron moves.*

At every MC step, we need to evaluate:

$$\Psi = (D_{\uparrow}^{\dagger}(C)D_{\downarrow}) \quad (5)$$

$$\nabla_i \Psi = \nabla_i D_{\uparrow}^{\dagger} \cdot (CD_{\downarrow}) \text{ or } (D_{\uparrow}^{\dagger}C) \cdot \nabla_i D_{\downarrow} \quad (6)$$

$$\Delta_i \Psi = \Delta_i D_{\uparrow}^{\dagger} \cdot (CD_{\downarrow}) \text{ or } (D_{\uparrow}^{\dagger}C) \cdot \Delta_i D_{\downarrow} \quad (7)$$

$$V_{\text{pseudo}}^{\text{non-loc}} \Psi = V_{\text{pseudo}}^{\text{non-loc}} D_{\uparrow}^{\dagger} \cdot (CD_{\downarrow}) \text{ or } (D_{\uparrow}^{\dagger}C) \cdot V_{\text{pseudo}}^{\text{non-loc}} D_{\downarrow} \quad (8)$$

(\uparrow electrons and \downarrow electrons)

Vector-matrix product $\mathbf{D}_\uparrow^\dagger \cdot \mathbf{C}$: $N_{\text{det},\downarrow}$ operations, returns a $N_{\text{det},\downarrow}$ vector

1. Dot product with \mathbf{D}_\downarrow : $N_{\text{det},\downarrow}$ operations, produces a scalar
2. Matrix product with $\nabla \mathbf{D}_\downarrow$: $3N_{\text{elec},\downarrow} \times N_{\text{det},\downarrow}$ operations, produces a $3N_{\text{elec},\downarrow}$ vector
3. Matrix product with $\Delta \mathbf{D}_\downarrow$: $N_{\text{elec},\downarrow} \times N_{\text{det},\downarrow}$ operations, produces a $N_{\text{elec},\downarrow}$ vector
4. Matrix product with $V_{\text{pseudo}}^{\text{non-loc}} \mathbf{D}_\downarrow$: $N_{\text{elec},\downarrow} \times N_{\text{det},\downarrow}$ operations, produces a $N_{\text{elec},\downarrow}$ vector

If one localizes the pseudo-potential on the CI wave function :

$$\hat{T} + \hat{V} + \hat{V}_{\text{loc}}^{\text{pseudo}} + \sum_{lm} V_{\text{non-loc},l}^{\text{pseudo}} \frac{|Y_{lm}\rangle \langle Y_{lm} | \Psi_{\text{CI}}\rangle}{\Psi_{\text{CI}}} \quad (9)$$

$$V_{\text{non-loc}}^{\text{pseudo},i} = \sum_j S_{ij}^{-1} \sum_{lm} Y_{lm}(\mathbf{r}_i) \langle Y_{lm} | \phi_j \rangle \quad (10)$$

- The pseudo-potential integrals can be computed analytically
- $\langle Y_{lm} | \phi_j \rangle$ is pre-computed (once) on a grid
- No quadrature is needed : $\times 16$, $\times 27$ speed-up
- Tested on G2 set : Same MAD with/without pseudo-potentials
- Cost : Same as Laplacian
- Also applicable to single-determinant

DETERMINANTS AND THEIR DERIVATIVES

- First determinant : Standard technique as in single-determinant
- All other determinants : Sherman-Morrison updates

Determinant D_i is updated wrt D_{i-1} by successive column updates. Depends on the ordering.

Cl atom, cc-pVDZ, 10^6 Slater determinants:

- Average excitation degree wrt HF : 3.46
- Average number of column substitutions : 1.78

Excitation	$D_{\uparrow}D_{\downarrow}$	D_{\uparrow}	D_{\downarrow}
0	1	1	1
1	38	90	88
2	2 177	1603	1520
3	43 729	7811	6507
4	308 045	8581	5071
5	351 182	2090	579
6	291 481	77	0
7	3 067	0	0
8	280	0	0

