Acta Crystallographica Section A
Foundations of Crystallography
ISSN 0108-7673

Received 23 July 2008
Accepted 22 January 2009

Symmetry-modified charge flipping

Alexander Eggeman, Thomas White and Paul Midgley*

Department of Materials Science and Metallurgy, University of Cambridge, Pembroke Street, Cambridge CB2 3QZ, UK. Correspondence e-mail: pam33@cam.ac.uk

The charge-flipping algorithm has been adapted to allow symmetry constraints to be included during the solution of structures from diffraction data. Rather than impose symmetry at the start of the algorithm, which is known to cause the process to stagnate, it is shown that the algorithm must be allowed to find an intermediate stable solution first. Although care is needed when using this modified algorithm, the improvement in the fidelity of the structure solution is considerable.

1. Introduction

Charge flipping is a relatively new technique, and is gaining in popularity and importance for solving the phase problem in X-ray, neutron and electron diffraction studies (Oszlányi & Sütő, 2004). Unlike conventional direct methods, where phase-invariant relationships are used as a basis for phasing reflections (Karle & Hauptman, 1950; Giacovazzo, 1998), the charge-flipping algorithm uses only a single basic principle: the requirement that the final structure solution, be it an electron density for X-rays or a projected crystal potential in the case of electron diffraction, is sparse and positive at all points. From this starting point, successful solutions have been found for a wide variety of crystal structures; in almost all cases these are from single-crystal X-ray diffraction data (Oszlányi & Sütő, 2005; Baerlocher *et al.*, 2007; Lidin, 2008). Charge flipping has been successfully used in diffractive imaging studies using electron diffraction data (Wu *et al.*, 2005); however, the difficulties presented by dynamical diffraction and in accurate recording of intensities in electron diffraction mean that it is sometimes difficult to apply it to the solution of complex structures.

One important feature of the algorithm is the requirement that at least one external constraint is applied in real space to ensure a reasonable solution to the diffraction data. For the standard algorithm this is applied through a threshold value δ; all pixels in the real-space density with an intensity below this value have their sign flipped during each iteration. The value of δ is important because structure factors that contribute to strong positive peaks (with an intensity greater than δ) in the real-space density will be unaffected by the charge-flipping process and so are effectively removed from the active part of the iteration. For the cases described here we take δ to be 5% of the maximum intensity.

Other constraints have been proposed as ways of improving the effectiveness of the charge-flipping algorithm. The *Superflip* program (Palatinus & Chapuis, 2007) describes several useful methods for improving the structure solution. Coelho (2007) has combined a charge-flipping algorithm with the tangent formula and has shown that this minimizes the likelihood of the real-space intensity being condensed into single peaks (the so-called 'uranium atom' solution), although the intensities of the reflections used in the algorithm have to be carefully controlled. Zhou & Harris (2008) have proposed a 'residue-based' charge-flipping technique which uses a residual of the calculated structure factors at each iteration of the algorithm. This reduces the number of phases that are altered in the flipping stage of the iteration and so greatly improves the convergence properties of the algorithm. In the present article, the use of known symmetry relationships is proposed as an additional constraint. This should be of particular benefit for electron diffraction data, in which typically only a two-dimensional net of reflections is available, but where the projection symmetry is often known.

2. The symmetry-modified algorithm

The application of symmetry to the recovered structure-factor moduli, $|\mathbf{F_h}|$, is a conventional first stage in preparing data for phase retrieval, and applying similar symmetry constraints to the recovered phases is simply using all of the available data to maximize the likelihood of obtaining a successful structure. In electron diffraction experiments, determination of the crystallographic space group can be achieved in a relatively straightforward fashion using, for example, convergent-beam electron diffraction (CBED) (Buxton *et al.*, 1976) independently of any subsequent structure solution. Charge flipping is normally undertaken on data sets with only Friedel symmetry imposed ($\mathbf{F}^*_{-\mathbf{h}} = \mathbf{F_h}$) because this allows the algorithm to move freely in the N-dimensional solution space (where N is the number of reflections to be phased) and so find a global minimum in solution negativity. Enforcing symmetry has been reported to reduce the likelihood of a successful solution (Oszlányi & Sütő, 2004) by restricting the algorithm to a small subset of solution space (as defined by the initial randomly applied phases) which may not contain a minimum corresponding to a strong solution. The present work supports the suggestion in more recent studies (Oszlányi & Sütő, 2008) that

the timing of the application of symmetry is an important factor in determining the success of the algorithm.

A schematic diagram of our symmetry-modified charge-flipping algorithm is shown in Fig. 1. The application of symmetry is performed in the reciprocal-space part of the Fourier cycle. After the Fourier transform of the real-space density, the solution exists as a list of structure factors. We considered it to be easier to apply symmetry to the phases in this list rather than to the individual pixels in the real-space density, although there is no reason why a real-space phase imposition should not be possible. The imposition of centro-

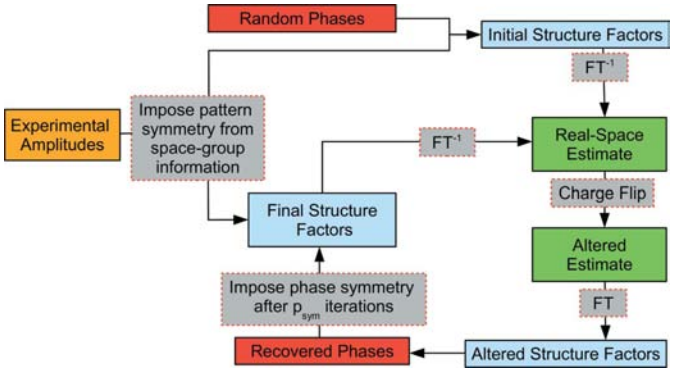

Figure 1
Schematic diagram of the new charge-flipping algorithm, including the application of symmetry constraints after p_{sym} iterations.

