IQA analysis of the two-particle density matrix: chemical insight and computational efficiency

Mark A. Vincent[1] · Paul L. A. Popelier[1]

Abstract

The interacting quantum atoms (IQA) method offers a rigorous and minimal route to calculate atomic electron correlation energies from the two-particle density matrix (2PDM). The price paid is that this method is very time-consuming. However, employing CCSD and CCSD(T), we explore several approaches to speed up such calculations. We make the pivotal observation that the removal, from the true 2PDM, of both the Hartree–Fock part of the 2PDM and an approximate 2PDM (Müller) dramatically reduces the size of the quadrature grid needed to obtain accurate energies.

1 Introduction

Understanding how a molecular system works energetically is a massively documented and important research activity, which caused the development and application of a good number of energy decomposition schemes over the last few decades. These schemes typically strive for a two-dimensional analysis: one dimension expresses the type of energy (e.g. electrostatic, exchange, dispersion, induction, correlation, kinetic, …) while the other expresses the locale in the (molecular) system. To obtain full resolution, the latter dimension should cover the atomic level. Some energy partitioning methods achieve this naturally, indeed by definition. A prime example is the interacting quantum atoms (IQA) [1], which is based on the calculation [2] of the potential energy of (quantum) topological atoms, independently of the virial theorem [3], which dominated and restricted the energy partitioning of the quantum theory of atoms in molecules (QTAIM) [4].

In contrast, energy decomposition schemes other and older than IQA, such as symmetry-adapted perturbation theory (SAPT), lack atomic resolution in their original design. In other words, SAPT faces the conceptual challenge of providing atomic information because the method has the idea of a molecule at its heart rather than that of an atom. Still, some time ago, an atomically decomposed version called A-SAPT [5] was formulated. However, A-SAPT immediately experienced difficulties producing chemically useful partitions of the electrostatic energy, due to the build-up of oscillating partial charges on adjacent functional groups. This defect triggered the introduction of F-SAPT [6], the functional-group SAPT partitioning. A critical review [7] by Skylaris and co-workers published in 2015 gives more examples of pitfalls and problems of several (more mature and often popular) energy decomposition schemes. Both the original versions and their variants suffer(ed) from problems while IQA, reviewed [8] by its originators 5 years later, is not discussed in the 2015 review. Finally, an unusually frank review [9] on the Hirshfeld family of partitioning schemes, published in 2018, states that it believes that "*every popular Hirshfeld-based partitioning method has at least one serious flaw*". In summary, IQA offers a robust and minimal framework on which to base both the understanding and prediction of the energetics of molecular systems.

Here, we are interested in a quantum topological representation of dispersion energy, but first we should mention that a "half-way house" is possible. One can apply [10] IQA with a dispersion "bolt-on" using one of the Dn ($n = 1,2,3,..$) schemes [11], which can then be denoted as IQA-D3, for example. Although fast and practical, we prefer to follow a route that is fully compatible with the IQA

✉ Paul L. A. Popelier
  pla@manchester.ac.uk

[1] Department of Chemistry, The University of Manchester, Manchester M13 9PL, Great Britain

framework itself. Indeed, IQA offers a well-defined energy contribution that it associates with pure electron correlation, alongside an energy term connected to electrostatic energy and another one to pure exchange energy. Very recently, the strategy behind this ultimately preferable route to dispersion energy has been laid out [12] with reference to the construction of a novel polarisable, multipolar, machine-learnt force field called FFLUX [13]. The calculation of IQA correlation energies is computationally extremely expensive because it handles the two-particle density matrix (2PDM). Unfortunately, the latter is a humongous object that contains all the precise correlation effects that the post-Hartree–Fock method at hand delivers. Atomic correlation energies were already calculated [1] at full CI level for H_2 and He_2 in 2005, but it took until 2016 before they were computed [14] for the first time, at MP2 level (as well as at MP3 and MP4SDQ level).

All our earlier work was confined to MPn wavefunctions, which provide interesting chemical insight such as transferability, through-space effects, covalent bond characterisation, hydrogen bonding, electron delocalisation, H…H dispersive cohesion and protobranching. These results were recently reported in a mini-review [15] of our own work carried on a variety of systems including water clusters, small inorganic molecules (hydrides), hydrocarbons and the He…H_2 complex. However, there is a conceptual issue with the MPn approach due to the fact that MPn correlation only affects electron–electron terms. Thus, the one-electron energy terms are not affected and hence correspond to those at Hartree–Fock level. This also means that the shapes (and volumes) of topological atoms are determined by the Hartree–Fock electron density. In contrast, this mismatch does not exist for coupled-cluster wavefunctions because its electron density now includes correlation effects. Hence, CCSD(T) wavefunctions are the way to go to obtain the definitive insight in how correlation energy distributes itself in a (molecular) system.

This is why in the current work we focus on CCSD(T) and CCSD wavefunctions only and systematically explore a variety of ideas, whether successful or not, to speed up the calculation of atomic IQA correlation energies. Commenting on our transition from MPn to CCSD(T) brings up three points. The first point is that recent collaborative work [16] of two external groups also considered MP2 wavefunctions and showed that efficiencies can be obtained due to the fact that 2PDM elements involve active occupied orbitals with the active virtual orbitals only. This situation is unlike that for CCSD or CCSD(T) wavefunctions where the 2PDM elements involve all the active orbitals. The second point is that our recent comparison [17] of IQA values obtained from the MP4SDQ and CCSD approaches (both with the HF component removed) reveals them to surprisingly different. However, if the HF component had not been removed, the values would have been much more similar. We concluded from our work on energetic transferability in water clusters that both CCSD and MP4SDQ uncover a remarkable additivity in the intra-atomic correlation energy of an oxygen, which drops by 25 kJ mol^{-1} for accepting a hydrogen when forming a hydrogen bond. However, only CCSD detected the (negative) increment of 15 kJ mol^{-1} for donating a hydrogen. The third point is that we have compared our MP2 IQA correlation energies with those obtained by the Müller approximation [18] (abbreviated as *M* throughout the article). We found a reasonable correlation between the two sets of values, although the statistical error varied from element to element, with, hydrogen having the smallest error, not unexpectedly. A similar comparison between CCSD(T) and *M* is currently in progress.

In summary, in this paper we explore a number of ways to speed up the calculation of atomic electron correlation energies using IQA and a CCSD(T)/CCSD 2PDM. We ask the central question if any of these ideas will ever make this preferred way of incorporating dispersion into the machine-learnt force field FFLUX feasible.

2 Theory and computational details

2.1 Background

We have given the theory behind our approach elsewhere [14] but briefly summarise it here. The energy V_{ee}^{AB} is the electron–electron (e-e) energy of a single atom ($A = B$) in a (molecular) system or the e-e energy between two atoms ($A \neq B$) and is given by Eq. 1,

$$V_{\text{ee}}^{\text{AB}} = \sum_{j=1}^{N_{\text{basis}}} \sum_{k=1}^{j} K_{jk} \sum_{l=1}^{N_{\text{basis}}} \sum_{m=1}^{l} K_{\text{lm}} d_{\text{jklm}} \int_{\Omega_A} d\mathbf{r}_1 G_{\text{jk}}(\mathbf{r}_1 - \mathbf{R}_{\text{jk}}) \int_{\Omega_B} d\mathbf{r}_2 \frac{1}{r_{12}} G_{\text{lm}}(\mathbf{r}_2 - \mathbf{R}_{\text{lm}}) \tag{1}$$

where N_{basis} is the number of primitive Gaussian basis functions, G_{jk} and G_{lm} each are the product of two Gaussian basis functions, Ω is the volume of an atom, and the 2PDM is designated by d_{Jklm}. Note that $V_{\text{ee}}^{\text{AB}}$ can become the electron correlation energy $V_{\text{ee,corr}}^{\text{AB}}$ if d_{Jklm} is restricted to the pure electron correlation part of the 2PDM, which can be achieved by subtracting from it the HF Coulomb and exchange part of the 2PDM. The subscripts j, k, l and m denote a basis function, while K (of course not to be confused with the index k) is a constant resulting from the product of Gaussian basis functions.