SCALING

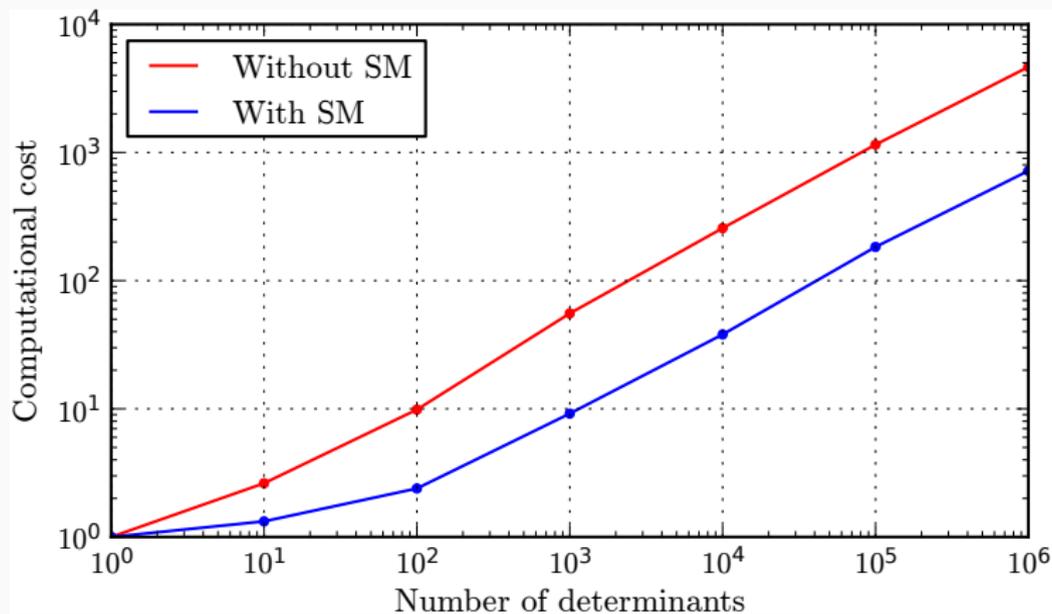
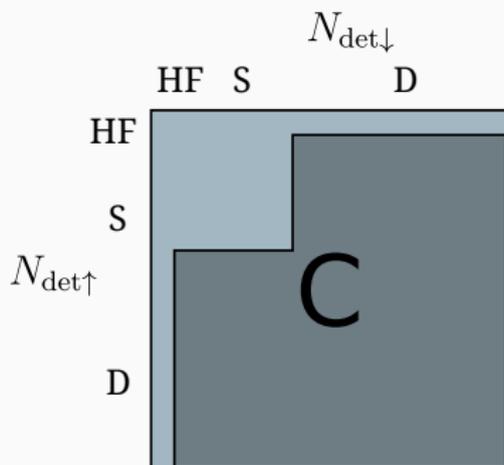


Figure 1: Cl : Scaling is $\sim \sqrt{N_{\text{det}}}$. SM updates $\times 7$.

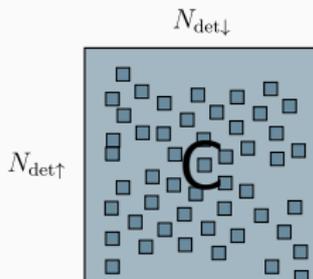
a priori truncation of the FCI space

- CAS : good for re-use of D_{\uparrow} and D_{\downarrow}
- CISD, CISDT, ...: Poor re-use



Usual truncation

- Based on the $|c_k|$ (or CSF coefficients)
- Removes mostly same-spin doubles in CISD \implies some gain is observed.
- Makes the matrix \mathbf{C} more sparse
- For CASSCF or FCI : Needs a severe truncation to remove D_{\uparrow} or D_{\downarrow} : the whole column must be zero

