Eur. Phys. J. E (2011) **34**: 110
DOI 10.1140/epje/i2011-11110-0

THE EUROPEAN PHYSICAL JOURNAL E

Regular Article

Efficiency of pseudo-spectral algorithms with Anderson mixing for the SCFT of periodic block-copolymer phases

P. Stasiak[a] and M.W. Matsen[b]

School of Mathematical and Physical Sciences, University of Reading, Whiteknights, Reading RG6 6AX, UK

Received: 23 June 2011 and Received in final form 18 August 2011
Published online: 11 October 2011 –

Abstract. This study examines the numerical accuracy, computational cost, and memory requirements of self-consistent field theory (SCFT) calculations when the diffusion equations are solved with various pseudo-spectral methods and the mean-field equations are iterated with Anderson mixing. The different methods are tested on the triply periodic gyroid and spherical phases of a diblock-copolymer melt over a range of intermediate segregations. Anderson mixing is found to be somewhat less effective than when combined with the full-spectral method, but it nevertheless functions admirably well provided that a large number of histories is used. Of the different pseudo-spectral algorithms, the 4th-order one of Ranjan, Qin and Morse performs best, although not quite as efficiently as the full-spectral method.

1 Introduction

The self-consistent field theory (SCFT) for structured polymers, which approximates molecular interactions by self-consistent mean fields [1–3], has proven to be an extraordinarily successful theory for predicting equilibrium behavior. However, the SCFT for block copolymers was not given much attention at the time of its introduction by Helfand [4] in 1975, due to the computational demands of applying it to multi-dimensional morphologies. The challenge was in solving a modified diffusion equation (MDE) for the propagators used to evaluate the polymer concentrations for a given set of mean fields. This difficulty was overcome with the introduction of a spectral algorithm in 1994 by Matsen and Schick [5], making it possible to treat triply periodic morphologies of arbitrary complexity, such as the gyroid phase.

Later, in 2002, Rasmussen and Kalosakas [6] introduced an alternative pseudo-spectral method for solving the MDE. This was, in turn, followed by two higher-order pseudo-spectral schemes introduced in 2006 by Cochran, Garcia-Cervera and Fredrickson [7] and in 2008 by Ranjan, Qin and Morse [8]. The pseudo-spectral methods are often assumed to be superior, because their computational cost scales as $M \log M$, where M is the number of wave vectors or equivalently number of grid points used to represent spatial functions. This can be contrasted by the cost of the full-spectral method, which scales as M^3. Cochran *et al.* [7] demonstrated the efficiency of their pseudo-spectral method by evaluating the position of the gyroid channel up to a segregation of $\chi N = 100$, far surpassing the earlier spectral calculation up to $\chi N = 40$ [9].

Shortly after, however, Matsen [10] demonstrated that the spectral method can reach similar levels of segregation, and is in fact orders of magnitude faster at the more experimentally relevant segregations (*i.e.*, $\chi N \approx 20$–40). The spectral method compensates for the M^3 scaling by the fact that it can readily incorporate the symmetry of the morphology to reduce M (by a factor of 96 for the gyroid phase). Furthermore, the spectral method achieved superior accuracy with far fewer planewaves (*e.g.*, $M = 76800$ as opposed to $M = 262144$ at $\chi N = 40$).

Another important consideration is the computational cost of adjusting the mean fields to satisfy the self-consistency conditions. For diblock copolymers, where there are two independent mean fields, this involves solving $2M$ simultaneous non-linear equations by some iterative scheme. The spectral calculation of Matsen [10] was able to obtain the self-consistent mean fields for a well-segregated gyroid morphology in less than a 100 iterations, whereas the pseudo-spectral calculation of Cochran *et al.* [7] required around 1000 iterations. However, the two studies performed the iterations differently; Cochran *et al.* used a semi-implicit scheme (SIS) [11] while Matsen used Anderson mixing [12]. Although the SIS proved to be efficient in the initial trials on the lameller phase [11], its performance dropped off considerably for the gyroid morphology. Similarly, the initial tests of Anderson mixing [13] were also impressive, but they too were limited to the less challenging lamellar and spherical morphologies at relatively weak segregations. Subsequent applications of Anderson mixing have encountered convergence diffi-

[a] e-mail: p.stasiak@reading.ac.uk
[b] e-mail: m.w.matsen@reading.ac.uk

culties [14,15], and the same could very well be true for a well-segregated gyroid phase. The fact that Anderson mixing works well in conjunction with the spectral method is no guarantee that the same will be true with the pseudo-spectral method. By including the symmetry, the spectral method greatly reduces the dimensionality of the problem, making it easier to locate the self-consistent solutions.

Although there is now a range of effective methods for solving SCFT, at the moment our judgment regarding the relative performance of the algorithms is highly clouded by speculation. Future progress will require benchmarking the algorithms in order to truly know which combination is best. So far, the only direct comparison of the pseudo-spectral and spectral algorithms has been by Jiang *et al.* [16]. Interestingly, they found that the SIS performed better with the pseudo-spectral algorithm than the spectral method, even though the latter had far fewer field equations to solve. Apart from that, however, they did not say anything about the relative performance of the two approaches. Furthermore, they only considered the original pseudo-spectral method of Rasmussen and Kalosakas.

Here we perform a far more thorough comparison, examining all three pseudo-spectral methods. Our study also tests out Anderson mixing to see if it performs better than the SIS used by Cochran *et al.* We finish by performing direct head-to-head comparisons of the pseudo-spectral and full-spectral algorithms, applied to the triply periodic gyroid and spherical phases of intermediately segregated diblock-copolymer melts.

2 SCFT and numerical algorithms

This section briefly describes self-consistent field theory (SCFT) and the numerical algorithms we use to solve it.

