



# Using mixed precision within DL\_POLY's force and energy evaluations: long-range interactions and fast Fourier transforms

HS Thorne

May 2018

©2018 Science and Technology Facilities Council



This work is licensed under a [Creative Commons Attribution 4.0 Unported License](https://creativecommons.org/licenses/by/4.0/).

Enquiries concerning this report should be addressed to:

RAL Library  
STFC Rutherford Appleton Laboratory  
Harwell Oxford  
Didcot  
OX11 0QX

Tel: +44(0)1235 445384  
Fax: +44(0)1235 446403  
email: [libraryral@stfc.ac.uk](mailto:libraryral@stfc.ac.uk)

Science and Technology Facilities Council reports are available online at: <http://epubs.stfc.ac.uk>

**ISSN 1358-6254**

Neither the Council nor the Laboratory accept any responsibility for loss or damage arising from the use of information contained in any of their reports or in any communication about their tests or investigations.

# Using mixed precision within DL\_POLY's force and energy evaluations: long-range interactions and fast Fourier transforms

H. Sue Thorne<sup>1,2</sup>

## ABSTRACT

DL\_POLY 4 is a molecular dynamics simulation code developed by CCP5. In this report, we consider the use of a simple mixed-precision methodology within the Smoothed Particle Mesh Ewald (SPME) method for systems with long-ranged ionic interactions. We show that such a method can be successfully used within DL\_POLY and that, for some of the test problems, the overall execution time was reduced by 14%.

This work forms part of the CoSeC-funded Software Outlook project.

**Keywords:** Molecular Dynamics, DL\_POLY, Ewald Method, FFT, Mixed Precision, Intel Xeon (IvyBridge), Software Outlook, CCP, CCP5

---

May 16, 2018

---

---

<sup>1</sup>The Hartree Centre, Science and Technology Facilities Council, Rutherford Appleton Laboratory, Harwell Campus, Didcot, OX11 0QX, UK.

<sup>2</sup>Correspondence to: sue.thorne@stfc.ac.uk

# 1 Background and motivation

In the last two decades, huge developments were made in molecular dynamics simulation methods, particularly in the ability to use large-scale parallel platforms. In this report, we consider the use of a mixed-precision methodology within the Smoothed Particle Mesh Ewald (SPME) method [4] for systems with long-ranged forces. In particular, our investigation hinges on the SPME implementation within DL\_POLY 4, which incorporates the domain decomposition method with a 3D Fast Fourier Transform, known as the Daresbury Advanced Fourier Transform (DAFT) [2]. In DL\_POLY 4, any real values are stored in double precision and operated on using double precision arithmetic (IEEE Standard 754 is assumed). On modern high performance computing architectures, there is normally the provision to use single precision arithmetic that is significantly faster than the corresponding double precision operation [6]. Additionally, the communication of real data between processors in a distributed system will normally be significantly faster when passing single precision data compared to double precision data. However, the desire to speed-up the code must be balanced with the requirements of getting an accurate solution. Indeed, if all real values are stored in single precision and operated on using single precision arithmetic, then DL\_POLY is unstable. Thus, it will be necessary to use single precision in a restricted manner, i.e., use mixed precision. In this report, we focus on using a mixed precision approach within part of the SPME calculation, namely, the DAFT component will use single precision as much as possible.

In Section 2, the Smoothed Particle Mesh Ewald Method is reviewed. The Discrete Fourier transform and DAFT are discussed in Section 3.1. Our mixed precision methodology is introduced in Section 4 and numerical results from a number of test problems are provided in Section 5. Section 6 contains the conclusions from this work.

## 2 The Smoothed Particle Mesh Ewald Method

In molecular dynamics, the conditionally convergent Coulomb sum is used to describe the system energy contribution from ionic interactions in periodic charged systems. The Ewald Sum replaces this sum by three distinct sums with guaranteed convergence:

$$E = E^S + E^L + E^{SELF},$$

where  $E^S$ , the *real space* sum, is cast in normal physical space and represents the short-range interactions;  $E^L$ , the *reciprocal space* sum, is cast in the reciprocal space of the unit cell and represents the long-range terms;  $E^{SELF}$  is the self interaction term. Suppose we have  $N$  ions in a vacuum at locations  $\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_N$  with point charges  $q_1, q_2, \dots, q_N$ , respectively. We define  $\epsilon_0 = 8.854 \times 10^{-12} \text{C}^2 \text{N}^{-1} \text{m}^{-2}$  to be the electric constant (or vacuum permittivity). Let the ions be subjected to periodic boundary conditions, which we describe using three repeat vectors  $\mathbf{c}_1, \mathbf{c}_2, \mathbf{c}_3$ : these form the *supercell*. Thus, if there is an ion with charge  $q_i$  at location  $\mathbf{r}_i$ , then there are also ions with charge  $q_i$  at  $\mathbf{r}_i + n_1 \mathbf{c}_1 + n_2 \mathbf{c}_2 + n_3 \mathbf{c}_3$ , where  $n_1, n_2$ , and  $n_3$  are arbitrary integers. We simplify the notation by writing an arbitrary repeat vector as  $\mathbf{n}L$ , where  $L$  represents the supercell. If we further assume that the charge distribution is of Gaussian form with standard deviation  $\sigma$  and the supercell has volume  $V$ , then  $E^S$ ,  $E^L$  and  $E^{SELF}$  are defined as