symmetry forces the structure-factor phases to assume the closest value of 0 or π. This defines the origin of the unit cell and so allows the position of the symmetry elements in the cell to be assigned consistently. The charge-flipping protocol was developed using the FFTW library for calculating Fourier transforms (Frigo & Johnson, 2005) and the potential maps were sampled to a scale of 500 pixels per nanometre to ensure that the full resolution of the map was achieved with respect to the input data; higher-resolution terms recovered by the Fourier cycle were not included in the calculation.

Notwithstanding the possibility of applying this algorithm to all types of diffraction data, we believe that this new version may be of considerable benefit for electron diffraction studies. Therefore, in this paper, we illustrate the improved performance of the charge-flipping algorithm with applied symmetry constraints by focusing on its application to simulated and experimental electron diffraction patterns.

2.1. Preliminary structure solutions

Theoretical zonal kinematical structure factors for 300 kV electron diffraction were calculated using the *JEMS* software package (Stadelmann, 2004) for a series of complex oxide crystals. Of these, several could be solved satisfactorily using the charge-flipping algorithm in $p1$. An example is kinoite, $Ca_2Cu_2(H_2O)_2[Si_3O_{10}]$, shown in Fig. 2, in which (*a*) the ideal and (*b*) the recovered potential maps are illustrated.

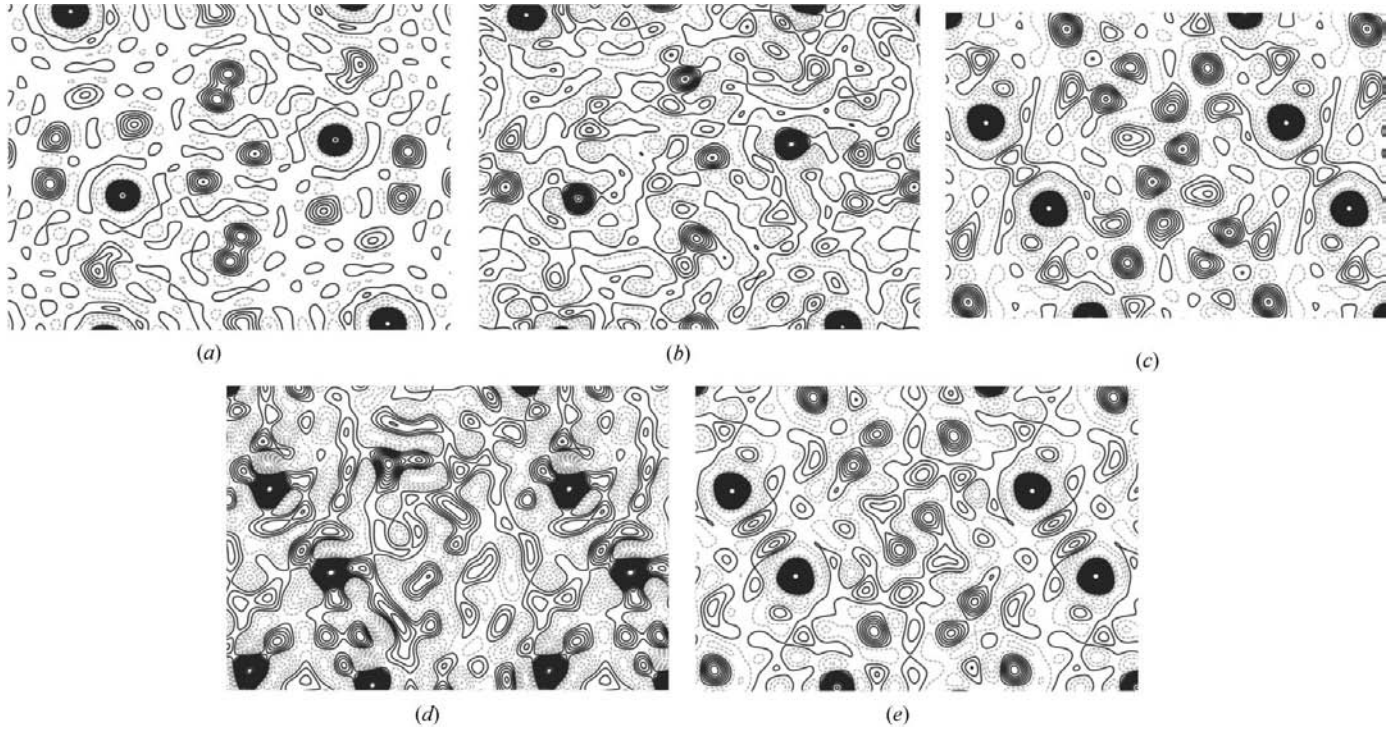

Figure 2
Potential maps showing (*a*) the ideal structure of kinoite and (*b*) the charge-flipping solution using $p1$ symmetry for kinoite after 50 iterations. (*c*) The ideal structure of epidote. (*d*) The charge-flipping solution using $p1$ symmetry for epidote after 50 iterations. (*e*) The solution for epidote using the new algorithm with $p2$ symmetry applied after 20 iterations; the total number of iterations was 50. All reflections in the range $-15 \leq h, k \leq 15$ were used for all solutions.