2.1 SCFT

Our application of SCFT [1] is for a melt of n AB diblock copolymers, each with an A-block of N_A segments joined to a B-block of N_B segments giving a total polymerization of $N \equiv N_\mathrm{A} + N_\mathrm{B}$ and an A-segment composition of $f = N_\mathrm{A}/N$. We assume equal statistical lengths, a, for the two segment types and a uniform segment concentration, ρ_0, such that the total volume of the melt is $V = nN/\rho_0$. The interaction between A and B segments is described by the usual Flory-Huggins χ parameter.

The first step in SCFT is to calculate the concentrations, $\phi_\mathrm{A}(\mathbf{r})$ and $\phi_\mathrm{B}(\mathbf{r})$, of the A and B segments with the molecular interactions replaced by external fields, $w_\mathrm{A}(\mathbf{r})$ and $w_\mathrm{B}(\mathbf{r})$, acting on the A and B segments, respectively. They are evaluated as

$$\phi_\mathrm{A}(\mathbf{r}) = \frac{V}{Q}\int_0^f q(\mathbf{r},s)q^\dagger(\mathbf{r},s)\mathrm{d}s, \tag{1}$$

$$\phi_\mathrm{B}(\mathbf{r}) = \frac{V}{Q}\int_f^1 q(\mathbf{r},s)q^\dagger(\mathbf{r},s)\mathrm{d}s, \tag{2}$$

where

$$Q = \int q(\mathbf{r},s)q^\dagger(\mathbf{r},s)\mathrm{d}\mathbf{r}. \tag{3}$$

The propagator, $q(\mathbf{r},s)$, satisfies the modified diffusion equation (MDE),

$$\frac{\partial}{\partial s}q(\mathbf{r},s) = \frac{a^2N}{6}\nabla^2 q(\mathbf{r},s) - w_\gamma(\mathbf{r})q(\mathbf{r},s), \tag{4}$$

with the initial value $q(\mathbf{r},0) = 1$. When the chain contour variable exceeds $s = f$, the field switches from $w_\mathrm{A}(\mathbf{r})$ to $w_\mathrm{B}(\mathbf{r})$. The backward propagator, $q^\dagger(\mathbf{r},s)$, satisfies the same MDE but with one side multiplied by -1, and it is integrated from $s = 1$ to 0 starting from $q^\dagger(\mathbf{r},1) = 1$.

The next step is to adjust the fields to satisfy the self-consistency conditions,

$$w_\mathrm{A}(\mathbf{r}) = \chi N\phi_\mathrm{B}(\mathbf{r}) + \xi(\mathbf{r}), \tag{5}$$

$$w_\mathrm{B}(\mathbf{r}) = \chi N\phi_\mathrm{A}(\mathbf{r}) + \xi(\mathbf{r}), \tag{6}$$

where $\xi(\mathbf{r})$ is used to enforce the incompressibility condition, $\phi_\mathrm{A}(\mathbf{r}) + \phi_\mathrm{B}(\mathbf{r}) = 1$. Once the fields have been determined, the free energy is evaluated as

$$\frac{F}{nk_BT} = -\ln\left(\frac{Q}{V}\right) + \frac{1}{V}\int[\chi N\phi_\mathrm{A}(\mathbf{r})\phi_\mathrm{B}(\mathbf{r}) - w_\mathrm{A}(\mathbf{r})\phi_\mathrm{A}(\mathbf{r}) - w_\mathrm{B}(\mathbf{r})\phi_\mathrm{B}(\mathbf{r})]\mathrm{d}\mathbf{r}. \tag{7}$$

2.2 Pseudo-spectral methods

The pseudo-spectral methods solve the MDE for the propagator at a set of discrete steps, s_i for $i = 0, 1, ..., N_s$, on a mesh of collocation points, $\mathbf{r}_j$ for $j = 1, 2, ..., M$. The original Rasmussen-Kalosakas (RK2) method steps the propagator forward in s using

$$q(\mathbf{r}, s_{i+1}) \approx \hat{\mathcal{R}}_{\Delta s}q(\mathbf{r}, s_i), \tag{8}$$

where $\Delta s = s_{i+1} - s_i$ and

$$\hat{\mathcal{R}}_t \equiv \exp\left(-\frac{t}{2}w_\gamma(\mathbf{r})\right)\exp\left(\frac{a^2Nt}{6}\nabla^2\right)\exp\left(-\frac{t}{2}w_\gamma(\mathbf{r})\right) \tag{9}$$

is a product of three operators. The application of the first operator on $q(\mathbf{r}, s_i)$ involves a multiplication at each mesh point, and thus its computational complexity is $\mathcal{O}(M)$. Then a fast-Fourier transform (FFT) is applied to the result, which has a complexity of $\mathcal{O}(M \log M)$. In Fourier space, the second operator becomes $\exp(-a^2N\Delta s|\mathbf{k}|^2/6)$ and thus its application again just involves one multiplication for each of the M reciprocal wave vectors, $\mathbf{k}$. Finally an inverse FFT is applied and the result is multiplied by the third operator to obtain $q(\mathbf{r}, s_{i+1})$. Note that we use slightly different step sizes, Δs, along the A and B blocks so that the junction point, $s = f$, falls on a grid point.

In the limit of small steps, the numerical error of the 2nd-order RK2 method decreases as $(\Delta s)^2$. Ranjan, Qin and Morse [8] extended this to a 4th-order method

(RQM4) using Richardson extrapolation. By comparing the estimation for a single step, $\hat{\mathcal{R}}_{\Delta s}q(\mathbf{r},s_i)$, to that of two half-sized steps, $\hat{\mathcal{R}}^2_{\Delta s/2}q(\mathbf{r},s_i)$, they obtain the improved estimation

$$q(\mathbf{r},s_{i+1}) \approx \left[\frac{4}{3}\hat{\mathcal{R}}^2_{\Delta s/2} - \frac{1}{3}\hat{\mathcal{R}}_{\Delta s}\right] q(\mathbf{r},s_i). \tag{10}$$