Note that in Eq. 1 there is no reference to which wavefunction the 2PDM came from. In other words, the equation is universally valid for any 2PDM. However, for convenience we introduce a notation that specifies the source of the 2PDM at hand. Most of our work so far drew $d_{\text{Jklm}}^{\text{corr}}$ from MP2 wave functions, and thus, the concomitant 2PDM will be denoted 2PDM/MP2. In the current work, we will work with 2PDM/CCSD and 2PDM/CCSD(T).

Of great importance in the current article is the Müller approximation [18], which is abbreviated with the letter *M*. This letter will mainly be used as a shorthand to mark a modification (by subtraction) of the 2PDM as discussed later in Sect. 2.4. However, "*M*" can also be used in 2PDM/M[CCSD] or 2PDM/M[CCSD(T)], for example, in the context of program checking (debugging). In other words, we rarely use this matrix 2PDM/M in its own right. Furthermore, in the limit, we also consider the role of the Hartree–Fock component of a 2PDM, which we thus denote 2PDM/HF. In the Molecular Orbital basis, this 2PDM/HF is extremely sparse, consisting only of zeroes, ones and twos.

2.2 Software

We employed the ab initio program PySCF to generate the relevant wavefunctions, matrices and so-called electrostatic potential integrals (ESP) integrals. Our own in-house program MORFI was used to carry out the IQA analysis of the 2PDM. The grids employed by MORFI are specified by two numbers: (i) a designation number (not to be confused with the number of grid points) for the angular Lebedev grid and (ii) the number of radial points. The grid designation number and the corresponding grid size can be obtained from Table S1 in the Supporting Information (SI). AIMAll was employed for other IQA analyses. Excel was used to generate the graphs and the program MOLDEN [19] for the figures.

2.3 Atomic integration by the ESP approach (hybrid analytical and three-dimensional quadrature)

It is possible to sum the $V_{\text{ee,Corr}}^{\text{AB}}$ over *B*, which gives the total energy of *A* with itself and all the other atoms. This sum is actually equivalent to a set of the ESP integrals, which can be evaluated analytically. Thus, two numerical integrations are replaced by one analytical integration and one numerical integration, leading [20] to a substantial reduction in numerical error. In addition, this algorithm drastically reduces CPU time. We refer to this approach by the shorthand *A-A'* and the double numerical integration approach as the *A-B* method.

2.4 2PDM Modification and error specification

Another approach we employ is to modify the 2PDM by subtracting another (approximate) 2PDM from it, such as the Hartree–Fock one (2PDM/HF). As will be demonstrated in Results and discussion section, the reason for executing this subtraction is that the numerical integration of the modified 2PDM needs fewer grid points to be accurately integrated than the original, full, 2PDM does. It is useful to focus on the notations, especially in connection with 2PDM modifications that we will use throughout this article. For example, after the Hartree–Fock component has been eliminated from the full 2PDM at CCSD level of theory, we refer to the resulting 2PDM as 2PDM-HF/CCSD. A second example would be 2PDM-M/CCSD(T), which refers to the subtraction of the Müller 2PDM from the original 2PDM at CCSD(T) level of theory.

The measure of accuracy of integration is termed the recovery error. It is obtained as the difference (typically in kJ mol^{-1}) between the true (original) energy and the energy "recovered" by numerical integration of the 2PDM (whether full or modified by subtraction). The recovery error can be determined accurately for the full 2PDM, or for the 2PDM-HF, because their true (original) two electron energies are easily obtainable from the total energy, one-particle density matrix (1PDM) and the nuclear repulsion energy. The error for the 2PDM-M approach is not determined in such a straightforward way as that for the 2PDM-HF approach, where we know it exactly. There are two ways of obtaining it: (i) use a large grid and determine the IQA terms from the 2PDM/M directly and then use these as reference energies, or (ii) employ AIMAll, which efficiently calculates the relevant energies employing the Müller approximation. This assumes we are using AIMAll with a large grid. We mainly employ the AIMAll approach.

3 Results and discussion

3.1 Substantial computational saving by subtracting Müller's approximation from the full 2PDM

Our early work [14, 15, 20–25] on the IQA analysis of the 2PDM involved MP2 wavefunctions, calculated by GAUSSIAN09 (abbreviated as G09). This ab initio program produced a 2PDM-HF/MP2 that contained *correlation only*, that is, without the Hartree–Fock component. In other words, subtracting the Hartree–Fock 2PDM (i.e. 2PDM/HF) from the 2PDM/MP2 is like filtering out, from the original 2PDM, the pure correlation part. Now, we made the pivotal observation that this correlation-only 2PDM-HF/MP2 can be accurately integrated with much smaller grids than those necessary for the (full) 2PDM/MP2. Here, "full" means having both HF and correlated parts of the 2PDM in one matrix, that is, 2PDM/MP2. Put differently, *the inclusion and thus presence of the Hartree–Fock component causes the integration to waste grid points.* Indeed, it is more efficient to integrate the Hartree–Fock component separately, with the large grid that it needs. The computational advantage is then based on the fact that the Hartree–Fock component is actually a two-dimensional object (2D) (i.e. HF is a one-electron theory) rather than the four-dimensional (4D) object that is 2PDM/MP2. Indeed, 2PDM-HF/MP2, being 4D in the number of basis functions, is a huge matrix but has the advantage of corresponding to (much) smaller energies such that a very small integration grid manages to obtain these energies accurately.

The question is now if this basic idea of "filtering" the 2PDM can be repeated such that more computational savings can be made. In other words, is there another low-dimensional component (most likely again 2D instead of 4D of course) that can be taken out from the full 4D 2PDM such that the latter can be integrated even faster, by use of smaller grids? Indeed, this component is called 2PDM/M, after Müller, who proposed it almost four decades ago as an approximate 2PDM that can be written as the product of a 1PDM with itself. After this subtraction, we are left with a 2PDM that contains essentially pure two-electron correlation, that is, the part that cannot be obtained by any 2D object such as the Müller approximation, nor by the Hartree–Fock component of course.

We have implemented the removal of the 2PDM/M from the full 2PDM, which we designate as 2PDM-M, in order to investigate if this action reduces the integration grid size even more than for 2PDM-HF. Before the success of this approach can be demonstrated, a number of tests needed to be carried out. This initial work aimed at proving the correctness of the implementation of the Müller 2PDM. This work compares the results, for H_2 and H_2O, originating from the program AIMAll with those generated by the in-house program MORFI. For this purpose, we use the CCSD method and a simple 4s basis set for H_2 and the uncontracted STO-3G basis set for water.