In contrast, the solution of the structure of epidote, $Ca_2(Fe,Al)_3(SiO_4)_3(OH)$, is poor and the weaker atom peaks are not recovered correctly. The ideal and recovered potential maps for epidote are shown in Figs. 2(*c*) and (*d*). When *p*2 symmetry was applied after 20 iterations of the Fourier cycle, a considerably more complete structure was recovered, as shown in Fig. 2(*e*). Both kinoite and epidote are monoclinic with space group $P2_1/m$. It was not clear from these examples why some structures are readily solved in *p*1 while others require higher-symmetry constraints. To investigate the algorithm further it was used on a more complex test structure.

2.2. Erbium pyrogermanate – a test case for the new algorithm

Erbium pyrogermanate ($Er_2Ge_2O_7$), EGO, is tetragonal with space group $P4_12_12$, $a = b = 0.6778$ nm and $c = 1.234$ nm. Using atomic coordinates determined by neutron diffraction studies (Smolin, 1970), a set of kinematical electron structure factors for $-15 \leq h, k \leq 15$ was calculated using *JEMS*. A potential map projected along [001] was generated from these structure factors and is shown in Fig. 3. The map shows a ring of eight strong peaks corresponding to the erbium-atom columns and four 'doublet' peaks, each comprising two germanium-atom columns. The oxygen-atom columns add only a small perturbation to this map and are not considered significant compared to the heavy-atom peaks. The displacement parameter B was set to zero for the the purpose of this study. The presence of a small number of negative-intensity contours is caused by the finite number of reflections introduced into the inverse Fourier transform used to construct the potential map.

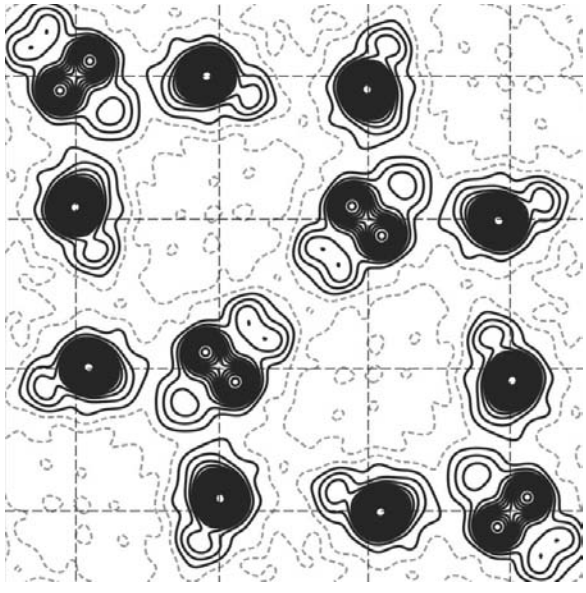

Figure 3
Projected potential map for the [001] zone axis of EGO. A 4 × 4 grid has been added to highlight the relationship to the strong 400 and 040 reflections in the kinematical diffraction pattern (shown in Fig. 4). All reflections in the range $-15 \leq h, k \leq 15$ were included.

The three different elements in EGO and its high space-group symmetry ensure that the diffraction pattern will contain a wide range of reflection intensities; because of the nature of the charge-flipping algorithm this serves as a useful test of the algorithm. When a pixel in the real-space solution has a large positive value, the reflections that contribute to it are considered to be correctly phased with respect to one another and these reflections are ignored by the active part of the iteration (the charge flipping). For a simple system where the diffracted beams are either very strong or very weak this will lead to a rapid convergence of all of the strong reflections to a phase relationship that will return a structure even if some of the strong reflections are phased incorrectly. In the case of EGO, which has a wide range of reflection intensities, the likelihood of medium- and weak-intensity reflections contributing to a high-intensity pixel is very small unless the phasing is exactly correct; this makes the structure sensitive to the effectiveness of the algorithm.

A second interesting feature of EGO comes from the diffraction by the constituent elements. If we exclude the weak contribution from the oxygen atoms, then the diffraction patterns from each of the two metal-atom sublattices show notable features. Kinematical diffraction patterns for the erbium and germanium sublattices for $-8 \leq h, k \leq 8$ were simulated using the *mulslice* program in the *TEMSIM* software package (Kirkland, 1998) and are shown in Figs. 4(*a*) and (*b*), respectively. The theoretical kinematical diffraction pattern for the complete structure (from the structure factors used in Fig. 3) is shown in Fig. 4(*c*). All the simulated diffraction patterns in this paper were generated as CBED patterns (with a convergence angle chosen such that the discs are just short of touching) to make a visual comparison of the spot intensities as easy as possible.

Comparison of Figs. 4(*a*) and (*b*) suggests that the erbium sublattice dominates the strong reflections (for example the 400, 440 and 110 reflections) and contributes significantly to only a few other low-order reflections. The germanium sublattice, however, contributes significantly to all the lower-order reflections in the diffraction pattern; significantly, this includes several strong reflections such as 210 and 320. These reflections are likely to have a strong influence on the position of peaks within the structure and so phasing these correctly with respect to both sublattices will be vital to the recovery of a good structure solution.