$$E^S = \frac{1}{4\pi\epsilon_0} \frac{1}{2} \sum_{\mathbf{n}} \sum_{i=1}^N \sum_{\substack{j=1 \\ j \neq i}}^N \frac{q_i q_j}{|\mathbf{r}_i - \mathbf{r}_j + \mathbf{n}L|} \operatorname{erfc} \left( \frac{|\mathbf{r}_i - \mathbf{r}_j + \mathbf{n}L|}{\sqrt{2}\sigma} \right), \quad (1)$$

$$E^L = \frac{1}{2V\epsilon_0} \sum_{\mathbf{k} \neq 0} \sum_{i=1}^N \sum_{j=1}^N \frac{q_i q_j}{k^2} e^{i\mathbf{k} \cdot (\mathbf{r}_i - \mathbf{r}_j)} e^{-\sigma^2 k^2 / 2}, \quad (2)$$

$$E^{SELF} = \frac{1}{4\pi\epsilon_0} \frac{1}{\sqrt{2\pi}\sigma} \sum_{i=1}^N q_i^2. \quad (3)$$

In this report, we focus on the calculation of  $E^L$ . The calculation of  $E^S$  is considered in [5]. The SPME method improves the performance of the Ewald reciprocal space sum calculation by using (complex) Fast Fourier Transforms (FFTs). See [2] for further details about the framework in which the FFT is used.

### 3 Discrete Fourier Transforms

Let  $y_0, y_1, \dots, y_{N-1}$  be complex numbers forming a one-dimensional array. The Discrete Fourier Transform (DFT) is defined as

$$Y_k = \sum_{n=0}^{N-1} y_n w_N^{k,n}, \quad k = 0, \dots, N-1, \quad (4)$$

where

$$w_N^{k,n} = e^{-i2\pi kn/N}. \quad (5)$$

Thus, a new one-dimensional array is formed. If we were to evaluate this  $Y_k$  directly for each  $k$ , then  $O(N^2)$  operations would be required. A Fast Fourier Transform (FFT) computes all  $Y_k$  using just  $O(N \log N)$  operations.

#### 3.1 Fast Fourier Transforms

Let  $a(x)$  be a polynomial of degree  $N-1$ , where  $a(x) = a_0 + a_1x + a_2x^2 + \dots + a_{N-1}x^{N-1}$ . Define the  $b(x)$  and  $c(x)$  to be polynomials of degrees  $\lfloor (N-1)/2 \rfloor$  and  $N-1 - \lceil N/2 \rceil$ , respectively, where  $b(x) = a_0 + a_2x + a_4x^2 + \dots$  and  $c(x) = a_1 + a_3x + a_5x^2 + \dots$ . Noting that  $a(x) = b(x^2) + xc(x^2)$  and that (4) is a polynomial with  $x = e^{-i2\pi k/N}$ , then the time  $T(N)$  to calculate all  $Y_k$  satisfies

$$\begin{aligned} T(N) &= 2T(N/2) + O(N) \\ &= O(N \log N). \end{aligned}$$

In DL\_POLY, a three-dimensional version of the DFT is used: for an array  $y_{n_1, n_2, n_3}$  with  $n_l = 0, 1, \dots, N_l - 1$  and  $l = 1, 2, 3$ , the three-dimensional DFT is defined as

$$Y_{k_1, k_2, k_3} = \sum_{n_1=0}^{N_1-1} w_{N_1}^{k_1 n_1} \sum_{n_2=0}^{N_2-1} w_{N_2}^{k_2 n_2} \sum_{n_3=0}^{N_3-1} w_{N_3}^{k_3 n_3} y_{n_1, n_2, n_3}, \quad k_l = 0, 1, \dots, N_l - 1, \quad l = 1, 2, 3. \quad (6)$$

Note that this is a sequence of 3 one-dimensional DFTs, i.e.,  $Y$  can be calculated by performing a DFT along the  $n_3$  dimension, then apply a DFT along the  $n_2$  directions and, finally, apply a DFT along the  $n_1$  direction. Let  $N = N_1 N_2 N_3$ , then it can be shown that the complexity of this approach is  $O(N \log N)$ .

The Daresbury Advanced Fourier Transform (DAFT) [2] is a fully distributed parallel implementation of the 3D FFT, which makes use of the underlying domain decomposition, and is designed to maximise performance on architectures with a large number of processors. Unfortunately, the nature of the DFT means that there needs to be a lot of communication between the processors and if this communication can be done using single precision data instead of double precision, then there should be an appreciable time/energy saving in communication costs. This motivates our desire to investigate the use of mixed precision within the reciprocal SPME calculation and DAFT.

## 4 Mixed precision within the reciprocal SPME calculation

Our method for using mixed precision within the reciprocal SPME calculation can be attributed to the work presented in [3]. If the SPME method is called by DL\_POLY, then the SPME reciprocal space sum is computed once during each time step. The DAFT is called twice during each calculation of the SPME reciprocal space sum. In the current version of DL\_POLY, the values of  $w_N^{k,n}$  are computed once in double precision, stored in double precision and reused during each call to DAFT. All calculations within DAFT are carried out using double precision arithmetic. Our proposed mixed precision method is as follows:

- Compute  $w_N^{k,n}$  using double precision arithmetic but store the resulting value using single precision. These values are reused by each call to DAFT.
- Convert the input FFT data to single precision (complex) data.
- Form the FFT using single precision arithmetic, producing single precision (complex) output.
- Convert the output to double precision on-the-fly as needed by the SPME calculation.