The vast majority of data and details are given in Sect. 1 of the SI. Preliminary results with these simple basis sets are shown in the SI (Tables S1 and S2 for *A-B* energies and Table S3 for *A-A'* energies). Despite different grids and β spheres (and indeed different integration algorithms), the calculated energies are sufficiently similar (errors of the order of ~ 0.01 to ~ 0.1 kJ mol^{-1}) to lead us to conclude that our results are correct. Secondly, Table S1 gives an idea of how much correlation energy is missed by the Müller approximation (2PDM/M[CCSD]) compared to the full 2PDM. For example, the total intra-atomic electron–electron (e-e) energy (thus containing Coulomb, exchange and correlation) of a hydrogen atom in H_2 is 402.2 kJ mol^{-1} according to 2PDM/CCSD but only 396.3 kJ mol^{-1} by 2PDM/M[CCSD]. In other words, the Müller approximation underestimates the intra-atomic *e-e* energy of one H by 5.9 kJ mol^{-1}. For the interatomic *e-e* energy in H_2, the Müller estimate is even better. Here, the latter overestimates the exact energy of 760.7 kJ mol^{-1} by only 1 kJ mol^{-1}.

The next preliminary tests are carried out in *A-A'* mode rather than in the previously used *A-B* mode. Table S3 shows A-A' energies for H_2 and H_2O obtained with various approach and integration grids. An enormous reference grid of more than half a million grid points allowed us to show the accuracy of a much smaller grid (about 2500 points, or more than 200 times smaller). Indeed, oxygen's *A-A'* energy differs by only 0.03 kJ mol^{-1}. Table S4 repeats this success now with the more realistic (uncontracted) basis set of aug-cc-pVDZ for a single water molecule: the 10–10 grid (1,700 points for one ß sphere only) generates an error for oxygen of only 0.03 kJ mol^{-1} compared to that of the largest grid (43,620 points for one ß sphere only). Table S5 reports similar success on the water dimer at the same level of theory: the energy errors on oxygen are very small (< 0.1 kJ mol^{-1}) even for the smallest grid of 10–15 (2,550 points for one ß sphere only).

To build on this successful traction, we now turn our attention to the water trimer whose results are given in Table 1 and whose labelling scheme in given in Fig. 1. This trimer is an important non-covalent system challenging FFLUX, the development of which motivates the current work, as explained in Introduction. This realistic non-covalent system exhibits the cooperative effect, which needs to be targeted by the machine learning behind FFLUX, *en route* to tackling bulk water. Here, we consider how the energy error with respect to the reference energy (largest grid) varies with grid size for the 2PDM-M/CCSD approach

Table 1 The *A-A'* energies (in Hartrees) for a water trimer, obtained with 177 basis functions of uncontracted aug-cc-pVDZ leading to 2PDM-M/CCSD

Grid	32–30	15–20	15–15	10–15	10–10
O1	0.052183	0.052186	0.052178	0.052165	0.052145
O2	0.052266	0.052276	0.052267	0.052259	0.052239
O3	0.052129	0.052128	0.052119	0.052128	0.052109
H4	0.006103	0.006110	0.006110	0.006090	0.006090
H5	0.007382	0.007382	0.007382	0.007383	0.007385
H6	0.006105	0.006095	0.006095	0.006105	0.006105
H7	0.007405	0.007405	0.007405	0.007405	0.007408
H8	0.006138	0.006146	0.006146	0.006108	0.006108
H9	0.007390	0.007390	0.007390	0.007390	0.007392
Error O (kJ mol^{-1})	0.00 (0.00)	0.02 (0.02)	0.03 (0.02)	0.05 (0.08)	0.10 (0.08)
no. of grid points O	43,260	7,000	5,250	2,550	1,700
% of grid points (O only)	100%	16%	12%	6%	4%
CPU (total run)	~73 h	~12 h	~8 h	~4 h	~3 h

Grid	10–12 & 10–8	10–11 & 10–7	9–11 & 9–7	8–11 & 8–7	7–11 & 7–7
O1	0.052192	0.052169	0.052122	0.052131	0.052064
O2	0.052286	0.052262	0.052204	0.052213	0.052289
O3	0.052156	0.052132	0.052164	0.052129	0.052064
H4	0.006089	0.006089	0.006080	0.006086	0.006109
H5	0.007371	0.007392	0.007391	0.007389	0.007386
H6	0.006104	0.006105	0.006094	0.006083	0.006105
H7	0.007394	0.007416	0.007416	0.007414	0.007414
H8	0.006107	0.006107	0.006130	0.006193	0.006184
H9	0.007378	0.007399	0.007399	0.007398	0.007406
Error O (kJ mol^{-1})	0.07 (0.08)	0.04 (0.08)	0.17 (0.06)	0.14 (0.14)	0.31 (0.12)
no. of grid points O&H	2040 & 1360 Tot = 3400	1870 & 1190 Tot = 3060	1606 & 1022 Tot = 2628	1210 & 770 Tot = 1980	946 & 602 Tot = 1548
% grid points (O only)	5%	4%	4%	3%	2%
CPU	~4 h	~3 h	~3 h	~2 h	~2 h

Grid	6–11 & 6–7	6–12 & 6–7	5–10 & 5–7	4–10 & 4–7	3–10 & 3–7
O1	0.052203	0.052226	0.051979	0.052673	0.052155
O2	0.052464	0.052489	0.052265	0.052155	0.051845
O3	0.052125	0.052148	0.052266	0.051868	0.052034
H4	0.006185	0.006185	0.006109	0.006151	0.006115
H5	0.007381	0.007381	0.007400	0.007362	0.007347
H6	0.006088	0.006088	0.006094	0.006079	0.005982
H7	0.007419	0.007419	0.007390	0.007496	0.007332
H8	0.006100	0.006100	0.006107	0.006200	0.005973
H9	0.007402	0.007402	0.007375	0.007399	0.007249
Error O (kJ mol^{-1})	0.52 (0.22)	0.58 (0.22)	0.54 (0.08)	1.29 (0.24)	1.11 (0.43)
no. grid points O&H	814 & 518 Tot = 1332	888 & 518 Tot = 1406	500 & 350 Tot = 850	380 & 266 Tot = 646	260 & 182 Tot = 442
% grid points (O only)	2%	2%	1%	1%	~1%
CPU	~2 h	~2 h	~1 h	~1 h	~1 h

The lower part of each of the 3 sub-tables gives the largest error, with respect to the reference, of any of the oxygen atoms and in brackets that of any of the hydrogens. In addition, a variable grid (marked by "&") is employed in the second and third sub-tables, where the larger radial grid is that of O and the smaller that of H. Each grid is labelled by two numbers (*m–n*), where *m* is a label for the angular grid (sizes in Table S1) and *n* the number of radial points. To obtain the total number of grid points, the number of grid points needs to be doubled because the *β* sphere is not included here
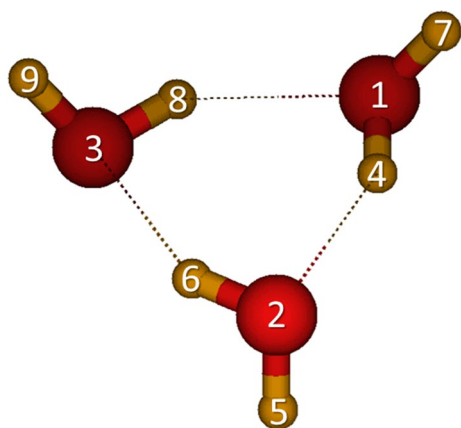
Fig. 1 The geometry and atomic labelling for the water trimer

employing the more contemporary uncontracted aug-cc-pVDZ basis set.