Cochran, Garcia-Cervera and Fredrickson [7] also proposed another 4th-order method (CGF4). By treating the Laplacian operator implicitly with a backward differentiation formula and the field operator explicitly using an Adams-Bashford formula, they obtain [17]

$$\begin{aligned}&\left[\frac{25}{12\Delta s} - \frac{a^2N}{6}\nabla^2\right] q_{i+1} \approx \\ &\frac{1}{\Delta s}\left[4q_i - 3q_{i-1} + \frac{4}{3}q_{i-2} - \frac{1}{4}q_{i-3}\right] \\ &-w_\gamma\left[4q_i - 6q_{i-1} + 4q_{i-2} - q_{i-3}\right],\end{aligned} \tag{11}$$

where, for brevity, q_i denotes $q(\mathbf{r},s_i)$. The right-hand side of the equation is evaluated and transformed to Fourier space, where the operator on the left-hand side becomes diagonal allowing $q(\mathbf{k},s_{j+1})$ to be computed by simple division. An inverse FFT then converts the solution back to real space. Again, the numerical complexity is dominated by the FFTs, but this time only one FFT pair is required. However, a disadvantage of the method is that the first three values of q_i must be obtained by a different algorithm. Here we compute them using the RQM4 method.

The contour integrations, eqs. (1) and (2), are performed by Simpson's method, which is also 4th-order accurate in Δs. The spatial integrations, eqs. (3) and (7), are approximated as simple sums over the mesh points, j, which in fact provides exponential accuracy for the case of periodic symmetry.

2.3 Anderson mixing

The self-consistency conditions for the fields are solved iteratively using a variation of the Anderson-mixing scheme first applied to SCFT by Thompson *et al.* [13]. The k-th iteration begins with the evaluation of new fields, $\bar{w}^{(k)}_{\mathrm{A},j}$ and $\bar{w}^{(k)}_{\mathrm{B},j}$, from eqs. (5) and (6) where $\phi^{(k)}_{\mathrm{A},j}$ and $\phi^{(k)}_{\mathrm{B},j}$ are evaluated from the old fields, $w^{(k)}_{\mathrm{A},j}$ and $w^{(k)}_{\mathrm{B},j}$, and $\xi_j = \frac{1}{2}(w^{(k)}_{\mathrm{A},j} + w^{(k)}_{\mathrm{B},j})$. Note that the SCFT is unaffected by additive constants to the fields, and therefore we adjust all the fields so that their spatial averages are zero. Next, we evaluate the deviations,

$$d^{(k)}_{\gamma,j} = \bar{w}^{(k)}_{\gamma,j} - w^{(k)}_{\gamma,j}, \tag{12}$$

for $\gamma = \mathrm{A}$ and B, and we define

$$\text{error} \equiv \left[\sum_{\gamma,j}\left(d^{(k)}_{\gamma,j}\right)^2 \Big/ \sum_{\gamma,j}\left(w^{(k)}_{\gamma,j}\right)^2\right]^{1/2} \tag{13}$$

as a measure of the numerical inaccuracy in the field equations (5) and (6).

If the error exceeds some tolerance, then improved estimates, $w^{(k+1)}_{\mathrm{A},j}$ and $w^{(k+1)}_{\mathrm{B},j}$, are obtained from the preceding n_r iterations. This is done by evaluating the symmetric matrix,

$$U_{mn} = \sum_{\gamma,j}(d^{(k)}_{\gamma,j} - d^{(k-m)}_{\gamma,j})(d^{(k)}_{\gamma,j} - d^{(k-n)}_{\gamma,j}), \tag{14}$$

and vector,

$$V_m = \sum_{\gamma,j}(d^{(k)}_{\gamma,j} - d^{(k-m)}_{\gamma,j})d^{(k)}_{\gamma,j}, \tag{15}$$

for $m,n = 1,...,n_r$. From these, we calculate the coefficients,

$$C_n = \sum_m (U^{-1})_{nm}V_m, \tag{16}$$

and the previous histories are combined as

$$W^{(k)}_{\gamma,j} = w^{(k)}_{\gamma,j} + \sum_{n=1}^{n_r} C_n(w^{(k-n)}_{\gamma,j} - w^{(k)}_{\gamma,j}), \tag{17}$$

$$D^{(k)}_{\gamma,j} = d^{(k)}_{\gamma,j} + \sum_{n=1}^{n_r} C_n(d^{(k-n)}_{\gamma,j} - d^{(k)}_{\gamma,j}). \tag{18}$$

Finally, the new fields are obtained from

$$w^{(k+1)}_{\gamma,j} = W^{(k)}_{\gamma,j} + \lambda D^{(k)}_{\gamma,j}, \tag{19}$$

where $0 < \lambda \leq 1$ is a mixing parameter [18,19]. The iterations converge more quickly as $\lambda \to 1$, but they also become unstable for large λ if there is an insufficient number of histories.

Thompson *et al.* [13] generated their initial histories using simple mixing [20], and then switched to Anderson mixing with $\lambda = 1$. This is fine when only a small number of histories is required, but it is not so great for large n_r. Not only are the simple mixing iterations costly and relatively ineffective at improving the solution, they do not provide a particularly good set of histories. To maximize the benefits of Anderson mixing, we apply it from the very start by gradually ramping up the mixing parameter while the number of histories builds up to some maximum, n_h. This is done by setting $\lambda = 1.0 - 0.9^k$ and $n_r = \min(k-1, n_h)$.

3 Results

Here we investigate the performance of the original 2nd-order pseudo-spectral algorithm of Rasmussen and Kalosakas (RK2) [6] and the 4th-order algorithms of Cochran, Garcia-Cervera and Fredrickson (CGF4) [7] and of Ranjan, Qin and Morse (RQM4) [8] for solving the modified diffusion equation (4). Then we examine the effectiveness of Anderson mixing for solving the self-consistent field equations (5) and (6). We conclude by performing the first ever direct head-to-head comparisons with the full-spectral method.

Table 1. Equilibrium domain size, D, and free energy, F, of the gyroid and spherical phases at several points in the phase diagram.