In the original double precision version of DL\_POLY, we note that the first call to the 3D FFT function within each reciprocal SPME call is preceded by the conversion of real double precision data to double precision complex data, which then forms the input data to the FFT. In our proposed method, the real double precision data is converted to single precision complex data.

## 5 Numerical results

All of our tests were run on the Hartree Centre’s Napier Compute System [1], which consists of 360 nodes containing 2 x12 core Intel Xeon processors (IvyBridge E5-2697v2 2.7GHz) and 64GB RAM. The interconnect is Infiniband from Mellanox (FDR Connect-IB 56 GB/s). In all our runs, 24 processes per node were requested. DL\_POLY 4.09 was used as our base code and all double precision runs use this version. DL\_POLY was compiled using Intel MPI (version 5.1.1) and its `mpif90` wrapper, which points to the GNU compiler `gfortran`. The flag `-pg` was used to enable profiling and we always set the environment variable `GMON_OUT_PREFIX` to `gmon.out-‘/bin/uname -n’` to save the separate profiles for each MPI process.

For each test problem, number of MPI processes,  $n_p$ , and choice of precision used within the 3D FFT method, five separate runs were performed and the `gprof` profiler was used to analyse the timing profiles across all of the processes. In the following, we report times that are averaged (mean) across all five runs and all MPI processes. We are particularly interested in

*fft\_3d* the total time spent performing the 3D FFT kernel by an MPI process;

*spme* the total time spent performing the Ewald SPME kernel by an MPI process;

*total* the total DL\_POLY execution time for an MPI process.

### 5.1 Standard DL\_POLY test cases

As part of DL\_POLY, 28 small test problems are provided. Of these, 14 use the Ewald SPME method and we also immediately discounted test problems 01 and 03 because total execution time for the SPME method on a single processor was a fraction of a second. In Table 1, we list the test problems used and their attributes: the problems are numbered as in the DL\_POLY Manual [7]. When comparing the use of the mixed precision and double precision versions of DL\_POLY, we found that the mixed precision version that there was very little or no loss in accuracy in the output from DL\_POLY.

In Table 2, we list the (mean) time spent by each MPI process in the various kernels (listed in Section 5) when either the mixed precision or double precision versions of the DaFT are used within the DL\_POLY run. The times in bold are those that are at most 2% larger than the best time.

Problem	Total system size	Temperature(K)	Description
02	51737 atoms	300	200 DMPC molecules in 9379 water molecules
04	99120 atoms	300	8 Gramacidin A molecules in 32096 water molecules
07	12428 atoms	300	Lipid bilayer in water
08 and 09	8000 charged points and 4000 shells	3000	MgO with adiabatic and with relaxed shell model
10	500 ions and 39000 atoms	300	Potential of mean force on K+ in water
18	34992 atoms	25	SPC IceVII Water with constraint bonds
19	34992 atoms	25	SPC IceVII Water with rigid bodies
20	28816 atoms	295	64 NaCl ion pairs with 4480 water molecules represented by constraint bonds and 4416 water molecules represented by rigid bodies
21	29052 particles	295	7263 TIP4P rigid body water molecules
22	44352 ions	400	Ionic liquid dimethylimidazolium
23	23712 ions	310	600 molecules of calcite in 6904 water molecules

Table 1: Standard DL.POLY test problems using the Ewald SPME Method and their attributes. Test problems 01 and 03 excluded due to them having very small execution times.