Table 1 shows the obtained A-A' energies of each atom, for a large number of grids presented in the order of decreasing number of grid points. In addition, the percentage of the number of grid points with respect to the reference is given along with the CPU time consumed for each of the runs. The largest grid (column labelled 32–30 in the upper sub-table of Table 1) serves as this reference against which we determine how well the smaller grids are performing in terms of energy accuracy. In total, 15 grids are listed by decreasing total grid size, starting at the top left and going down, from left to right. Each grid that is smaller (and indeed much smaller) than the reference grid manages to still produce the "exact" reference energy to within 0.6 kJ mol^{-1} for oxygen, with the exception of the two smallest grids where errors of just over 1 kJ mol^{-1} are reached. Secondly, the first four grids (upper sub-table) show a monotonic increase in error (up to 0.1 kJ mol^{-1}) as the grid size (for oxygen) decreases to 4% of that of the reference grid. The corresponding CPU time reduction is more than a factor of 20. We also observe that while the error mostly increases smoothly with grid size reduction, sudden bumps may occur. For example, when moving from the grid labelled **10–11** and **10–7** (middle sub-table, Table 1) to that labelled **9–11** and **9–7** (middle sub-table, Table 1) oxygen's error more than quadruples. In addition, the CPU time reduction from the larger grid to the smaller grids generally follows the grid size. This effect can be quantified by comparing the ratios of the grid sizes with that of the CPU times. For example, the ratio of grid sizes 2,550 to 43,260 (for grids 10–15 and 32–30, respectively, Table 1) is 0.059 (i.e. 6%), while the corresponding ratio of CPU times is 4/73 = 0.055. These ratios are indeed remarkably similar, with other grids seeming to also adhere to the corresponding quantities.

However, a word of caution is necessary as our hardware operates with heterogeneous nodes.

Finally, Table 1 reports on the effect of allowing the radial grid of hydrogen and oxygen to be different. Such mixed grids enable further CPU time savings. For example, it is almost spectacular that a grid as small as 500 and 350 points (column labelled **5–10** and **5–7** of the lower sub-table of Table 1) for O and H, respectively, yields a correlation energy for O that is in error by only ~ 0.5 kJ mol^{-1} while being obtained almost two orders of magnitude faster compared to the reference grid. Trimming the grids further suddenly bumps up the error to slightly above the psychological barrier of 1 kJ mol^{-1}. Although this error is still four times smaller than the oft-quoted chemical accuracy threshold of 1 kcal mol^{-1}, the tiniest grids cause alarm bells to ring with the 1 kJ mol^{-1} barrier in mind.

The SI shows six more tables (Tables S6 to S11 in Sect. 1 of the SI), which essentially reinforce the findings of Table 1. Table S6 shows that very small grids, 2% to 4% of the size of the same reference grid (32–30), generate errors of the order of 0.1 kJ mol^{-1} for a stretched cyclic and a linear water trimer configuration. With these distorted trimer geometries, we are attempting to show the effects of geometric distortions that one might expect from a dynamics calculation. Tables S7 and S8 confirm similar performance of small grids on HF and the halogen-bonded complex HF…F$_2$, respectively. Table S9 reports on LiH, where a sudden jump in error, from ~ 0.1 kJ mol^{-1} to 0.5 kJ mol^{-1}, occurs for grids smaller than 6–10. Remarkably, there is no grid size dependence for the Li ion's energies. Table S10 again shows the effectiveness of small grids in methane, with a sudden deterioration for C from the 7–10 grid size onward, but still contained within 0.16 kJ mol^{-1}. The hydrogen A-A' energies of CH$_4$ start showing small deviations from practical point group symmetry, from only 0.01 kJ mol^{-1} for larger grids to ~ 0.1 kJ mol^{-1} for grid 6–10. Finally, Table S11 shows that the eventual symmetry deterioration upon use of smaller grids is smaller in ammonia compared with methane, with the heavy atom error also being similar for both molecules (0.16 kJ mol^{-1}).

Before we can study the effect of grid size reduction involving CCSD(T) wavefunctions instead of CCSD ones, we first look at the differences between CCSD and CCSD(T) energies themselves. The test set of 17 molecules consists of 7 s period hydrides, N$_2$, CO, NO$^+$, CN$^-$, C$_2$H$_2$, F$_2$… FH, (H$_2$O)$_2$, H$_2$O$_2$, C$_2$H$_4$ and C$_2$H$_6$. Table S12 presents the results of CCSD(T) and CCSD calculations enabling a comparison between the two. The basis set is uncontracted 6–31 + +G(2d,2p), and the grid is fixed to 15–30 (i.e. 350 angular Lebedev points and 30 radial ones). The conclusions are complex and listed in the SI, but an overall message that can be extracted is that moving from CCSD to CCSD(T) can increase or decrease the electron–electron (e-e) energy

of an atom in the *A-B* approach and similarly for the *A-A'* (ESP) approach.

Table S13 lists energy errors obtained by the CCSD(T) method and the uncontracted $6-31+G(d,p)$ basis set. The errors are smaller for the A-A' 2PDM-M/CCSD(T) (i.e. the modified 2PDM resulting from the 2PDM/CCSD(T) less the corresponding 2PDM/M) than for any other approach. The smallest grid, unsurprisingly, has the largest recovery error of 3.10 kJ mol^{-1}, while the next one is 1.23 kJ mol^{-1}, which does not correspond to the next smallest grid but the 6–12/7 grid. The latter recovery error is actually acceptable from the point of view of force field development. The corresponding values, where the subtraction of 2PDM/M is replaced by the subtraction of 2PDM/HF (2PDM-HF/CCSD(T)), have much larger values for the A-A' approach, that is, of the order of tens of kJ mol^{-1}. However, when comparing the corresponding *A-B* values, the errors are often smaller for the 2PDM-HF/CCSD(T) method compared to 2PDM-M/ CCSD(T). Overall, two observations are clear: (i) removing an approximate 2PDM from the full 2PDM significantly reduces the size of grid needed to integrate to a satisfactory error and (ii) *A-A'* benefits from small grids much more than *A-B*, if not uniquely so.

The results in Table S13 present the root-mean-square errors for a given grid and thus give important data for the question of what is an appropriate grid to employ. These arise from the individual recovery errors of all the first-row 'hydrides' (i.e. LiH to HF), determined with reference to the 47–60 grid energies and then squared and summed together and then averaged and then square rooted. It can be seen that for 2PDM-M/CCSD(T) with the uncontracted $6-31+G(d,p)$ basis almost all grids provide a reasonable recovery error (~ 1 kJ mol^{-1}), with the exception of the smallest grid. Thus, which grid is chosen is a function of what one wants to achieve, with 4–10 on non-H and 4–7 on H (referred to as 4–10/7) being acceptable for A-A' analysis of a 2PDM-M/CCSD(T), but for the A-B equivalent where none of the given grids is acceptable. In considering 2PDM-HF/CCSD(T) 2PDM with the A-A' analysis, an 8–11/7 grid seems to be acceptable (Table S13), which is repeating the significance of the point made before concerning the removal of the Müller-approximate 2PDM from the true 2PDM. We will return to the issue of grids in Sect. 3.4.

3.2 Reducing the number of 2PDM matrix elements based on the removal of small elements

We have just presented one way to reduce CPU time for integrating the 2PDM by subtracting off an approximate 2PDM from it. Here, we present the potential of a completely different method, that of matrix element removal, to see if we can gain CPU time this way. Section 2 of the SI reports

all the data and technical details behind this attempt to speed up the computation of atomic electron correlation energies.