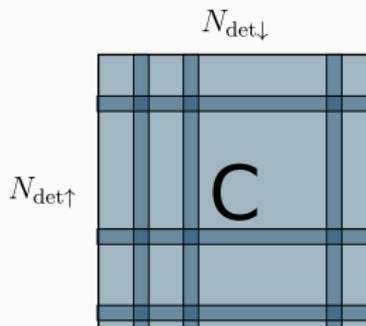


Alternate truncation scheme

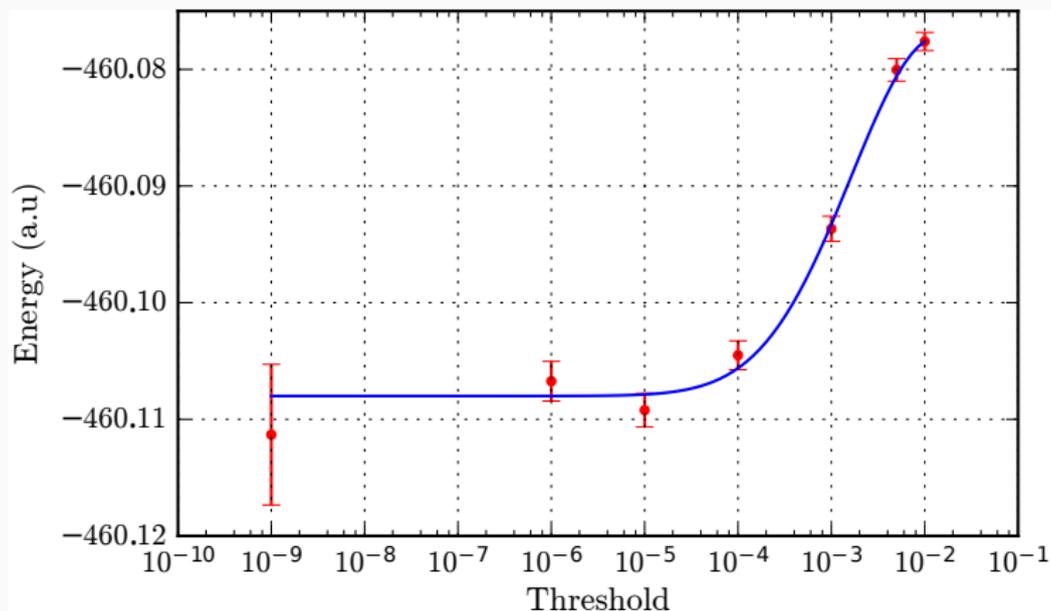
- Contribution of D_{\uparrow} and D_{\downarrow} to the norm of Ψ

$$\mathcal{N} = \sum_{i=1}^{N_{\text{det}\uparrow}} \sum_{j=1}^{N_{\text{det}\downarrow}} C_{ij}^2 = \sum_{i=1}^{N_{\text{det}\uparrow}} \mathcal{N}_i^{\uparrow} = \sum_{j=1}^{N_{\text{det}\downarrow}} \mathcal{N}_j^{\downarrow} \quad (11)$$

- Remove $D_{i\uparrow}$ when $\mathcal{N}_i^{\uparrow} = \sum_{j=1}^{N_{\text{det}\downarrow}} C_{ij}^2 < \epsilon$

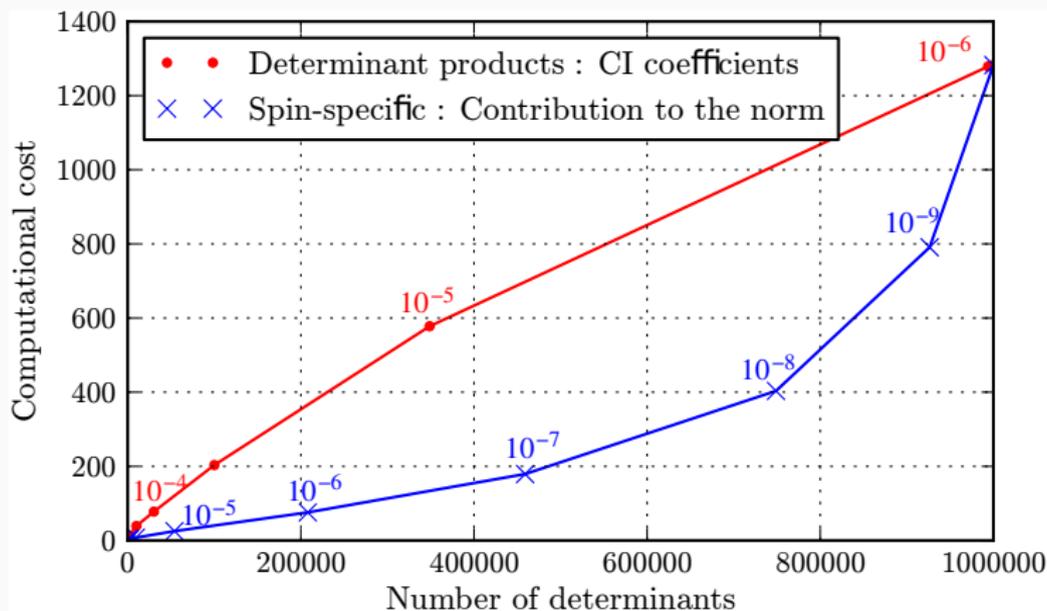


CONVERGENCE OF THE FN-DMC ENERGY



Cl FCI/cc-pVTZ

COMPUTATIONAL COST



Converged FCI/cc-pVTZ FN-DMC energy with $\epsilon = 10^{-5}$: 7.1 \times the cost of a single-determinant