Problem	$n_p$	<i>fft_3d</i>			<i>spme</i>			<i>total</i>			<i>fft_3d/total</i>	
		mixed	double	ratio	mixed	double	ratio	mixed	double	ratio	mixed	double
02	1	<b>4.15</b>	5.26	0.79	<b>8.59</b>	9.67	0.89	<b>45.65</b>	46.85	0.97	0.091	0.112
02	2	<b>1.97</b>	2.31	0.85	<b>4.19</b>	4.55	0.92	<b>23.63</b>	<b>24.03</b>	0.98	0.083	0.096
02	4	<b>0.87</b>	1.10	0.79	<b>2.05</b>	2.29	0.90	<b>12.16</b>	12.44	0.98	0.072	0.088
02	8	<b>0.45</b>	0.52	0.87	<b>1.00</b>	1.08	0.92	<b>6.25</b>	<b>6.34</b>	0.99	0.072	0.082
04	1	0.77	<b>0.75</b>	1.03	<b>6.46</b>	<b>6.43</b>	1.00	<b>54.25</b>	<b>53.80</b>	1.01	0.014	0.014
04	2	0.37	<b>0.36</b>	1.03	<b>3.21</b>	<b>3.19</b>	1.01	<b>28.16</b>	<b>28.06</b>	1.00	0.013	0.013
04	4	0.18	<b>0.16</b>	1.13	<b>1.70</b>	<b>1.69</b>	1.01	<b>14.61</b>	<b>14.51</b>	1.01	0.012	0.011
04	8	0.09	<b>0.08</b>	1.13	<b>0.78</b>	<b>0.78</b>	1.00	<b>7.51</b>	<b>7.50</b>	1.01	0.012	0.011
07	1	<b>3.31</b>	<b>3.31</b>	1.00	<b>10.01</b>	<b>10.11</b>	0.99	<b>31.69</b>	<b>31.67</b>	1.00	0.104	0.105
07	2	<b>1.71</b>	<b>1.72</b>	0.99	<b>5.06</b>	<b>5.14</b>	0.98	<b>16.80</b>	<b>16.88</b>	1.00	0.102	0.102
07	4	<b>0.85</b>	0.95	0.89	<b>2.65</b>	2.78	0.95	<b>9.36</b>	<b>9.45</b>	0.99	0.091	0.101
07	8	<b>0.45</b>	0.57	0.79	<b>1.41</b>	1.54	0.91	<b>4.90</b>	5.02	0.98	0.092	0.114
08	1	<b>0.12</b>	0.14	0.86	<b>0.85</b>	0.89	0.96	<b>2.81</b>	<b>2.81</b>	1.00	0.043	0.050
08	2	<b>0.05</b>	0.06	0.83	<b>0.40</b>	0.42	0.95	<b>1.51</b>	1.55	0.97	0.033	0.039
08	4	<b>0.02</b>	0.03	0.67	<b>0.21</b>	<b>0.21</b>	1.00	<b>0.83</b>	<b>0.84</b>	0.99	0.024	0.036
08	8	<b>0.01</b>	0.02	0.50	<b>0.12</b>	<b>0.12</b>	1.00	<b>0.43</b>	<b>0.44</b>	0.98	0.023	0.046
09	1	<b>1.19</b>	1.49	0.80	<b>9.05</b>	9.29	0.97	<b>28.74</b>	<b>28.95</b>	0.99	0.0414	0.0515
09	2	<b>0.62</b>	<b>0.63</b>	0.98	<b>4.16</b>	<b>4.23</b>	0.98	<b>15.96</b>	<b>15.94</b>	1.00	0.039	0.040
09	4	<b>0.28</b>	0.33	0.85	<b>2.18</b>	2.24	0.97	<b>8.73</b>	<b>8.77</b>	1.01	0.032	0.038
09	8	0.15	<b>0.13</b>	1.15	<b>1.19</b>	1.21	0.98	<b>4.34</b>	<b>4.39</b>	0.99	0.035	0.030
10	1	0.24	<b>0.22</b>	1.09	<b>2.49</b>	<b>2.50</b>	1.00	<b>19.65</b>	<b>19.70</b>	0.99	0.012	0.011
10	2	0.12	<b>0.11</b>	1.09	<b>1.32</b>	<b>1.34</b>	0.99	<b>10.52</b>	<b>10.57</b>	0.99	0.011	0.010
10	4	<b>0.07</b>	0.07	1.00	<b>0.62</b>	<b>0.61</b>	1.02	<b>5.38</b>	<b>5.42</b>	0.99	0.013	0.013
10	8	<b>0.03</b>	0.04	0.75	<b>0.33</b>	<b>0.33</b>	1.00	<b>2.85</b>	<b>2.85</b>	1.00	0.011	0.014
18	1	1.59	<b>1.34</b>	1.19	<b>14.84</b>	<b>14.59</b>	1.02	<b>56.95</b>	<b>56.42</b>	1.01	0.028	0.024
18	2	0.77	<b>0.70</b>	1.10	<b>7.77</b>	<b>7.77</b>	1.00	<b>30.09</b>	<b>30.04</b>	1.00	0.026	0.023
18	4	0.38	<b>0.34</b>	1.12	<b>3.62</b>	<b>3.58</b>	1.01	<b>15.27</b>	<b>15.24</b>	1.00	0.025	0.022
18	8	0.19	<b>0.17</b>	1.12	<b>1.90</b>	<b>1.88</b>	1.01	<b>8.38</b>	<b>8.40</b>	1.00	0.023	0.020
19	1	<b>0.34</b>	<b>0.34</b>	1.00	3.37	<b>3.31</b>	1.02	<b>14.27</b>	<b>14.21</b>	1.00	0.024	0.024
19	2	0.19	<b>0.14</b>	1.35	1.78	<b>1.73</b>	1.03	<b>7.56</b>	<b>7.49</b>	1.01	0.025	0.019
19	4	0.09	<b>0.08</b>	1.13	<b>0.81</b>	<b>0.82</b>	0.98	<b>3.81</b>	<b>3.81</b>	1.01	0.024	0.021
19	8	0.05	<b>0.03</b>	1.67	0.44	<b>0.43</b>	1.02	<b>2.00</b>	<b>1.99</b>	1.01	0.025	0.015
20	1	<b>0.30</b>	0.35	0.86	<b>4.33</b>	<b>4.40</b>	0.98	<b>24.65</b>	<b>24.71</b>	1.00	0.012	0.014
20	2	<b>0.13</b>	0.16	0.81	<b>2.17</b>	<b>2.18</b>	1.00	<b>12.82</b>	<b>12.78</b>	1.00	0.010	0.013
20	4	<b>0.07</b>	0.08	0.88	<b>1.06</b>	<b>1.08</b>	0.98	<b>6.63</b>	<b>6.63</b>	1.00	0.011	0.012
20	8	<b>0.03</b>	0.03	1.00	<b>0.59</b>	<b>0.58</b>	1.02	<b>3.52</b>	<b>3.51</b>	1.00	0.009	0.009
21	1	0.51	<b>0.44</b>	1.16	<b>7.38</b>	<b>7.28</b>	1.01	<b>41.43</b>	<b>41.32</b>	1.01	0.012	0.011
21	2	0.27	<b>0.25</b>	1.08	<b>3.37</b>	<b>3.36</b>	1.00	<b>21.10</b>	<b>21.10</b>	1.00	0.013	0.012
21	4	<b>0.12</b>	0.14	0.86	<b>1.79</b>	<b>1.81</b>	0.99	<b>11.09</b>	<b>11.11</b>	1.00	0.011	0.013
21	8	0.07	<b>0.06</b>	1.17	<b>1.00</b>	<b>1.00</b>	1.00	<b>5.83</b>	<b>5.82</b>	1.00	0.012	0.010
22	1	<b>0.40</b>	0.43	0.93	<b>3.15</b>	<b>3.18</b>	0.99	<b>13.67</b>	<b>13.63</b>	1.01	0.029	0.032
22	2	0.21	<b>0.20</b>	1.05	<b>1.58</b>	<b>1.58</b>	1.00	<b>7.17</b>	<b>7.13</b>	1.01	0.029	0.028
22	4	<b>0.10</b>	0.11	0.91	<b>0.75</b>	<b>0.75</b>	1.00	<b>3.78</b>	<b>3.79</b>	1.00	0.027	0.029
22	8	<b>0.06</b>	0.06	1.00	<b>0.40</b>	<b>0.40</b>	1.00	<b>2.08</b>	<b>2.07</b>	1.00	0.029	0.029
23	1	<b>0.04</b>	0.05	0.80	<b>1.53</b>	<b>1.52</b>	1.01	<b>12.80</b>	<b>12.93</b>	0.99	0.003	0.004
23	2	<b>0.02</b>	0.04	0.50	<b>0.82</b>	<b>0.83</b>	0.99	<b>6.77</b>	<b>6.81</b>	0.99	0.003	0.006
23	4	<b>0.02</b>	<b>0.02</b>	1.00	<b>0.37</b>	<b>0.37</b>	1.00	<b>3.46</b>	<b>3.48</b>	0.99	0.006	0.006
23	8	<b>0.01</b>	<b>0.01</b>	0.99	<b>0.20</b>	0.21	0.95	<b>1.82</b>	<b>1.83</b>	0.99	0.006	0.006