A first set of computational experiments was carried out on MP2 rather than CCSD wavefunctions of three systems: glycine, $(H_2O)_3$ and Ne_2 with all data listed in Tables S14 to S22. A cut-off criterion was systematically increased, eventually in steps of ten, and the corresponding recovery error monitored. A sudden dramatic change occurred in the recovery error at 0.1×10^{-5} for glycine and at 0.1×10^{-6} for $(H_2O)_3$. Increasing these respective cut-off values results in a recovery error (far) above 1 kJ mol^{-1}, which is undesirable because we strive for a sub-kJ mol^{-1} error. At those respective cut-off values, we retain about half of the total number of elements in the 2PDM. This saving is welcome because it corresponds to the elimination of billions of entries. The neon dimer behaves differently in that the 0.1×10^{-10} cut-off already slashes 70% of the elements without affecting the recovery error much. The next step was to find out the speed-up gained by using these reduced sizes of 2PDM. Increasing the cut-off criterion to 0.1×10^{-7} leads to a significant improvement in the recovery error.

The second and final test pertains to CCSD wavefunctions. Table S22 gives the number of matrix elements of various sizes for solvated zwitterionic glycine (2 waters, 358 basis functions). We carried out MORFI calculations of the electron correlation energies in the AO basis although we obtain our 2PDM in the MO basis when using PySCF but in the AO basis when using G09. Table S22 lists the number of MO and AO matrix elements by size range. The degradation of the energy terms, calculated by MORFI, from throwing away small matrix elements is significant once one employs a cut-off of around 10^{-7} for matrix element removal. However, most of the matrix elements are larger than this cut-off and thus the gain in removing matrix elements is not great. This is why we decided not to remove matrix elements.

3.3 Basis set extrapolation

Another strategy to obtain electron correlation energies possible faster is to find a relationship between the energies calculated using basis sets of increasing size. In particular, we wonder if one can extrapolate energies to some sort of a limit while systematically increasing the basis set size. In other words, given a quantum method, can we find the basis set limit of the energies? Many details are given in Sect. 3.1 of the SI.

Dunning et al. have introduced the aug-cc-pVXZ ($X = D$, T, Q, 5 and 6) basis sets, and various authors have used them to extrapolate to limits, often by means of the extrapolation [26, 27] approach of the co-workers of Helgaker. We have adopted the aug-cc-pVXZ basis sets, removed the angular functions higher than *d*, uncontracted them and calculated the IQA energies (*A-A'*) for each of the basis sets. Figures S1

to S14 show the results for all 7 hydrides in the second period, with two successive figures for each compound: one for the varying element (Li to F) and one for the hydrogen. These energies are for CCSD wavefunctions with the cores included and Cartesian d-functions.

In addition, the Hartree–Fock component of 2PDM has been removed from the CCSD equivalent (i.e. 2PDM-HF/CCSD). As observed above, removing an approximate 2PDM from the true 2PDM enables a smaller grid to be used, when integrating. In addition, we employed the aug-cc-pVDZ IQA values as a reference and subtracted this value from all other larger basis set values. Hence, we extrapolate a difference rather than an absolute value. These obtained energy differences are denoted Δ(A-A') and use the energy generated at aug-cc-pVDZ as the zero reference. In Figures S1 to S14, the X is the value given by aug-cc-pVXZ, with $X = D = 2$, $X = T = 3$, $X = Q = 4$, $X = 5$ and $X = 6$, with 7 being the extrapolated point. For Li and Be, 6 refers to the mixed basis sets aug-cc-pV5Z on Li/Be and aug-cc-pV6Z on H. The extrapolation involved the last three points and employed Aitkin's δ^2 process [28], which is a three-point extrapolation.

Considering the data in Figs. S1 to S14 as a whole, we see that the diagrams for ammonia cannot be usefully extrapolated (Figs. S9 and S10). However, given the results for the other systems, it is perhaps not too difficult to see that the results for the final point would be close to the limit for both N and H.

Considering LiH (Figures S1 and S2) shows that lithium's Δ(A-A') energy difference smoothly tends to a value of about -41 kJ mol^{-1}. However, for H, the curve fitted to the energies undulates, but, given the very small energy scale, the actual energies practically converge to 0.1 kJ mol^{-1}. For BeH$_2$ (Figs. S3 and S4), both Be and H tend to converge to -27 and -2.2 kJ mol^{-1}, respectively. For B in BH$_3$, the extrapolated point's energy (8.4 kJ mol^{-1}) is more similar to the value of aug-cc-pV5Z (8.5 kJ mol^{-1}) than that of the larger aug-cc-pV6Z with 7.9 kJ mol^{-1} (Fig. S5). For H in BH$_3$, the extrapolated value is about -6.9 kJ mol^{-1} (Fig. S6). Turning to CH$_4$ (Figs. S7 and S8), the data behave similarly to those for B in BH$_3$. The C atom shows an extrapolated value of 43 kJ mol^{-1}, which is more like the aug-cc-V5Z result (44 kJ mol^{-1}) than the aug-cc-pV6Z result of 40 kJ mol^{-1}. The Δ(A-A') energy difference for H in CH$_4$ is -9 kJ mol^{-1}. The results for NH$_3$ (Figs. S9 and S10) have been briefly mentioned already as being linear, thereby lacking asymptotic behaviour and thus not being extrapolatable in a meaningful way. The best energies are 62.8 and -14.0 kJ mol^{-1} for N and H, respectively. For water, we see that the hydrogen extrapolation is excessive rather than then ending up very to the aug-cc-pV6Z value. The reason for this extrapolation error is that the data used in the extrapolation are approximately linear. Finally, Figs.

S13 and S14 show the data for HF. The extrapolation of the Δ(A-A') energy for F (21.6 kJ mol^{-1}) is reasonable given that the small energy tempers the potential adverse effect of the oscillation (Figs. S13). Figure S14 offers the same conclusion for H with a proposed converged energy difference of -5.4 kJ mol^{-1}.

The data for the 'hydrides' given above do point to the fact that the Δ(A-A') are converging to a constant value, although the data do not appear universal on this point, with NH$_3$ and H$_2$O being notable exceptions. The extrapolation procedure seems to yield mixed quality of results, as some extrapolated points resemble smaller basis set results than the larger ones, while visual inspection of these points indicates otherwise. Generally, it does seem that the s-, p- and d-functions in the uncontracted aug-cc-pV6Z basis set are reasonably close to the basis set limit for most of the first-row elements.

We can study again these first-row hydrides just as we have discussed above, but with only the s- and p-functions of the Dunning aug-cc-pVXZ ($X = D$, T, Q, 5 and 6) basis sets retained. Many details are given in Sect. 3.2 of the SI. The results of these calculations are graphically displayed in Figures S15 to S21 for the heavy atoms only. With the exception of BH$_3$, the data are fortunately 'plateauing'. In the case of BH$_3$, the energy differences for the three points used in the extrapolation are almost linear and the extrapolation gives a meaningless or wrong result.

Taken as a whole, the data displayed in Figures S1 to S21 generally indicate an advance to the basis set limit for these hydrides. Although more research needs to be carried into questions such as to why nitrogen in NH$_3$ seems to be converging to a limit for an sp basis but not for the equivalent spd basis set.