3. Structure solutions using charge flipping

Fig. 5 shows a typical solution obtained by the charge-flipping algorithm. This solution was found using the structure-factor moduli used in Fig. 3 with *p*1 plane-group symmetry and with Friedel symmetry applied such that $\mathbf{F}^*_{hk} = \mathbf{F}_{\overline{hk}}$. The charge-flipping algorithm allows the structure-factor phases to take any value in the initial stage of the algorithm. The threshold value δ was set to 5% of the peak pixel intensity in the density; this value was found through experimentation to constrain the positivity requirement of the algorithm and also to allow sufficient freedom to explore the solution space and find good

solutions. For this reason, and for consistency between the results, this threshold value was used for all of the solutions described in this article.

The solution in Fig. 5 bears little resemblance to the ideal projected potential map in Fig. 3. At best, it seems that the positions of the double germanium-atom columns have been revealed by the algorithm, even if the 'doublets' are not resolved as two peaks, and some of the erbium-atom columns appear to be present. However, there are additional peaks in the potential that do not correspond to atom columns from the ideal structure, and all the strong peaks appear to lie on a 4 × 4 grid defined by the strong 400 and 040 reflections. It seems likely that the majority of the strong reflections have been phased to the germanium-atom positions (which do satisfy the structure-factor moduli well) and this incorrect phasing has led to a solution that satisfies the remaining intensity distribution, despite being a relatively poor match to the true structure.

Despite the weakness of this solution, the algorithm has found a minimum in solution space, which suggests that some of the phase relationships between reflections must have relevance to the correct solution. Furthermore, the [001] projection of a $P4_12_12$ structure has $p4gm$ plane-group symmetry, which establishes several significant phase relationships between groups of reflections (shown in Fig. 6) that should improve the final solution. So, rather than impose symmetry on the phases from the start, the new algorithm was developed to enable symmetry to be imposed on the solution after a regular number of iterations (denoted by $p_{\rm sym}$). Examples of the solutions recovered from the algorithm with $p_{\rm sym}$ = 2, 20 and 40 iterations are shown in Fig. 7.

When $p_{\rm sym}$ is large (Figs. 7*b* and *c*), the algorithm produced a solution which contained all the heavy-atom peaks and agreed well with the ideal structure shown in Fig. 3. The structure found using a small value of $p_{\rm sym}$ contains strong peaks that match the position of the germanium-atom doublets but the rest of the structure was not recovered. Much like the $p1$ structure in Fig. 5, the algorithm seems to have phased all the reflections to the 'germanium solution' rather than to the real structure.

The behaviour of the algorithm can be described using the concept of the 'feasible-set approach' described by Marks *et al.* (1999). Here a solution space is created in which every real-space potential map can exist; a subset of this space (S_1) contains all of the solutions that satisfy the structure-factor

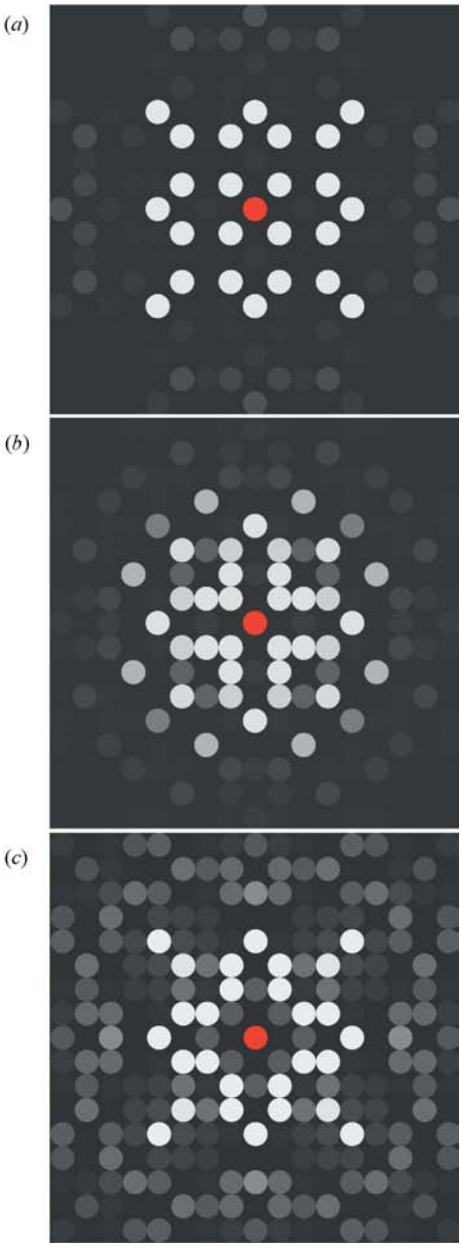

Figure 4
Calculated kinematical diffraction patterns for the [001] zone axis of EGO, showing the contributions from (*a*) the erbium-atom columns alone, (*b*) the germanium-atom columns alone and (*c*) the complete EGO structure.