Gyroid phase			
χN	f	D	F/nk_BT
20	0.36	3.88925	3.933754694901
30	0.34	4.27332	4.716730826140
40	0.33	4.54110	5.278334267792
Spherical phase			
χN	f	D	F/nk_BT
20	0.24	1.89009	3.554220947205
30	0.19	1.90470	4.247976365310
40	0.17	1.94618	4.796511447944

3.1 Pseudo-spectral methods

First we start by examining how the accuracy of the pseudo-spectral calculations is affected by the finite number of steps along the chain contour, N_s, and of spatial mesh points, M. For the cubic unit cells of the gyroid and spherical phases, $M = L^3$, where L is the number of grid points in each spatial direction. Our focus is on the free energy, F, since this is the quantity involved in the calculation of phase diagrams. To determine its inaccuracy, $|\Delta F|$, we have computed some ultra-precise solutions of the SCFT using the spectral method with up to 5000 basis functions (see table 1).

Figure 1 shows the computational error in the free energy of the gyroid phase as a function of N_s for different values of L. (Note that the field equations (5) and (6) are solved to an ultra-high accuracy so as not to significantly affect ΔF.) The error in F initially decreases with N_s and then plateaus once it becomes dominated by the finite spatial resolution associated with L. As expected, the initial decrease scales as N_s^{-2} for the 2nd-order RK2 method and as N_s^{-4} for the 4th-order CGF4 and RQM4 methods. In some cases, the error changes sign producing a rapid drop and rise in $|\Delta F|$. Of course, we are only aware of this because of our *a priori* knowledge of the accurate free energy. In general, we must wait for the free energy to plateau before being confident that the error associated with N_s is small.

It is already clear from fig. 1 that the 4th-order methods are superior to the 2nd-order method, and therefore we omit the latter from our subsequent study. It may also appear that the RQM4 method is substantially more accurate than the CGF4 one, but we must remember that it requires three sets of FFTs for each step in s. For a fairer comparison, fig. 2 plots $|\Delta F|$ *versus* νN_s, where $\nu = 3$ for RQM4 and $\nu = 1$ for CGF4. This time, $|\Delta F|$ is plotted for three values of χN with L chosen large enough to prevent any significant contribution due to the finite spatial resolution. Compared this way, the CGF4 method seems to provide slightly higher accuracy, but this is not universal. The situation is reversed for the spherical phase (not shown).

Now we isolate the error due to the spatial discretization, L, by selecting large values of N_s. Figure 3a plots the resulting $|\Delta F|$ as a function of L for different values

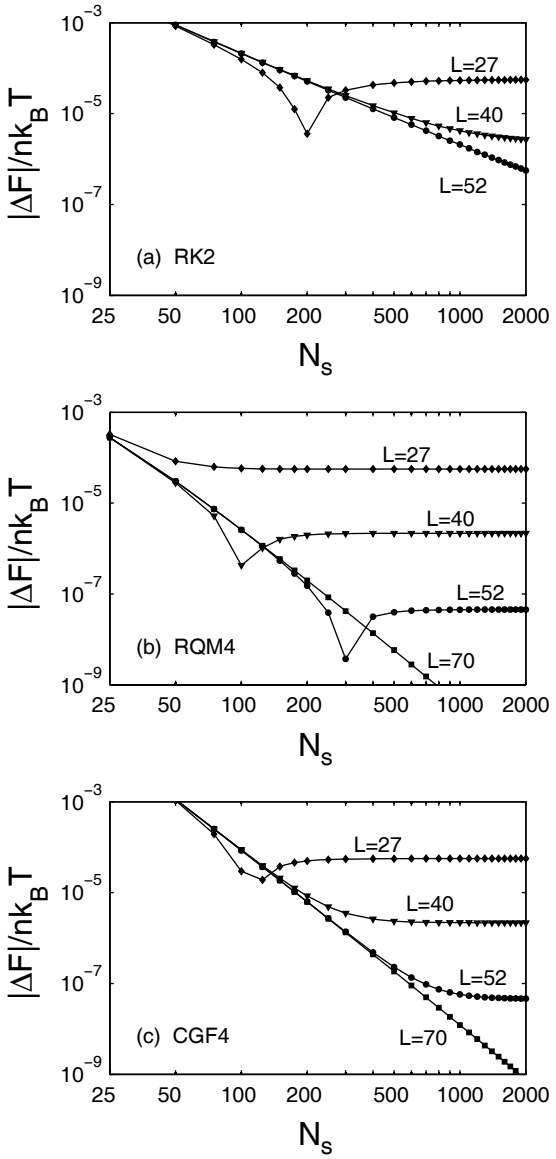

Fig. 1. Inaccuracy in the free energy, $|\Delta F|$, of the gyroid phase at $\chi N = 20$ resulting from the (a) RK2, (b) RQM4 and (c) CGF4 pseudo-spectral methods plotted as a function of chain-contour steps, N_s, for different numbers of spatial mesh points, $M = L^3$.

of χN. In this limit, all the pseudo-spectral methods produce equivalent results. The dashed lines are shown so as to emphasis the exponential decay characteristic of specular accuracy. In addition to the decay, the error jumps every time L is a multiple of 4, upwards for $\chi N = 20$ and 30 and downwards for $\chi N = 40$. The error for the spherical phase (not shown) also exhibits a similar behavior, but with a periodicity of 2.

Figure 3b shows analogous results for the spectral method, where $|\Delta F|$ is plotted as a function of the number of basis functions, M. To make the comparison equivalent in terms of the number of planewaves, the horizontal axis is labeled by $(96M)^{1/3}$, where the factor of 96 accounts the $Ia3d$ symmetry built into the spectral basis functions

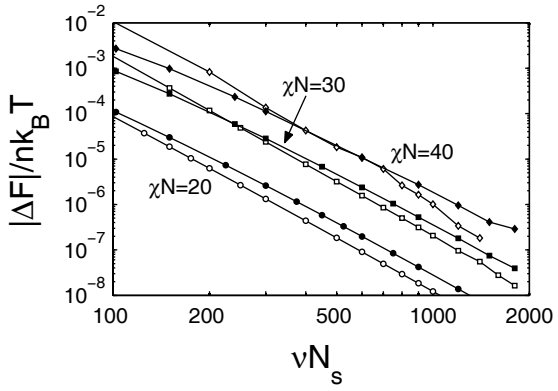

Fig. 2. Inaccuracy $|\Delta F|$ resulting from the RQM4 (solid symbols, $\nu = 3$) and the CGF4 (open symbols, $\nu = 1$) pseudo-spectral methods plotted as a function of νN_s in the large-L limit for various levels of segregation, χN.