We now consider the effect of adding f-functions to the Dunning basis sets for non-hydrogen atoms and d-functions for H atoms, as previously we only considered the spd/sp (i.e. spd on the heavy atoms and sp on H) or sp/s components of these basis sets. Many details are given in Sect. 3.3 of the SI. Tables S23 to S26, respectively, give the relevant energies for LiH, BeH$_2$, H$_2$O and HF with the sp/s, spd/sp and spdf/spd sets of functions of the Dunning basis sets. We note that aug-cc-pVDZ has no f-functions by definition. The energies given (in Hartrees) are for the 2PDM-HF/CCSD approach and thus reflect the effects of correlation only. For the heavy atoms, the effect adding f-functions to an spd basis is minor compared to adding d-functions to an sp basis set, which is as expected. However, the tables do highlight that odd results can occur with the 2PDM-HF/CCSD approach, as seen in Table S25 (water), where the correlation energy changes sign on the addition of d-functions on the hydrogens and f-functions on O. It seems that Hartree–Fock places too much energy on H, for the sp/s and spd/sp basis sets, and correlation corrects this, while too little energy is placed

on H with the spdf/spd basis sets and correlation makes up for the shortfall. In addition, the correlation energy of the F atom in HF changes by 68.7 kJ mol^{-1} with the deployment of f-functions (aug-cc-pV6Z). It seems that f-functions are generally vital to the accurate description of an atom, whose importance, not surprisingly, grows as one moves from Li to F.

We now consider what happens in the water trimer if we extrapolate the individual A-A' energies together with the dimer and monomer. Many details are given in Sect. 3.4 of the SI. Tables S27, S28 and S29, respectively, contain the results for the A-A' values of the water monomer, dimer and trimer (see Fig. 1 for labelling) with the "code 20" angular grid (i.e. 590 Lebedev points, Table S1) and a 20-point radial grid. Included in these tables are the extrapolated values by the method of Helgaker and co-workers and that of Aitkin. While the extrapolated values do not differ wildly, with one exception (dimer in Table S28), they are not really consistent with each other most of the time. The results do indicate that we can obtain atomic energies to about 2.6 to 5.2 kJ mol^{-1} in accuracy. However, these values are perhaps a little too large for our force field work. In order to improve on these rather inaccurate values, we see if differences, as opposed to absolute values, extrapolate better. The results are given in Tables S30, S31 and S32. These tables give the energy change upon forming a dimer compared to a monomer, the trimer compared to the monomer, and the trimer compared to the dimer. They thus represent changes in intra-atomic energies due to hydrogen bond formation. From Tables S27-S32, the section labelled 'range' gives how much the values in the above column (A-A' terms) range over (i.e. the difference between the largest value and the smallest value). Comparing these 'range' values for Tables S30-S32 with those of Tables S27-S29 demonstrates that the values in Tables S30-S32 are generally smaller because the values in Tables S31-S32 are differences due to hydrogen bond formation, while those in Tables S27-S29 are absolute values. We conclude that considering how much an A-A' value changes on forming the hydrogen bond(s), rather than considering the absolute value, leads to a narrowing of the spread of determined values. This in turn implies that the error in a given value is reduced.

We now return to the question if a small basis set can be used to predict the energies that a larger basis set gives. In particular, we consider here how the various atomic correlation energy terms for H and O of water vary with geometry. Here, the aim is to find a pattern enabling a small basis set to be used, which can then have a correction factor applied to it, in order to return the equivalent to a larger basis set result. Initially, we changed the bond angle by $\pm 10°$ and lengthened and shortened of one of the bonds by 0.01, 0.05 and 0.10 Å. The method used was 2PDM-HF/CCSD alongside various truncated basis sets aug-cc-pVXZ ($X = D, T, Q,$

5 and 6) used before in this work. The results for oxygen in water are given in Table S33. In Table S34 are given results that derive from Table S33, although the units are now kJ mol^{-1} as opposed to in Hartrees. These tables show that the difference between the oxygen correlation energy with the aug-cc-pV6Z and the aug-cc-pVTZ basis sets is relatively constant (from a maximum of 33.13 to a minimum of 29.87 kJ mol^{-1}). Although this range is perhaps larger than one would like, the variations follow a set of trends. In particular, for angle bending the change is moderate ($+0.20$ or -0.43 kJ mol^{-1}), while the bond compression the change is about 1 kJ mol^{-1} per 0.05 Å. The situation is a bit different for bond lengthening, where the 0.05 Å extension yields a very small change of 0.02 kJ mol^{-1} and a 0.31 kJ mol^{-1} energy change when the bond is extended by 0.1 Å. The reduction in the number of basis functions in employing aug-cc-pVTZ over aug-cc-pV6Z is 133 (i.e. 254−121). Effectively, this means reducing computer time from 3.5 days down to a few hours for each geometry.

We now look at the effect of the basis set upon geometry change for the hydrogen atoms in H$_2$O. We have considered each hydrogen separately, as one of them has its bond length changed (except in geometries that involve bends only), while the other does not have its bond length changed at all. The determined correlation energy is given in Table S35 for each geometry. Note that generally the sign changes when f-functions(O)/d-functions(H) are added to the basis set, but for bond compression this is not always the case (Table S35 given in red). However, the largest basis set value is always negative. Figure S22 shows that the change in energy, upon for geometry variation, between aug-cc-pV6Z (truncated) and aug-cc-pVTZ (truncated) behaves smoothly. Thus, at least for modest geometry variations it should be possible to predict with aug-cc-pV6Z accuracy while using only the aug-cc-pVTZ basis set. This result is similar to that observed for O above and gives confidence that one can correct aug-cc-pVTZ basis set results to yield aug-cc-pV6Z accuracy.

Finally, we consider the hydrogen whose bond length remains fixed during the bending, as well as for the extension and compression of the other hydrogen. However, it is interesting to see how the correlation energy of this atom changes due geometry changes elsewhere in water. We note, as we did for the other hydrogen, that the addition of f-functions on O and d-functions on H changes the sign of the correlation. There are no exceptions this time. The f-functions on O, or the d-functions on H, are the second level of polarisation functions, and thus, one expects subtle effects arising from these, rather than an unsubtle change in sign. However, the actual change in energy in going from an spd basis on O, and an sp basis on H, to an spdf basis on O, and an spd basis on H, is moderate, that is, of the order 10–11 kJ mol^{-1}, although this is not perhaps the small subtle effect we were expecting. We note that the change in energy

in going from an sp basis O, and an s basis on H, to an spd basis on O, and an sp basis on H, is 16–18 kJ mol^{-1}. When we consider these two changes in basis set and as the associated correlation energy increases, it is apparent that the basis set is not complete and thus the correlation energy is not converged to a constant value with respect to higher angular momentum basis functions. In other words, 16–18 kJ mol^{-1} and 10–11 kJ mol^{-1}, respectively, represent the addition of the first and second levels of polarisation functions. The fact that the latter energy interval is not much smaller than the former indicates that further levels will be needed to reach a limit. Figure S23 shows the variation of the energy difference (between aug-cc-pV6Z and aug-cc-pVTZ) with geometry change of the other hydrogen. Hence, we investigate energy changes due to remote geometry change. The range of energy changes is small. Thus, as we concluded previously, prediction of the H atom's energy from the aug-cc-pVTZ basis to obtain an energy of aug-cc-pV6Z quality is possible. To this end, we have fitted these energy changes to a polynomial (near Table S37 in the SI) and thus possess a mathematical expression to convert from aug-cc-pVTZ energies to aug-cc-pV6Z energies.

3.4 Which grid should be employed?

Previously, we have mentioned the grid size when trying to judge the effect of removing approximate 2PDM from the true 2PDM. We observed then that one could obtain good recovery errors with smaller grids when not using the full 2PDM. Thus, this is the way forward. However, this route begs the question: which grid should be used? In Tables S38 to S43, we consider this question in some detail. Many of the tables in the SI contain results for more than one grid applied to the same molecular system. However, some of these are large grids used to assess the accuracy of the energies that we obtain with smaller grids.