Table 2: Average time (seconds) per MPI process spent doing FFT calculations (*fft\_3d*), doing SPME calculations (*spme*) and executing DL\_POLY (*total*). Timings for both the mixed precision and double precision versions are given and for 1,2,4 and 8 MPI processes.



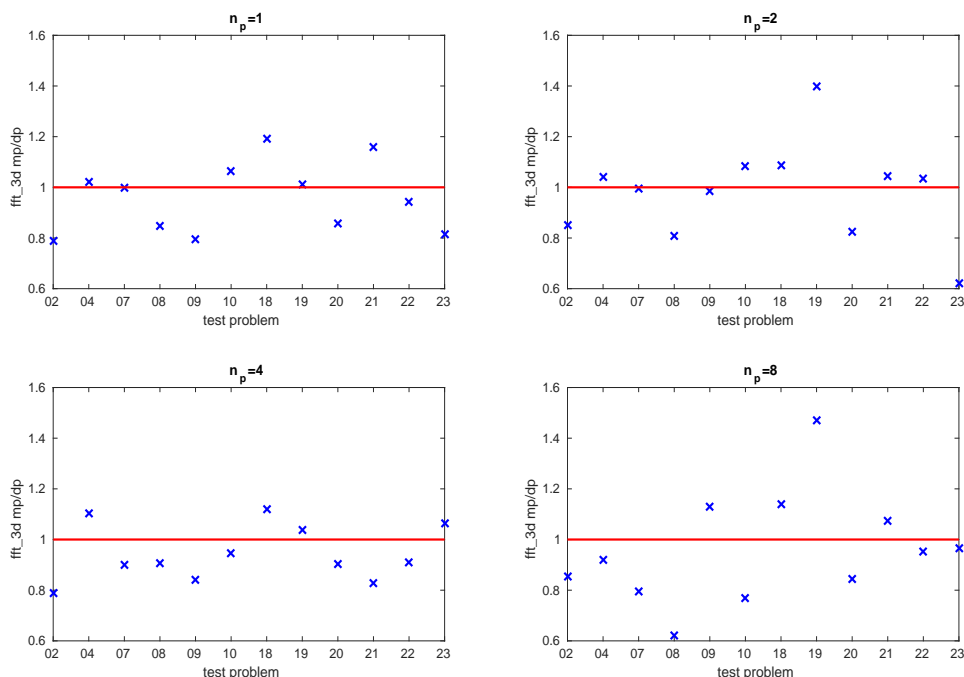


Figure 1: Ratio of the time spent performing the FFT calculations using mixed precision versus double precision.

In Figure 1, we look at the ratio of mean time spent by a process on the DaFT calculation when using mixed precision versus double precision. We observe that, for some problems using one MPI process, the time dropped by upto 20%. For one test problem, Test 08 with 8 MPI processes, there was a 40% drop in execution time when mixed precision was used instead of double precision. However, there were also problems where the reverse happened, for example, Test 19. For test problems 02 and 07, the DaFT kernel accounted for roughly 45% and 35% of the time in performing the Ewald SPME calculations. For the rest of the test problems, the DaFT kernel accounted for roughly 5-15% of the Ewald SMPE time. As a result, the savings or gains in using mixed precision for the SPME component are far more moderate when we consider the mean execution time across an MPI process, see Figure 2. In fact, the maximum savings is close to 10% (Test 02) and the largest increase is 3%.