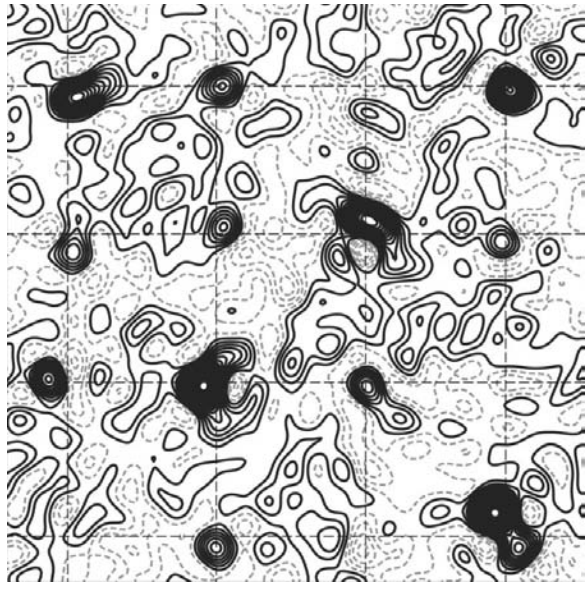

Figure 5
Output structure from the charge-flipping algorithm for theoretical EGO diffracted intensities using $p1$ symmetry after 80 iterations. Note that the peaks sit approximately at the corners of a 4 × 4 grid (*cf.* Fig. 3).

moduli $|F|$ input into the algorithm. A second subset (S_2) contains all of the solutions that satisfy the constraint applied by the algorithm (in this case, sparseness and a positive pixel intensity). A perfect set of structure-factor data for a crystal with a unique solution would therefore lead to a single contact between the sets, although limitations in the accuracy of the structure-factor data will result in two non-contacting sets. During each iteration of the algorithm the solution will move from the start point in S_1 to the nearest point in S_2 and then to the nearest point in S_1 as the different parts of the Fourier cycle are completed. In the conventional charge-flipping algorithm, this should converge towards a solution with the minimum possible negativity. This is shown schematically in Fig. 8(*a*). The application of symmetry to the algorithm forces an additional constraint and effectively reduces the accessible part of the solution set (S_1). This is shown in Fig. 8(*b*) and

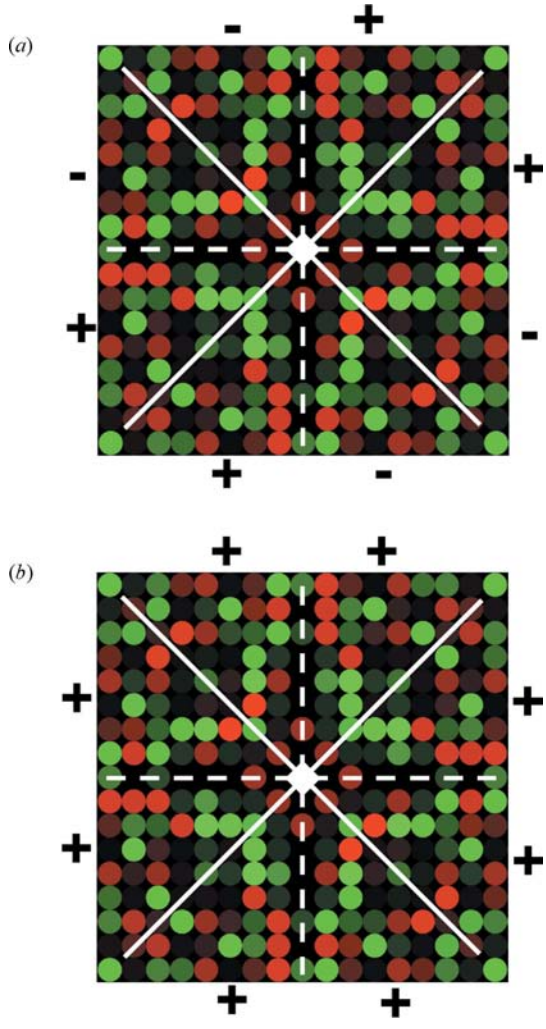

Figure 6
Phase relationships for reflections corresponding to the *p*4*gm* plane group. The dashed lines show the positions of glide planes and the solid lines show the positions of mirror planes in the reciprocal lattice; reflections in green have a phase of π while those in red have a phase of 0. In (*a*) for $h + k = 2n + 1$ the glide planes act as 'anti-mirrors' producing $p4'm'm$ symmetry, while in (*b*) for $h + k = 2n$ the glide planes act as mirrors, producing *p*4*mm* symmetry.