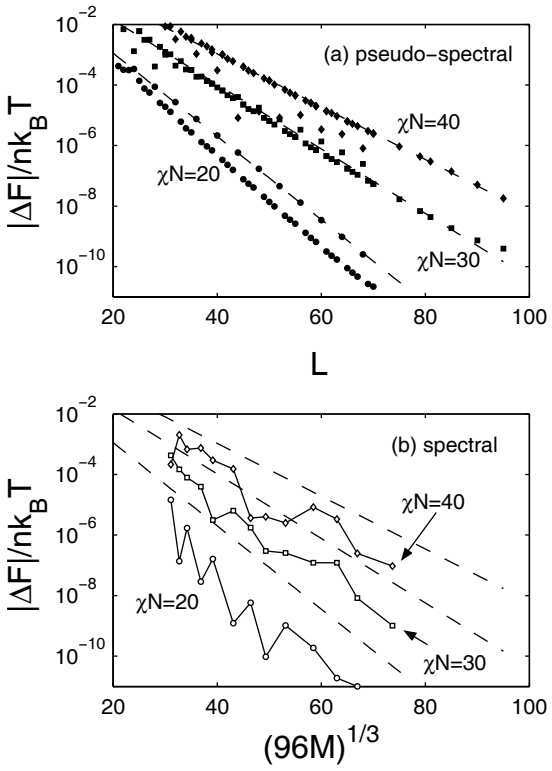

Fig. 3. Inaccuracy $|\Delta F|$ resulting from (a) the pseudo-spectral methods in the large-N_s limit plotted in terms of mesh points, L, and (b) the spectral method plotted in terms of basis functions, M. The dashed lines are included for comparison purposes and to demonstrate the exponential convergence.

for the gyroid phase. To aid the comparison, we also repeat the dashed lines from fig. 3a. In this case, the error is rather erratic, but it still displays a general exponential decay. It is also smaller, implying that the spectral method requires fewer planewaves than the pseudo-spectral methods to achieve equivalent accuracy. There could be many reasons for this, but one contributing factor is the way the planewave expansion is truncated (*i.e.*, the cutoff in reciprocal space is a sphere for the spectral method and a cube for the pseudo-spectral method). Nevertheless, the reduction is not nearly as substantial as suspected by ref. [10]; as we will show later, the accuracy of the pseudo-spectral calculations of Cochran *et al.* was limited by their choice of N_s rather than their finite L.

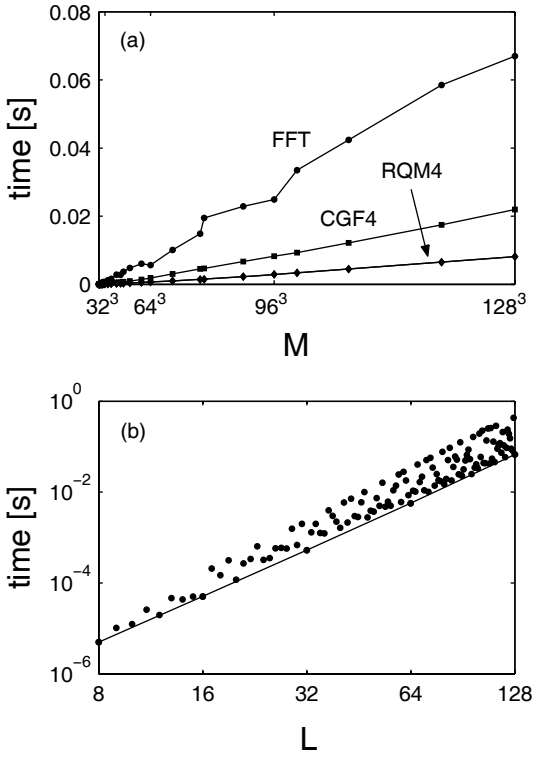

Fig. 4. (a) CPU time to perform a FFT and its inverse (circles) and the remaining time for a CGF4 step (squares) or a third of a RQM4 step (diamonds) plotted as a function of spatial mesh points, $M = L^3$, where L only contains factors of 2, 3 and 5. (b) Logarithmic plot of CPU time to perform a pair of FFTs for arbitrary L. The line denotes the maximum efficiency, which occurs when $L = 2^n$.

Next we investigate the computational cost of solving the MDE with the 4th-order pseudo-spectral methods. Figure 4 plots the CPU time taken to perform a single Δs step divided by ν ($\nu = 1$ for CGF4 and $\nu = 3$ for RQM4), so that the comparison is done on a per FFT basis, as was the case in fig. 2 where we compared numerical accuracies. We separate out the dominant cost of performing a FFT and its inverse, which is common to both algorithms. The remain calculations amount to about 10% and 25% of the total computational cost for the RQM4 and CGF4 methods, respectively. Considering that the numerical accuracy of the two methods was approximately equivalent when compared in this way, we conclude that the RQM4 method is slightly more efficient.

The Fourier transforms were performed using the freely available FFTW subroutines, version 3.2.2 [21]. Contrary to what one might expect, the computational cost of the FFTs is an erratic function of L, as illustrated in fig. 4b. The efficiency is best for $L = 2^n$ (solid curve),

and becomes poorer when L includes more prime factors, particularly large ones. For this reason, we restrict L to factors of 2, 3 and 5, in which case the time for the FFTs increases monotonically with L, apart from a slight decrease between $L = 2^2 \cdot 3 \cdot 5 = 60$ and $L = 2^8 = 64$ (see fig. 4a).