The Ewald SPME calculations are just part of the overall force calculations being performed within a run of DL\_POLY and, hence, the time spent performing the FFT computations is generally a small fraction of the total time. Therefore, the execution time gains/losses for the FFT computation will have an even smaller effect on the total execution time for DL\_POLY. For test problems 02 and 07, roughly 10% of the time is spent on each MPI process is the FFT calculation; for five of the problems (04, 10, 20, 21 and 23), less than 2% of the time is spent in the FFT calculations.

It should be noted that, whilst these standard DL\_POLY tests are not generally FFT dominant, for applications where the FFT is dominant, any savings in the FFT execution could be significant. Additionally, the size of these test problems do not replicate the typical nature of the molecular dynamics problems run using DL\_POLY.

## 5.2 Enlarged DL\_POLY test cases

In this section, we enlarge some of the test problems to enable us to examine the effect of communication between processes for the 3D FFT component and give a better idea of the mixed precision behaviour. To enlarge the test problems, we use DL\_POLY's `nfold` facility. For test problem 02, we had to additionally alter the CONTROL

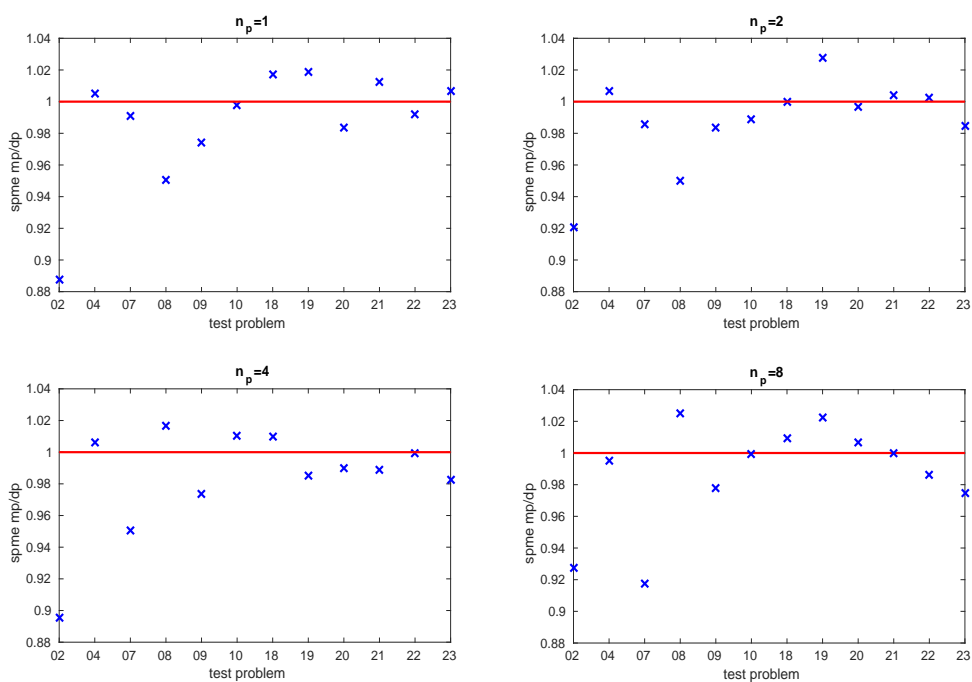


Figure 2: Ratio of the time spent performing the Ewald SPME calculations using mixed precision versus double precision.

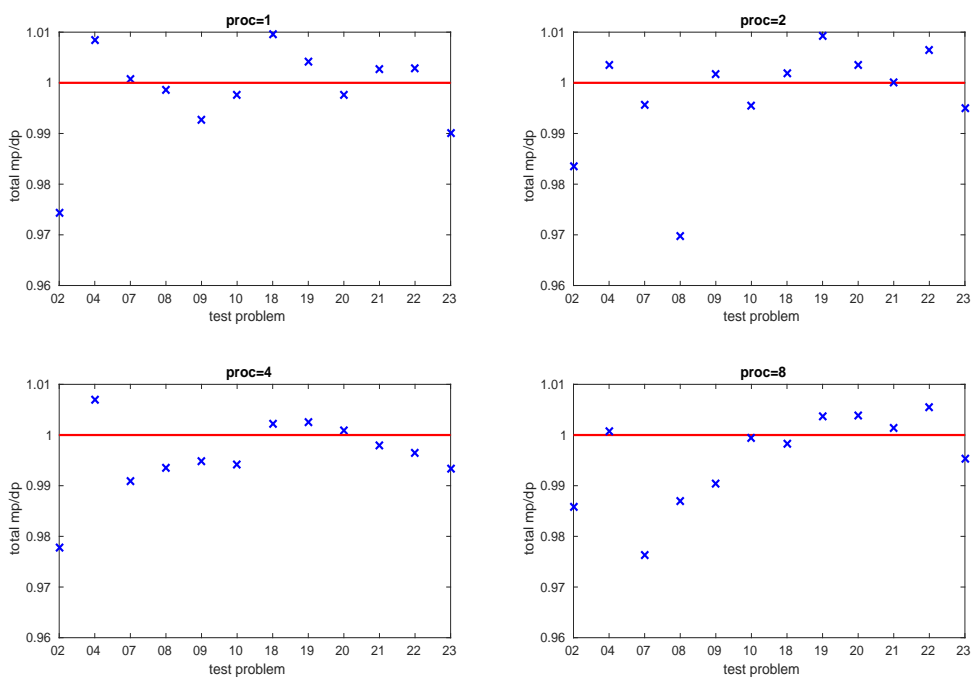


Figure 3: Ratio of the time spent running DL\_POLY using mixed precision versus double precision.