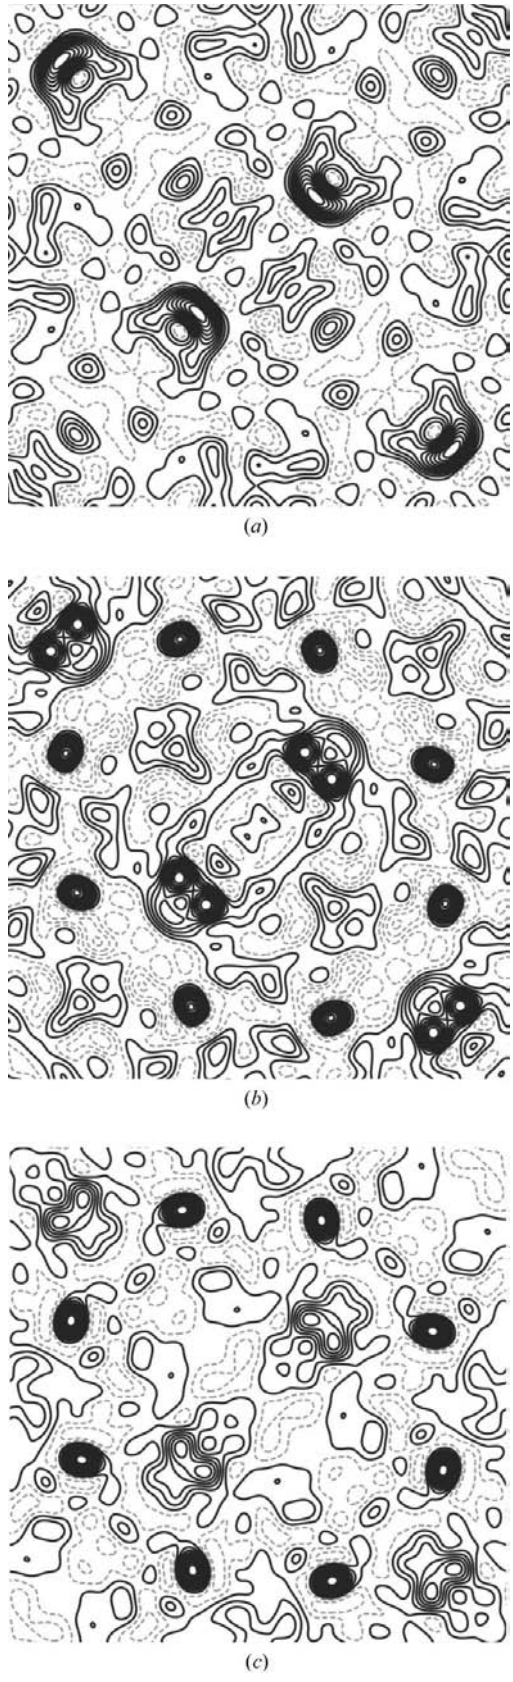

Figure 7
Solutions from the symmetry-applied charge-flipping algorithm for theoretical kinematical EGO data with p_{sym} of (*a*) 2, (*b*) 20 and (*c*) 40 iterations. All reflections in the range $-15 \leq h, k \leq 15$ were included in all solutions.

helps to explain why the process can stagnate if symmetry is applied from the outset of the algorithm. In this case, the solution becomes trapped in the first local minimum encountered in the symmetry-constrained solution set, while other potentially better minima are unexplored.

The application of symmetry after a number of iterations allows the algorithm to explore the negativity constraint more completely before the additional symmetry constraint alters the solution set. This is shown in Fig. 8(*c*). In this case, the solution is able to explore minima that satisfy the negativity constraint and then explore the symmetry-constrained set that offers the most improved negativity. In this sense, the algorithm is exploring the full set of solution space before the additional constraint is added in the expectation of an improved solution. Of course, once the symmetry is applied, the solution will become trapped in a local minimum in the constrained solution set, but this minimum should offer a better solution with respect to both negativity and symmetry constraints than would be found otherwise.

This behaviour can be seen by investigating the statistical entropy of the solution at each step of the algorithm. Statistical entropy can be considered as a measure of the 'peakiness' of a data set (Jaynes, 1957). A high entropy indicates a solution with well defined strong peaks rather than low-intensity noise and so is likely to indicate a solution that represents atomic potentials rather than Fourier ripples. The entropy, S, was calculated using

$$S = -\sum_{R} \rho_R \ln(\rho_R), \qquad (1)$$

where ρ_R is the normalized intensity of pixel R in the real-space solution and the sum is over all the pixels in the density. As an independent measure of the fidelity of the solution, the total intensity of all the negative pixels in the solution was also calculated after each iteration. These two quantities were measured for a range of $p_{\rm sym}$ from 2 to 40 and are shown in Figs. 9(*a*) and (*b*), respectively.

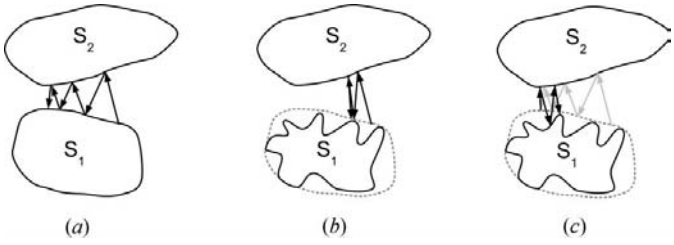

Figure 8
Schematic diagrams of the operation of the charge-flipping algorithm under different conditions: (*a*) $p1$ symmetry, (*b*) higher symmetry applied from the outset of the algorithm, (*c*) higher symmetry applied after a number of iterations. S_1 represents a set containing all solutions that satisfy the structure-factor data and S_2 represents all solutions that satisfy the positivity criterion. The black arrows indicate the migration of the algorithm from a random starting solution towards one that satisfies the constraints applied by the algorithm. In (*b*) and (*c*) the dotted lines indicate the extent of S_1 under $p1$ symmetry while the solid line represents the extent of the set under higher applied symmetry. The light grey arrows in (*c*) indicate the migration of the algorithm under $p1$ symmetry and the black arrows indicate the further development of the solution with the higher-symmetry constraint.

The evolution of entropy and negativity for $p_{\rm sym} = 2$ shows that when symmetry is applied very early in the algorithm, the solution settles rapidly into a stable minimum in the constrained solution space (*cf.* Fig. 8*b*). While this solution satisfies the criteria by locating a minimum in the local solution space, the entropy indicates that the structure is comprised predominantly of weak ripples rather than strong atomic peaks. This is evident in the solution (Fig. 7*a*), where there are only four peaks and a large amount of the unit cell is filled with weak background. This solution may represent a local minimum in negativity but may not be particularly strong when compared with other global solutions.