The memory requirements of the pseudo-spectral methods are dominated by the need to store $q(\mathbf{r}, s)$, which involves $N_s M$ double precision numbers (8 bytes each). Thus, the CGF4 method will generally require 3 times the memory of the RQM4 method on account of the fact that it needs N_s to be 3 times larger to achieve similar accuracy. Nevertheless, even for quite large values of $N_s = 600$ and $L = 60$, this works out to be about 1GB, which is not too onerous for modern computers. Note that there is no need to store $q^\dagger(\mathbf{r}, s)$, provided that one calculates the segment concentrations while stepping back though the values of s.

3.2 Anderson mixing

The performance of Anderson mixing is demonstrated on the gyroid phase in fig. 5 by plotting the error in the self-consistent field conditions, eq. (13), after each iteration, k. The convergence slows with increasing segregation, χN, and is sensitive to the method used to solve the MDE. In particular, Anderson mixing converges faster with the spectral method (dashed curves) than the two pseudo-spectral methods (solid lines), but this might be expected given that it has two orders of magnitude fewer equations to solve on account of the fact that it uses the symmetry. In any case, Anderson mixing performs admirably well in conjunction with the pseudo-spectral methods.

The first application of Anderson mixing by Thompson *et al.* [13] was kept to a small number of histories (*e.g.*, $n_h = 3$) so as not to slow down the calculation. That was fine for relatively weak segregations, but the well-segregated gyroid phase requires far more histories. The results in fig. 5 were obtained with a limit of $n_h = 50$ histories. Although the spectral method performs similarly well with only $n_h = 20$, the pseudo-spectral methods need the additional histories. Figure 6 demonstrates how n_h affects the convergence when the MDE is solved with the RQM4 method at $\chi N = 40$.

With so many histories, it is imperative to minimize the computational cost of the Anderson mixing. The most costly part is the evaluation of the matrix, U_{mn}, defined by eq. (14). The time required to evaluate it from scratch scales as $n_h^2 M$, which becomes prohibitive for large n_h. Fortunately, by saving and reusing the scalar products, $\sum_{\gamma,j} d_{\gamma,j}^{(k)} d_{\gamma,j}^{(k')}$, from previous iterations, the computational cost can be reduced to $\mathcal{O}(n_h M)$. With that, the time taken for the Anderson-mixing iterations remains a relatively small fraction of the overall SCFT calculation, even when n_h is as large as 100. In fact, the more significant issue becomes the storage of $w_{\gamma,j}^{(k)}$ and $d_{\gamma,j}^{(k)}$, which requires $4n_h M$ double precision numbers. Even still, this is no more demanding than storing the propagator provided that $n_h \leq N_s/4$, which was always the case for us.

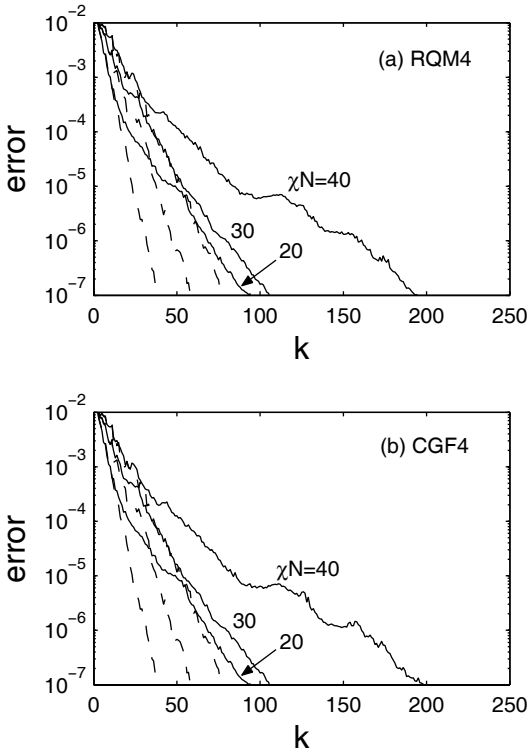

Fig. 5. Convergence of the field equations as a function of the number of Anderson-mixing iterations, k, with n_h=50. The solid curves are obtained with the (a) RQM4 and (b) CGF4 methods whereas the dashed curves are for the spectral method.

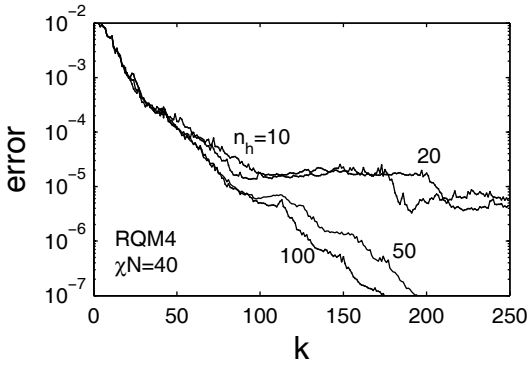

Fig. 6. Convergence of the Anderson-mixing iterations with different limits on the number of histories, n_h, obtained with the RQM4 method.

3.3 Comparison of methods

Our comparisons between the two 4th-order pseudo-spectral methods and the full-spectral method begin with the gyroid phase at the three points listed in table 1. In order to demonstrate the relative improvement in performance of the pseudo-spectral methods with increasing segregation, the algorithms are also compared at $\chi N = 50$ and $f = 0.32$. We aim for an accuracy of $|\Delta F|/nk_BT \lesssim 10^{-4}$, which is enough to calculate phase boundaries to

an accuracy of $|\Delta f| \lesssim 10^{-4}$. For triply periodic phases, the computational cost of increasing L far outweighs that of increasing N_s, and thus it is prudent to be a bit generous with N_s. Therefore, we first select N_s to give $|\Delta F|/nk_BT \lesssim 10^{-5}$ referring to fig. 2, and then we select L to give $|\Delta F|/nk_BT \lesssim 10^{-4}$ referring to fig. 3a. For the spectral method, we only need to concern ourselves with selecting M, for which we refer to fig. 3b. In all cases, the fields are iterated to an error $< 10^{-5}$.