In Table S38, we consider the effect of angular and radial grid change on uncontracted STO-3G water (A-A'). It is clear that with 1 kJ mol^{-1} as recovery threshold all these energies fall within it. However, the radial grid needs to be greater than 20 and the angular grid greater than angular grid code 10 (170 points) in order to obtain a recovery error of less than 0.1 kJ mol^{-1}. A quick test looks at the effect of using one grid (29–60) on two larger basis sets, which substantially increases the error with size of basis set, but none of the errors were large. Table S39 presents a systematic change in radial and angular grids for the first-row hydrides, with the 2PDM-HF/CCSD and A-A' approaches. All grids perform reasonably well to very well, adopting the FFLUX criterion of 1 kJ mol^{-1} as an acceptable cut-off for the recovery error. However, the recovery error

rises sharply when the radial grid drops from 20 to 10 points and the angular grid drops to that of code 10 (170 points).

Presented in Table S40 are the results for the Lebedev 20 and 17 grids with 10 and 20 radial points (with CCSD-HF and A-A'). Some additional molecules with extra elements (S, Cl) were added to obtain a more balanced picture of the effect of the grid on the recovery error. It is clear from this table that one definitely requires 20 radial points for S and Cl (A-A'). For the negative ions (OF^{-} and OH^{-}), the smaller grids often give surprisingly good recovery errors.

In Table S41, we consider the effect of basis size again with grid size, having briefly considered it vide supra. These results seem to contradict the one given just before in that the larger basis set has the smaller recovery error. This may be because the previous comparison was not part of a systematic basis set increase in size. However, the last line of Table S41 does make clear that what works for one basis set may not work for a related one.

The significance of a smaller grid for H than for other atoms and for Li to F compared to sulphur needs to be addressed. Generally, it can be observed that H needs a smaller grid to for accurate integration than any other atom, not surprisingly. Thus, in Table S10 we employ a 32–30 grid as a reference and note that the recovered energy compared to this reference energy, for hydrogen, is ~ 0.01 kJ mol^{-1} (5–10 grid), while for C, the same comparison yields an error of about an order of magnitude larger. However, in the context of FFLUX, none of these errors are unduly large. We note that LiH is an exception and that the diffuse nature of the electron density around H creates problems (Table S9) with the hydride generally having a larger recovery error than Li^{+}. The molecules BeH$_2$ and BH$_3$ do not to show this behaviour, and the heavier atom has a larger recovery error than the hydride group. We assume that the different behaviour of H bound to Li, rather than to Be or B, arises from the fact that the higher charge on the Be or B compared to Li modifies the nature of the hydride. Although hydride transfers are part of the biochemistry of life, it is unlikely that behaviour similar to that of the extreme case of H^{-} in LiH will be much observed. In Tables S42 and S43, a systematic study of the A-A' energy change of H (2PDM-HF/CCSD) with grid is given for the first-row hydrides. The conclusion is similar to previous that the energy does not change with grid very much, with the exception of LiH. In Table S13, we give the root-mean-square error hydrogen energies for the A-A' 2PDM-M/CCSD(T) method with uncontracted 6–31 + G(d,p) basis set for the first-row 'hydrides', with various grids. The conclusion is much as before: almost any grid gives an acceptable error. However, we stress that this conclusion materialises using the 2PDM-M/CCSD(T) with the A-A' approach. The main conclusion is that almost any grid, except for the smallest, gives an acceptable error of

hydrogen while using 2PDM-HF/CCSD, 2PDM-M/CCSD, 2PDM-HF/CCSD(T) or 2PDM-M/CCSD(T) with the A-A' method.

The presence of sulphur in biological molecules is well known, where it can be divalent, as in amino acids or oxidised, as in taurine (2-aminoethane-1-sulfonic acid). We ask how its oxidation state affects the grid needed for integration. Of course, the presence of two shells of core electrons, compared to oxygen's one shell, is expected to affect the recovery error. If we work with 2PDM-M/CCSD(T) and A-A' as the preferred method, and consider CH_3SH and HSO_3^-, then we see how the grids of 10–10 and 15–15 perform. For CH_3SH, the 10–10 grid has a recovery error of 3.2 kJ mol^{-1}, while the corresponding error for the 15–15 grid is 0.4 kJ mol^{-1}. These values are obtained by reference to an AIMAll calculation. We next considered the bisulphite ion, which yields errors of $+0.4$ and -0.3 kJ mol^{-1}, for the 10–10 and 15–15 grids, respectively.

3.5 Transferability study: capped histidine

The details of this computational experiment are in Sect. 3.7 of the SI. The purpose is to quantify the effect on the atomic correlation energies of substituting a small part of a molecule. This type of test provides information on the transferability of these energies. In other words, what kind of loss in accuracy is caused if a molecule is truncated? Can one safely calculate the electron correlation energy for a given larger molecule from a truncated, smaller molecule?

The test system is a 29-atom molecule representing histidine as if it were part of a protein. This amino acid is capped by two methyl groups, so that it represents histidine in a peptide chain. We optimised capped histidine at uncontracted MP2/6-31G(d,p) level and obtained the geometry displayed in Figure S23. As during a dynamics run the geometry is expected to change, we made no attempt to locate the global minimum for this system. Two substitutions are applied in succession: one methyl cap is replaced by a hydrogen at one side and then the other methyl cap, at the other side, is also substituted by a hydrogen. The α carbon is the backbone atom that is most remote from the two substitutions and the effect of these is

to change its energy by 0.6 and 0.4 kJ mol^{-1}, successively over the two substitutions. The non-hydrogen atoms of imidazole group are affected by less than 0.2 kJ mol^{-1}, while the hydrogen atoms suffer changes of less 0.1 kJ mol^{-1}. The same energy difference patterns are seen between two geometries (one geometry-optimised, the other not) of the same molecule: double-capped histidine. Although the geometry did not change that much, sub-kJ mol^{-1} changes emerge with the exception of a backbone nitrogen.

3.6 Non-standard bonds as found in S_N2 transition states

So far we have considered covalent bonds that have not been stretched, but in transition states this may not be the case. To find out if such stretched bonds still conform to the treatment presented so far, we considered the classic, symmetric S_N2 reaction of the fluoride ion with methyl fluoride. We employed the 2PDM-M/CCSD(T) with the A-A' approach and the uncontracted 6–31+G(d,p) basis set. Geometries were determined at the CCSD/uncontracted 6–31+G(d,p) level. Table 2 presents the results, with the final row displaying the sum of the entries of the respective column above, and also shows the recovery error with respect to the largest grid (32–30), in brackets. It is clear that the 20–20 and 15–15 grids have modest recovery errors of -0.04 and -0.29 kJ mol^{-1}, respectively. Hence, our approach of using the 15–15 grid is still valid for stretched bonds.

4 Conclusions

It is possible to study electron correlation effects using the quantum topological energy partitioning method interacting quantum atoms (IQA). This is our preferred route because of its rigour and minimality. We believe that these attributes assist in making this route more future-proof than ad hoc and approximate schemes to include dispersion energy in ab initio calculations. Yet the latter

Table 2 The 2PDM-M/CCSD(T) A-A' energies for the symmetric transition state F-CH$_3$-F$^-$

Atom/Grid	32–30	20–20	15–15	10–10
C	0.060856	0.060864	0.060844	0.060527
F1	0.042569	0.042569	0.042701	0.042529
F2	0.060856	0.012427	0.060844	0.060527
H1	0.012427	0.012427	0.012427	0.012430
H2	0.012427	0.012427	0.012427	0.012430
H3	0.012427	0.012427	0.012427	0.012430
sum	0.201561 (0.00)	0.201578 (-0.04)	0.201671 (-0.29)	0.200874 (1.80)

The final row displays sum of the atomic energies and gives the recovery error, in brackets, with respect to the largest grid. Units: Hartrees and in brackets kJ mol^{-1}

are very popular, helped by their very low computational cost. In contrast, the IQA atomic electron correlation energies, calculated from the humongous two-particle density matrix (2PDM), need orders of magnitude more CPU time to obtain. Yet, these energies benefit from the conceptual advantages that IQA offers, being part of quantum chemical topology.