The entropy curve also shows a fundamental effect that must be considered: once symmetry has been applied to a solution there is no mathematical way to remove or alter the constraint. The action of the Fourier cycle will always preserve symmetry that exists in the solution; only numerical artefacts are able to introduce any variation. This is the reason why applying the symmetry constraint from the outset of the

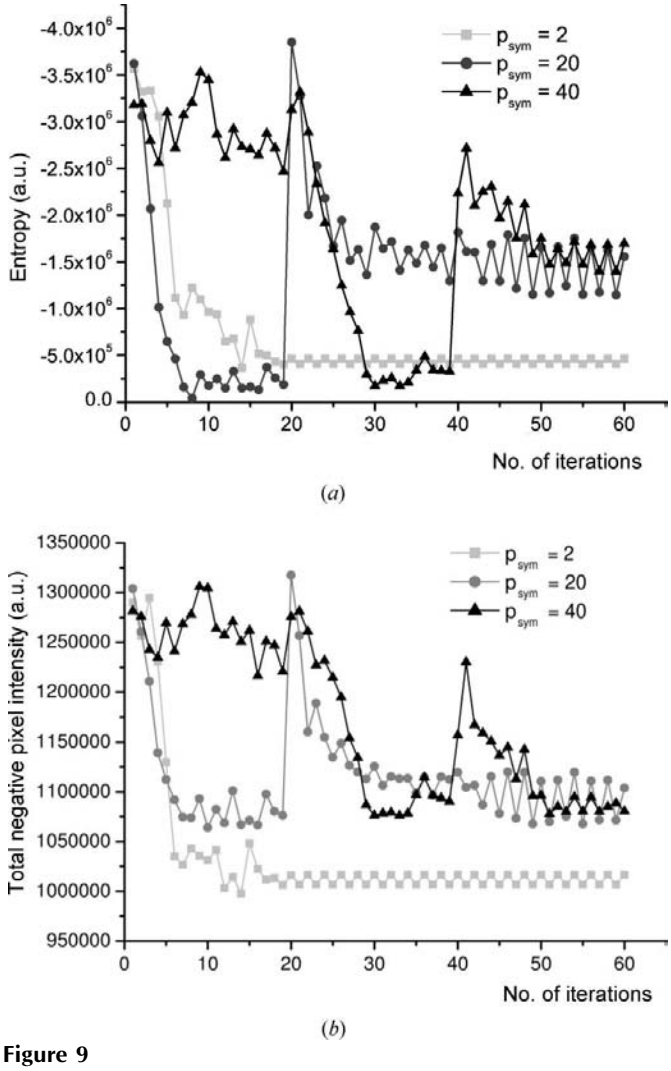

Figure 9
Error metrics for the EGO structure solutions using different values of $p_{\rm sym}$ showing (*a*) the statistical entropy (in arbitrary units) and (*b*) the total negative pixel intensity (in arbitrary units) as a function of the number of iterations.

charge-flipping algorithm is so harmful to the likely outcome; it also alters the original idea of the algorithm, which was to apply symmetry periodically. Clearly, the subsequent application of symmetry will not alter the set of structure factors that are accessed by the algorithm. However, for large values of $p_{\rm sym}$ small numerical artefacts (such as rounding errors) could lead to the symmetry breaking and the algorithm reverting to $p1$ symmetry, so the periodic application of symmetry every $p_{\rm sym}$ cycles was retained in the algorithm to ensure that the symmetry is preserved. This supports the conclusion drawn in the original study of the technique by Oszlányi & Sütő (2004).

When a large number of iterations are carried out before the application of symmetry ($p_{\rm sym} \geq 10$), something important occurs: the solution is able to approach a stable minimum in $p1$ symmetry (shown by the low-entropy plateau) which has less well defined strong peaks; but after symmetry is applied the entropy increases dramatically before decreasing again to a new stable solution. The notable feature is that the second stable entropy is considerably larger than the first, indicating a solution with more strongly defined atomic peaks. The reason for the short period of change after the application of symmetry is shown in the negativity curve. The stable $p1$ solution satisfies the reduced-negativity criterion of the algorithm, but after symmetry is imposed it is not unexpected that more pixels will be made negative as the phases of the solution are altered. A few iterations are needed for this negativity to be removed from the solution; the final stable solution shows increased entropy with respect to the intermediate solution, but there is insignificant change in the overall negativity after the application of symmetry. This is analogous to the relocation in the constrained solution space to the secondary minimum (as described by Fig. 8*c*).

The entropy and negativity curves for the solution with $p_{\rm sym} = 40$ show why care is needed when using this algorithm. While the solution recovered in this instance is good when compared with the ideal structure, the entropy curve shows that the intermediate stable solution did not form until more than 25 iterations had been carried out (compared with fewer than 10 iterations in the $p_{\rm sym} = 20$ example). While this lag time is a common feature of charge flipping, the choice of too small a value of $p_{\rm sym}$ could have affected the outcome of the algorithm. It is therefore important to run the algorithm many times with different starting (random) phases to ensure consistency.