Table 2 shows the number of Anderson-mixing iterations, k, required to obtain self-consistent solutions at $(f, \chi N)$ starting from neighboring solutions at $(f + 0.004, \chi N)$. The table also lists the resulting time and memory requirements in each case. All three methods perform admirably well, but nevertheless the CGF4 method is outperformed by the RQM4 method, which in turn is outperformed by the spectral method, apart from the one exception at $\chi N = 50$. Table 3 shows similar comparisons for the spherical phase at the three points listed in table 1 plus an extra point at $\chi N = 50$ and $f = 0.15$; in these cases, the Anderson mixing was started from solutions at $(f + 0.002, \chi N)$. This time the spectral method performs best on all accounts, followed by the RQM4 pseudo-spectral method. We do not concern ourselves with doubly periodic or lamellar phases, because their computational demands are negligible.

These calculations were all performed with a fixed domain size, D. In general, however, D needs to be adjusted so as to minimize the free energy, F. While the SIS is able to do this in conjunction with the field iterations [7], the same option is not available with Anderson mixing. Nevertheless, a single Newton-Raphson iteration is usually all it takes to adjust the domain size to an adequate accuracy, which merely doubles the computational cost of the calculation [10].

Table 2. Time and memory required to solve SCFT for the gyroid morphology using the pseudo-spectral (RQM4 and CGF4) and the spectral methods. L, N_s and M were chosen large enough to obtain an accuracy of $|\Delta F|/nk_BT \lesssim 10^{-4}$ and the number of Anderson-mixing iterations, k, were sufficient to solve the field equations to an error $< 10^{-5}$.

RQM4					
χN	L	N_s	k	time	memory
20	27	66	49	16 s	44 MB
30	40	140	55	106 s	180 MB
40	54	200	84	639 s	510 MB
50	72	300	119	3061 s	1500 MB
CGF4					
χN	L	N_s	k	time	memory
20	27	200	49	21 s	66 MB
30	40	420	55	143 s	320 MB
40	54	600	85	838 s	1000 MB
50	72	900	119	4194 s	3100 MB
Spectral					
χN	M		k	time	memory
20	150		24	1.0 s	2.8 MB
30	350		38	16 s	10 MB
40	850		52	433 s	45 MB
50	1600		62	3360 s	168 MB

Table 3. Analogous results to those of table 2, but for the spherical phase.

RQM4					
χN	L	N_s	k	time	memory
20	12	50	22	0.23 s	5.4 MB
30	18	110	31	4.4 s	17 MB
40	24	180	50	26 s	46 MB
50	32	270	60	84 s	125 MB
CGF4					
χN	L	N_s	k	time	memory
20	12	150	22	0.41 s	7.0 MB
30	18	330	32	5.9 s	26 MB
40	24	540	58	34 s	83 MB
50	32	810	60	144 s	261 MB
Spectral					
χN	M		k	time	memory
20	25		19	0.01 s	1.1 MB
30	50		26	0.07 s	1.3 MB
40	100		28	0.39 s	2.0 MB
50	160		32	1.53 s	3.2 MB

4 Discussion

The original 2nd-order pseudo-spectral method of Rasmussen and Kalosakas does not perform nearly as well as the 4th-order methods. Although both 4th-order methods are quite similar in performance, the RQM4 one has a number of distinct advantages. Even though the RQM4 method requires three FFTs per Δs step, it compensates for this by achieving similar accuracy with three times the step size. This combined with the fact that it involves fewer additional calculations gives the RQM4 method a noticeable speed advantage. The reduced number of contour steps also results in substantially lower memory requirements. Furthermore, it is a one-step scheme, and thus it does not require another method to start it off like the CGF4 method does.

In general though, the spectral method proves to be best for periodic morphologies. In our experience, most applications of the spectral method require at most a few hundred basis functions, for which the calculation is on the order of seconds. The accurate solution of the gyroid phase at the relatively high segregation of $\chi N = 50$ requires an unusually high 1600 basis functions because of its large unit cell, and only then does the speed of the pseudo-spectral methods become comparable. Even still, the memory requirements remain much lower for the spectral method. The most important advantage of the spectral method, however, is that it performs the integration over s analytically, which eliminates one source of numerical inaccuracy. This also makes it easier to handle generalizations of the system, such as the addition of high-molecular-weight homopolymer or polydispersity [22].

Key to the efficiency of the spectral method is that it readily incorporates the symmetry of the block-copolymer phases. For the example of the gyroid phase, M would

be 2 orders of magnitude larger without the symmetry, which would make the calculation 6 orders of magnitude slower. Our pseudo-spectral calculations have taken in to account the fact that the gyroid and spherical phases are centro-symmetric (*i.e.*, spatial functions were expanded in terms of $\cos(\mathbf{k}_j \cdot \mathbf{r})$), but this only affects the speed of the calculation by a factor of 2. The main reason for including the symmetry was to avoid any spatial drift of the solution; however, if need be, there are other ways of preventing drift [23].

The performance of the CGF4 method in this study is considerably better than in the previous demonstration by Cochran *et al.* [7], where it took about 2000 CPU minutes [24] to solve the SCFT for the gyroid phase at $\chi N = 40$ as opposed to our 14 minutes. Not only that, their solution was less accurate as evident by an inaccuracy of $\Delta f = -0.0007$ in their gyroid/lamellar phase boundary [10]. This is because they used $L = 64$ and $N_s = 250$, which does not balance the numerical inaccuracies. By increasing the number of chain-contour steps to $N_s = 600$ we are able to gain an extra digit of accuracy, and then by decreasing the spatial mesh resolution to $L = 54$ we are able to save CPU time without compromising the accuracy. Admittedly, the time savings is not terribly significant, because the FFT operates more efficiently for $L = 2^6$ than for $L = 2 \cdot 3^3$. As one might imagine, it requires some effort to determine the best combination of L and N_s, particularly with the complicated efficiency of the FFT subroutines. In contrast, it is much easier to manage the numerical inaccuracy of the spectral method.