TOTAL CPU TIME PER MC STEP

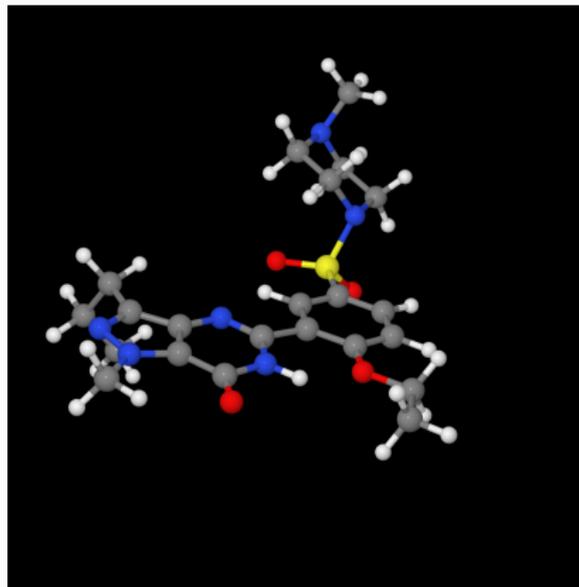
ϵ	N_{det}	$N_{\text{det}}^{\uparrow}$	$N_{\text{det}}^{\downarrow}$	CPU time(ms)
10^{-2}	1	1	1	0.02450
5.10^{-3}	3	3	2	0.02688
10^{-3}	86	21	17	0.03899
5.10^{-4}	214	29	28	0.04636
10^{-4}	1361	93	74	0.08808
5.10^{-5}	2424	120	89	0.1054
10^{-5}	9485	234	166	0.1855
10^{-6}	54016	772	523	0.5960
10^{-7}	207995	2279	1389	1.740
10^{-8}	459069	5797	3291	4.196
10^{-9}	748835	14456	8054	9.724
10^{-10}	926299	30320	16571	19.18
0	1000000	52433	26833	31.92

Order of magnitude :

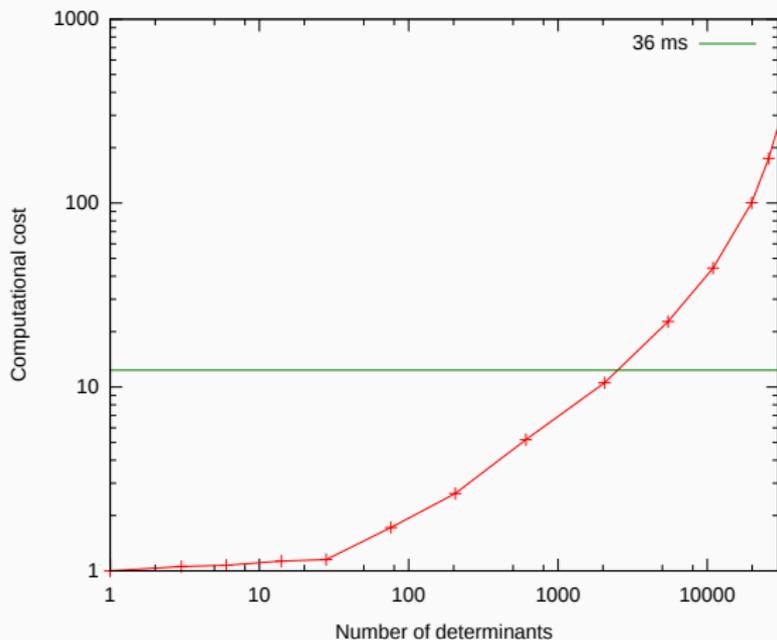
36 ms/step $\equiv 10^{11}$ steps in 10 hours on 100 000 cores

Sildenafil

- 63 atoms
- 126 \uparrow and 126 \downarrow electrons
- 649 basis functions
- MP2 wave function