In the recent past, we have embarked on various successful initiatives to speed up the calculation of atomic 2PDM energies: (i) the OpenMP (but not MPI) parallelisation of the program MORFI, (ii) the ESP A-A' method, and (iii) the observation that machine learning (which underpins the force field FFLUX) needs much fewer training points compared to Coulomb and exchange energies. Here, we explore (i) the potential sparseness of the 2PDM, (ii) molecular truncation based on transferability (e.g. capped histidine), (iii) basis set extrapolation and (iv) quadrature grid optimisation. The latter initiative is the most successful and indeed dramatically so.

In more detail, we have shown that our form of the Müller 2PDM is correct when compared to AIMAll's equivalent, to within numerical accuracy. With this proof, we were then able to remove the Müller 2PDM from the full 2PDM and obtain a matrix that requires a smaller grid to obtain the desired energies than expected had we used another form of the 2PDM (A-A' approach). In particular, a proof-of-concept calculation on a water trimer showed that a grid of merely a couple of hundred quadrature points generates an energy that differs only ~ 1 kJ mol^{-1} from that generated by a grid more than 150 times larger. Hence, it is clear that the pure electron correlation is slowly varying "ripple" that is easy to integrate over atomic volumes.

We then considered the number of matrix elements that can be neglected from the 2PDM when integrating it. However, the initially chosen cut-off point of 0.1×10^{-6} turned out to be too large and a smaller value was required for a set of diverse molecular systems. This finding, and the extra time needed to test if a matrix element was small, resulted in us keeping all the matrix elements.

The energies generated by the A-A' approach combined with 2PDM-M/CCSD(T) compared to the corresponding 2PDM-HF/CCSD(T) energies generally tend to be smaller, although we stress that this is only a general observation and contrary examples are known. The results for MORFI-generated energies for the 2PDM-HF/CCSD(T) method compared to the corresponding 2PDM-HF/CCSD energies can be quite different for molecules with triple bonds, indicating a need for triple excitations in these cases. The A-B energies are inconsistent as to whether 2PDM-M/CCSD(T) or 2PDM-HF/CCSD(T) is the best, unlike in the A-A' case where 2PDM-M/CCSD(T) has the lowest recovery error.

Generally, it is possible to extrapolate the A-A' energies of an atom in a molecule to a limit, when employing an spd basis set (Li to F). However, there are a small number of exceptions, where extrapolation was not possible or was very bad. In these cases, we expect that the final, determined point (aug-cc-pV6Z) will provide an answer close to the expected limit, based on the other extrapolatable hydrides. Exactly which angular moment basis functions are needed to reach an absolute limit of these A-A' energies has not been determined. However, the extrapolated energies were not converged when including f-functions. There is a relationship between the atomic energies determined with the aug-cc-pVTZ and the aug-cc-pV6Z basis set results, leading to the hope that calculations with a small basis set can be scaled up to yield the equivalent larger basis set energies. However, when extrapolating the A-A' energies in the case of water oligomers, we found that extrapolating the changes, on hydrogen bond formation, rather than the absolute values, reduces the anticipated error. The question as which grid should be employed in determining the A-A' energies has been considered. It is clear that hydrogen requires very small grids, with the 2PDM-M/XXXX (XXXX = CCSD, CCSD(T) or MP2) methods, but larger ones are needed other atoms, with sulphur needing larger grids than Li to F.

Overall, the most important observation is that the computation of IQA electron correlation energies benefits enormously from the use of tiny grids. The knack is to subtract the Hartree–Fock part from the two-particle density matrix, as well as a one-particle-based approximation of this matrix (such as Müller's). Put differently, the presence of the Hartree–Fock component causes the integration to waste grid points. In other words, the pure (real two-particle) electron correlation allows itself to be integrated accurately with very few quadrature points.

Author contributions MV and PLAP wrote the manuscript; MV prepared the figures and the tables; MV carried out the calculations.

Declarations

References

1. Blanco MA, Martín Pendás A, Francisco E (2005) J Chem Theor Comp 1:1096–1109
2. Popelier PLA, Kosov DS (2001) J Chem Phys 114:6539–6547
3. Bader RFW, Beddall PM (1973) J Am Chem Soc 95:305–315
4. Bader RFW (1990) Atoms in Molecules. A Quantum Theory. (Oxford Univ. Press, Oxford, Great Britain)
5. Parrish RM, Sherrill CD (2014) J Chem Phys 141:044115
6. Parrish RM, Parker TM, Sherrill CD (2014) J Chem Theor Comp 10:4417–4431
7. Phipps MJS, Fox T, Tautermann CS, Skylaris C-K (2015) Chem Soc Rev 44:3177–3211
8. Guevara-Vela JM, Francisco E, Rocha-Rinza T, Martín Pendás A (2020) Molecules 25:4028
9. Heidar-Zadeh F, Ayers PW, Verstraelen T, Vinogradov I, Vöhringer-Martinez E, Bultinck P (2018) J Phys Chem A 122(17):4219–4245
10. Suarez D, Diaz N, Francisco E, Martin Pendas A (2018) Chem Phys Chem 19:973–987
11. Grimme S, Hansen A, Brandenburg JG, Bannwarth C (2016) Chem Rev 116:5105–5154
12. Popelier PLA (2022) J Mol Model 28:276
13. Popelier PLA (2016) Phys Scr 91:033007
14. McDonagh JL, Vincent MA, Popelier PLA (2016) Chem Phys Lett 662:228–234
15. Silva AF, Duarte LJ, Popelier PLA (2020) Struct Chem 31:507–519
16. Casals-Sainz JL, Guevara-Vela JM, Francisco E, Rocha-Rinza T, Martín Pendás A (2020) J Comp Chem 41:1234–1241
17. Vincent MA, Silva AF, Popelier PLA (2020) Z Anorg Allg Chemie 646:1–9
18. Müller AMK (1984) Phys Lett A 105(9):446–452
19. Schaftenaar G, Noordik JH (2000) J Comput Aided Mol Design 14:123–134
20. Vincent MA, Silva AF, Popelier PLA (2019) J Comp Chem 40:2793–2800
21. McDonagh JL, Silva AF, Vincent MA, Popelier PLA (2018) J Chem Theor Comput 14:216–224
22. McDonagh JL, Silva AF, Vincent MA, Popelier PLA (2017) J Phys Chem Letts 8:1937–1942
23. Silva AF, Vincent MA, McDonagh JL, Popelier PLA (2017) ChemPhysChem 18:3360–3368
24. Silva AF, Popelier PLA (2018) J Mol Mod 24:201
25. Vincent MA, Silva AF, McDonagh JL, Popelier PLA (2018) Int J Quant Chem 118:e25519
26. Halkier A, Helgaker T, Jorgensen P, Klopper W, Olsen J (1999) Chem Phys Lett 302:437–446
27. A Halkier, T Helgaker, P Jorgensen, W Klopper, H Koch, J Olsen, AK Wilson (1998). Chem Phys Lett **286**, 243–252
28. Noble B (1964) Numerical Methods: Volume 1. (Oliver and Boyd, Edinburgh and London, Great Britain)