file by appropriately scaling the maximum number of  $k$ -vector indices in the Ewald specification line. We append “b” to the test problem number for problems that were expanded by doubling the  $x$ ,  $y$  and  $z$  dimensions; “c” is appended to the problem number for problems where the  $x$ ,  $y$  and  $z$  dimensions are increased by a factor of 5. The original problem is replicated 8 (125) times and shifted appropriately to fill the enlarged space. With the enlarged systems, it is necessary to increase the number of MPI processes used to solve the problems.

Problem	$n_p$	<i>fft_3d</i>			<i>spme</i>			<i>total</i>			<i>fft_3d/total</i>	
		mixed	double	ratio	mixed	double	ratio	mixed	double	ratio	mixed	double
02b	8	<b>5.10</b>	6.77	0.75	<b>9.99</b>	11.76	0.85	<b>49.93</b>	51.69	0.97	0.102	0.132
02b	16	<b>2.75</b>	4.13	0.67	<b>5.48</b>	6.93	0.79	<b>26.36</b>	27.77	0.95	0.104	0.149
02b	24	<b>2.63</b>	5.05	0.52	<b>4.66</b>	7.17	0.65	<b>19.31</b>	21.77	0.89	0.136	0.232
02b	32	<b>1.38</b>	2.30	0.60	<b>2.81</b>	3.79	0.74	<b>13.78</b>	14.77	0.93	0.100	0.156
04b	8	<b>1.00</b>	1.25	0.80	<b>6.72</b>	7.00	0.96	<b>60.37</b>	<b>60.58</b>	1.00	0.017	0.021
04b	16	<b>0.49</b>	0.84	0.58	<b>3.45</b>	3.81	0.91	<b>32.49</b>	<b>32.72</b>	0.99	0.015	0.026
04b	24	<b>0.46</b>	1.06	0.43	<b>2.55</b>	3.16	0.81	<b>23.85</b>	<b>24.32</b>	0.98	0.019	0.043
04b	32	<b>0.24</b>	0.44	0.55	<b>1.80</b>	2.02	0.89	<b>16.76</b>	<b>16.89</b>	0.99	0.014	0.026
07b	8	<b>4.10</b>	4.72	0.87	<b>11.22</b>	11.94	0.94	<b>40.04</b>	<b>40.76</b>	0.98	0.102	0.116
07b	16	<b>2.51</b>	3.71	0.68	<b>6.46</b>	7.72	0.84	<b>23.3</b>	24.42	0.95	0.108	0.152
07b	24	<b>2.72</b>	4.84	0.56	<b>5.63</b>	7.85	0.72	<b>18.89</b>	21.04	0.90	0.144	0.230
07b	32	<b>1.40</b>	2.31	0.61	<b>3.55</b>	4.51	0.79	<b>12.74</b>	13.66	0.93	0.110	0.169
10b	8	<b>0.42</b>	0.47	0.88	<b>2.76</b>	2.83	0.98	<b>24.2</b>	<b>24.4</b>	0.99	0.017	0.019
10b	16	<b>0.23</b>	0.30	0.79	<b>1.50</b>	1.57	0.95	<b>13.1</b>	<b>13.2</b>	0.99	0.018	0.023
10b	24	<b>0.16</b>	0.25	0.62	<b>1.03</b>	1.13	0.91	<b>8.77</b>	<b>8.84</b>	0.99	0.018	0.028
10b	32	<b>0.11</b>	0.13	0.87	<b>0.72</b>	0.74	0.98	<b>6.65</b>	<b>6.76</b>	0.98	0.017	0.019
18b	8	<b>1.11</b>	1.15	0.96	<b>8.64</b>	<b>8.73</b>	0.99	<b>38.98</b>	<b>39.12</b>	1.00	0.029	0.029
18b	16	<b>0.58</b>	0.65	0.90	<b>4.66</b>	<b>4.75</b>	0.98	<b>20.76</b>	<b>20.88</b>	0.99	0.028	0.031
18b	24	<b>0.44</b>	0.63	0.69	<b>3.25</b>	3.44	0.94	<b>14.95</b>	<b>15.15</b>	0.99	0.029	0.042
18b	32	<b>0.31</b>	<b>0.34</b>	0.91	<b>2.29</b>	<b>2.32</b>	0.99	<b>11.12</b>	<b>11.13</b>	1.00	0.028	0.031
19b	8	<b>0.31</b>	0.33	<b>0.95</b>	<b>2.48</b>	<b>2.46</b>	1.01	<b>11.82</b>	<b>11.60</b>	1.02	0.026	0.028
19b	16	<b>0.17</b>	0.18	0.92	<b>1.29</b>	<b>1.30</b>	0.99	<b>6.19</b>	<b>6.16</b>	1.00	0.026	0.029
19b	24	<b>0.12</b>	0.18	0.66	<b>0.88</b>	0.95	0.92	<b>4.24</b>	4.3466	0.97	0.028	0.041
19b	32	<b>0.09</b>	0.10	0.86	<b>0.65</b>	0.67	0.98	<b>3.17</b>	<b>3.19</b>	0.99	0.028	0.031
02c	48	<b>27.66</b>	46.10	0.60	<b>45.00</b>	63.67	0.71	<b>174.11</b>	192.76	0.90	0.159	0.239
02c	96	<b>14.62</b>	29.07	0.50	<b>23.05</b>	38.65	0.60	<b>91.84</b>	106.61	0.86	0.159	0.273
02c	144	<b>9.83</b>	17.22	0.57	<b>15.00</b>	22.58	0.66	<b>62.68</b>	70.22	0.89	0.157	0.245
04c	96	<b>3.06</b>	5.32	0.58	<b>11.07</b>	13.42	0.82	<b>107.20</b>	109.20	0.98	0.029	0.049
04c	144	<b>2.37</b>	4.27	0.56	<b>7.76</b>	9.71	0.80	<b>72.34</b>	73.84	0.98	0.033	0.058
04c	192	<b>1.34</b>	2.60	0.52	<b>5.40</b>	6.70	0.81	<b>53.36</b>	54.42	0.98	0.025	0.048
18c	48	<b>2.16</b>	4.01	0.54	<b>8.76</b>	10.69	0.82	<b>56.48</b>	58.48	0.97	0.038	0.069
18c	96	<b>0.87</b>	1.66	0.53	<b>4.17</b>	5.02	0.83	<b>27.36</b>	28.22	0.97	0.032	0.059
18c	144	<b>0.57</b>	1.38	0.41	<b>2.81</b>	3.64	0.77	<b>17.07</b>	17.89	0.95	0.033	0.077