Our improvement in computational speed can be largely attributed to our implementation of Anderson mixing. Cochran *et al.* required about 1000 iterations to converge their fields at $\chi N = 40$, whereas we required less than 100. Although Anderson mixing did not perform as quickly as it did with the spectral method [10], we were still able to recover much of the performance by simply increasing the number of histories. Although the convergence is impressive for our periodic diblock-copolymer morphologies, we must remember that problems have arisen in some previous applications on other systems [14,15]. Given how incredibly well Anderson mixing can perform, there is good reason to investigate the convergence more closely in hopes of extending its applicability. Meanwhile, it is good to have alternative methods available, such as the SIS [11].

In order to make forward progress, it is important to start benchmarking the different competing algorithms rather than simply speculating on their relative performance by, for example, the way their computational complexity scales. Hopefully our results in tables 2 and 3 for the gyroid and spherical diblock-copolymer phases, respectively, will serve this purpose. Of course, future tests will be performed on faster processors. Indeed, the 2.66 GHz Intel Xeon x5650 processor used in this study is already more than two times faster than the 2.4 GHz AMD Opteron 2216 processor used in the previous study of the spectral method [10]. Nevertheless, one can easily account for the relative speed improvement of a new processor by simply checking how the time for two $M = 128^3$ FFTs [21] compares to the 0.067s taken on our current processor (see fig. 4).

5 Conclusions

This study investigated the computational demands of the SCFT for diblock-copolymer melts in the intermediate-segregation regime. First we compared three different pseudo-spectral algorithms for solving the MDE. The two 4th-order (CGF4 and RQM4) methods are similar in performance, and both are much better than the original 2nd-order (RK2) method. Optimum performance requires sensible choices for the spatial and chain-contour mesh resolutions (*i.e.*, L and N_s). For triply periodic phases, it is important to use a generous value for N_s. It is also advisable to use values of L composed of small prime factors, but not necessarily just factors of 2. Next we examined the Anderson-mixing iteration for solving the self-consistent field conditions. It performs well in conjunction with the pseudo-spectral methods provided that enough histories are used, although not as quite as well as with the full-spectral method.

Although two 4th-order pseudo-spectral methods perform similarly, lower memory requirements, a modest speed improvement, and the extra simplicity of a one-step scheme leads us to recommend the RQM4 method. The CPU time required to solve the SCFT for triply periodic morphologies (*e.g.*, gyroid and spherical) ranges from just seconds to minutes over the intermediate-segregation regime. Still the spectral method is generally faster and requires less memory, particularly for the weaker segregations. This combined with the fact it has one less source of numerical error, since the integration over the chain contour is done analytically, leads us to conclude that the spectral method is the best algorithm for solving periodically ordered morphologies. On the other hand, the pseudo-spectral method is far superior for problems involving little or no symmetry. We hope that this study will help researchers choose the appropriate algorithm for their needs, and that it will also provide a useful benchmark for testing future developments.

We are grateful to Eric Cochran, Glenn Fredrickson, Dave Morse and Russell Thompson for helpful discussions. This work was supported by the EPSRC (EP/F068425/1 and EP/G026203/1).

References

1. M.W. Matsen, J. Phys.: Condens. Matter **14**, R21 (2002).
2. G.H. Fredrickson, *The Equilibrium Theory of Inhomogeneous Polymers* (Oxford University Press, New York, 2006).
3. M.W. Matsen, in *Soft Matter*, Vol. 1: *Polymer Melts and Mixtures*, edited by G. Gompper, M. Schick (Wiley-VCH, Weinheim, 2006).

4. E. Helfand, J. Chem. Phys. **62**, 737 (1975).
5. M.W. Matsen, M. Schick, Phys. Rev. Lett. **72**, 2660 (1994).
6. K.Ø. Rasmussen, G. Kalosakas, J. Polym. Sci., Part B **40**, 1777 (2002).
7. E.W. Cochran, C.J. Garcia-Cervera, G.H. Fredrickson, Macromolecules **39**, 2449 (2006).
8. A. Ranjan, J. Qin, D.C. Morse, Macromolecules **41**, 942 (2008).
9. M.W. Matsen, F.S. Bates, Macromolecules **29**, 1091 (1996).
10. M.W. Matsen, Eur. Phys. J. E **30**, 361 (2009).
11. H.D. Ceniceros, G.H. Fredrickson, Multiscale Model. Simul. **2**, 452 (2004).
12. D.G. Anderson, J. Assoc. Comput. Mach. **12**, 547 (1965).
13. R.B. Thompson, K.Ø. Rasmussen, T. Lookman, J. Chem. Phys. **120**, 31 (2004).
14. R.B. Thompson, Phys. Rev. E **74**, 041501 (2006).
15. J.U. Kim, M.W. Matsen, Soft Matter **5**, 2889 (2009).
16. K. Jiang, Y. Huang, P. Zhang, J. Comput. Phys. **229**, 7796 (2010).
17. E.W. Cochran, C.J. Garcia-Cervera, G.H. Fredrickson, Macromolecules **39**, 4264 (2006).
18. V. Eyert, J. Comput. Phys. **124**, 271 (1996).
19. F. Schmid, Phys. Rev. E **55**, 5774 (1997).
20. Simple mixing equates to Anderson mixing without any histories (*i.e.*, $n_r = 0$), which requires the mixing parameter to be relatively small (*e.g.*, $\lambda \approx 0.1$).
21. M. Frigo, S.G. Johnson, *The Design and Implementation of FFTW3*, Proc. IEEE **93**, 216 (2005). The FFT subroutine, version 3.2.2, was obtained from `www.fftw.org`.
22. M.W. Matsen, Phys. Rev. Lett. **99**, 148304 (2007).
23. B. Vorselaars, J.U. Kim, T.L. Chantawansri, G.H. Fredrickson, M.W. Matsen, Soft Matter **7**, 5128 (2011).
24. E.W. Cochran, private communication.