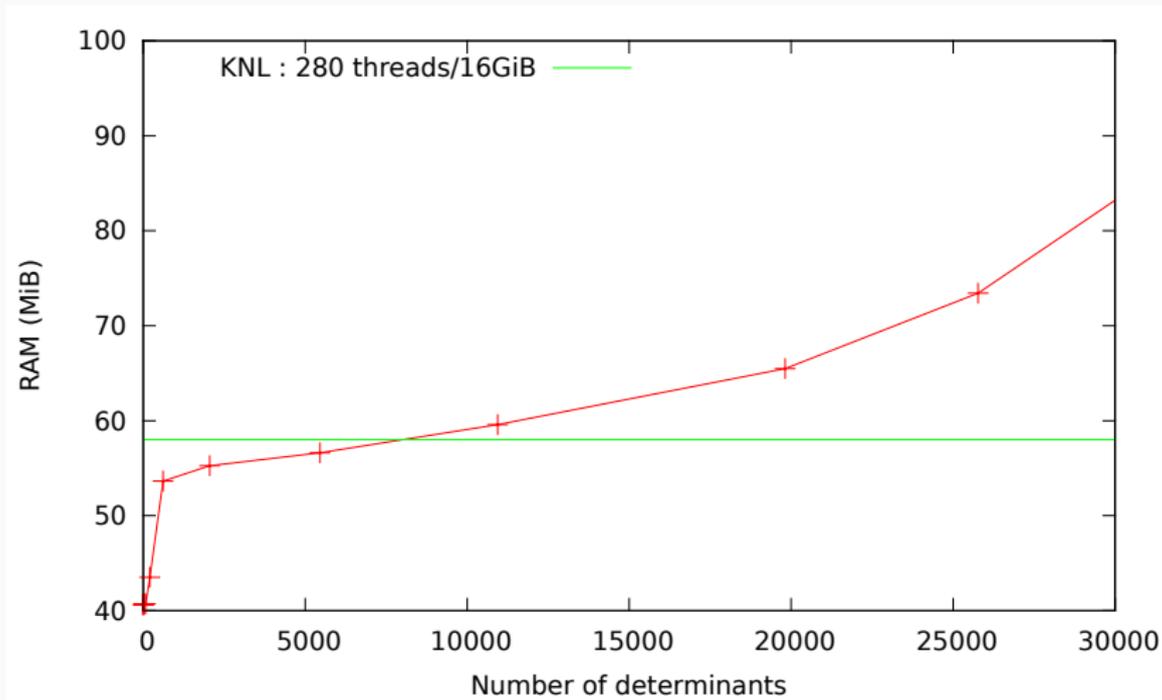
TOTAL TIME



Sildenafil : 31 ms/step $\Leftrightarrow \epsilon = 2 \cdot 10^{-5}$ ($N_{\text{det}} = 2049$)

100 dets only 2× more expensive than 1 det.

MEMORY PER CPU CORE



SUMMARY

- Bi-linear expression : scaling as $\sqrt{N_{\text{det}}}$
- The only N_{det} -scaling step is a matrix-vector product : tiny pre-factor, independent of N_{elec} , negligible
- Sherman-Morrison : $\sim \times N_{\text{elec}\uparrow}$
- Ordering $\Rightarrow D_i / D_{i-1}$: less work to do than D_i / D_0 : $\sim \times 2$
- Localizing the pseudo-potential on the determinantal component : $\times 16 - \times 27$
- Memory : scales as $N_{\text{elec}\uparrow} \times N_{\text{det}\uparrow}$
- 10^6 determinants are possible for small systems (pseudo and all-electron)
- ~ 2000 determinants are possible for a 250-electron system using a (difficult) CISD basis
- More details here : [arXiv:1510.00730v2](https://arxiv.org/abs/1510.00730v2) [*physics.chem-ph*]

COMPARISON WITH TABLE METHOD

Table method : store a 2D array where each element is the result of a dot product

- Dot product : $N_{\text{elec}} \uparrow$ FMA operations
- Modern CPUs can do 16 FMAs ($\mathbf{a} = \mathbf{a} + \mathbf{b} * \mathbf{c}$) per CPU cycle
- Replace FMA by a random memory access : L1 (4 cycles), L2 (12 cycles), L3 (44 cycles), RAM (357 cycles)
- And then compute a small determinant
- L1 \Leftrightarrow 64 \uparrow electrons
- L2 \Leftrightarrow 192 \uparrow electrons
- L3 \Leftrightarrow 704 \uparrow electrons
- RAM \Leftrightarrow 5712 \uparrow electrons

COMPARISON WITH TABLE METHOD

Measurements on CI, 1M determinants:

	D_{\uparrow}	D_{\downarrow}
Table method	2110 cycles	2431 cycles
This work	544 cycles	359 cycles

But:

The table method allows to avoid the storage of all Slater matrices when doing *single-electron moves* : very good!

We prefer all-electron moves because they use more efficiently the hardware (same as BLAS2 vs BLAS3).