Table 3: Average time (seconds) per MPI process spent doing FFT calculations (*fft\_3d*), doing SPME calculations (*spme*) and executing DL\_POLY (*total*). Timings for both the mixed precision and double precision versions are given and for 1,2,4 and 8 MPI processes.

In Table 3, we list the (mean) time spent by each MPI process in the various kernels (listed in Section 5) when either the mixed precision or double precision versions of the DaFT are used within the DL\_POLY run. The times in bold are those that are at most 2% larger than the best time. As the size of the test problem is scaled up by a factor 8 and then a factor of 125 (with respect to the original test problem), we observe that the mixed precision value of *fft\_3d* approaches 0.5 times that when double precision is used. As noted in Section 5.1, the affect on the overall execution time depends on how dominant the FFT component is in the overall DL\_POLY run. We note that, increasing the problem size has the effect of increasing the dominance of the FFT component. For the “b” problems, switching from 24 to 32 processes leads to a decrease in the dominance but this corresponds to

moving from using one to using two nodes. This need for off-node communication results in differing behaviours for different components within DL\_POLY and, hence, different ratios. For the “c” test problem where the FFT is most dominant (Test 02c), the (wall-clock) execution time is decreased by 14% when mixed precision is used in the SPME calculation compared to using double precision; for the test problem where it is least dominant (Test 04c), there is a 2% decrease in execution: for large molecular dynamics problems taking many hours to run, this can still be a useful gain.

## 6 Conclusions

We conclude that a simple modification to the DAFT algorithm allowed us to turn it into a mixed precision method and halved the fast Fourier transform calculation in the larger test problems. There was little or no loss in accuracy when comparing the output from DL\_POLY. For some of our test problems where the fast Fourier transform is a dominant component within DL\_POLY, we were able to reduce the overall execution time of DL\_POLY by 14%. The exact nature of the problem being solved will dictate the possible gains from using this mixed precision approach.

**Acknowledgements** The author would like to thank Alin Marin Elena and Michael Seaton for the discussions we had with regards using mixed precision within DL\_POLY, and for proposing this work.

## References

- [1] See <http://community.hartree.stfc.ac.uk/wiki/site/admin/resources.html>.
- [2] I. BUSH, I. TODOROV, AND W. SMITH, *A DAFT DL\_POLY distributed memory adaptation of the Smoothed Particle Mesh Ewald method*, Computer Physics Communications, 175 (2006), pp. 323 – 329.
- [3] S. QI, S. SHI, AND X. WANG, *Mixed Precision Method for GPU-based FFT*, 2013 IEEE 16th International Conference on Computational Science and Engineering, 00 (2011), pp. 580–586.
- [4] W. SMITH AND I. T. TODOROV, *DL\_POLY\_3: the CCP5 national UK code for molecular-dynamics simulations*, Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences, 362 (2004), pp. 1835–1852.
- [5] H. S. THORNE, *Using mixed precision within DL\_POLY’s force and energy evaluations: short-range interactions*, Tech. Rep. RAL-TR-2018-004, 2018.
- [6] H. S. THORNE, L. MASON, AND A. D. TAYLOR, *Mixed precision: Is it the holy grail for software efficiency?*, in 2nd Conference of Research Software Engineers (RSE17), Manchester, U.K., September 7-8 2017. <http://purl.org/net/epubs/work/34840371>.
- [7] I. TODOROV AND W. SMITH, *The DL\_POLY\_4 User Manual (version 4.08)*. March 2016.