The algorithm was then applied to experimental precession electron diffraction data (Vincent & Midgley, 1994) from a single EGO crystal 55 nm thick. The precession angle was 42 mrad and reflections with $-9 \leq h, k \leq 9$ were included. The result is shown in Fig. 10. In this solution the positions of the atom columns that are recovered are accurate compared with the ideal structure. The widths of the peaks are somewhat larger than in the ideal structure and there is also significant background negativity (indicated by the dotted contours in the potential map). EGO is tetragonal with a large c axis, so that for the [001] zone axis the first-order Laue-zone reflections appear quite close to the centre of the diffraction pattern. The large precession angle exacerbates this problem by increasing the number of higher-order Laue-zone reflections that are sampled by the Ewald sphere, and thus reduces the range of reflections that can be used to $-9 \leq h, k \leq 9$ to avoid overlap with the zero-layer reflections; it is this reduced data set that causes the peaks to be wider than in the ideal structure. Further experimental data for this sample will be published in a forthcoming article (Eggeman *et al.*, 2009).

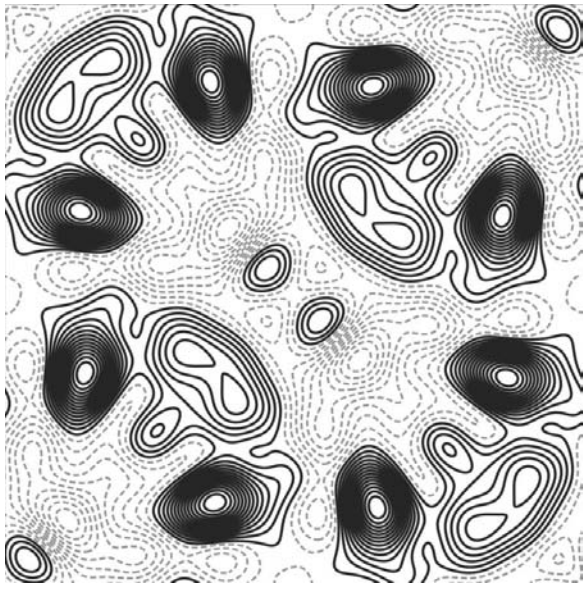

Figure 10
Structure solution recovered from experimental precession electron diffraction data for a 55 nm-thick EGO crystal with a beam precession angle of 42 mrad and $p_{\rm sym} = 40$. All reflections in the range $-9 \leq h, k \leq 9$ were included (*cf.* the ideal structure in Fig. 3).

4. Conclusion

In the case of electron diffraction it seems clear that the symmetry-modified charge flipping algorithm is extremely useful for structure solution. As shown by the cases of kinoite and epidote, it is not clear why some structures can be readily solved in $p1$ and why some need a higher-symmetry constraint. Unlike X-ray diffraction data, often very little (if any) three-dimensional data are obtained from electron diffraction experiments. We surmise that three-dimensional data will strongly constrain the possible locations of peaks within the real-space solution, preventing spurious peaks from forming and forcing real peaks to become more evident. The absence of this constraint in two-dimensional electron diffraction patterns means that some solutions may form where the majority of the potential is forced into a few strong peaks that still satisfy the algorithm: this can be seen in both EGO, where the germanium doublet positions tend to distort the $p1$ solution, and in epidote, where the weaker peaks in the centre of the cell are lost compared with the stronger peaks at the unit-cell edges.

The inclusion of a symmetry constraint can be highly advantageous; however, some provisos are needed. Firstly, the initial structure 'guess' must have no symmetry imposed greater than that provided by Friedel's law, as this will reduce the freedom of the algorithm to find a suitable minimum. Secondly, the algorithm must be allowed to run for a sufficient number of iterations to establish a stable (even if incomplete) intermediate structure, as this will establish a set of viable phase relationships between the stronger reflections in the data set. Finally, once the symmetry is applied the algorithm may need several more iterations for the correct phasing of the majority of reflections to occur, and once a stable structure has formed further iterations will not allow the solution to develop further. As with all iterative algorithms, the stable structure recovered by this process is only correct in that it satisfies the structure-factor moduli input into the algorithm and seeks to maximize the conditions of the algorithm itself; there can be 'correct' solutions that differ from the ideal structure. From this study it seems that monitoring the entropy of the solution can help in the selection of a structure, although careful interpretation is always required.

This work was supported by EPSRC grant number EP/E037275. The authors are grateful to FEI for financial support.

References

Baerlocher, C., Gramm, F., Massuger, L., McCusker, L. B., He, Z. B., Hovmoller, S. & Zou, X. D. (2007). *Science*, **315**, 1113–1116.
Buxton, B. F., Eades, J. A., Steeds, J. W. & Rackham, G. M. (1976). *Proc. R. Soc. London Ser. A*, **281**, 171–194.
Coelho, A. A. (2007). *Acta Cryst.* A**63**, 400–406.
Eggeman, A., White, T. A. & Midgley, P. A. (2009). *Ultramicroscopy*. Submitted.
Frigo, M. & Johnson, S. G. (2005). *Proc. IEEE*, **93**, 216–231.
Giacovazzo, C. (1998). *Direct Phasing in Crystallography*. Oxford University Press.
Jaynes, E. T. (1957). *Phys. Rev. A*, **106**, 620–630.
Karle, J. & Hauptman, H. (1950). *Acta Cryst.* **3**, 181–187.
Kirkland, E. (1998). *Advanced Computing in Electron Microscopy*. New York: Plenum.
Lidin, S. (2008). *Turning Points in Solid State, Materials & Surface Science: A Book in Celebration of the Life & Works of Sir John Meurig Thomas*. Cambridge University Press.
Marks, L. D., Sinkler, W. & Landree, E. (1999). *Acta Cryst.* A**55**, 601–612.
Oszlányi, G. & Sütő, A. (2004). *Acta Cryst.* A**60**, 134–141.
Oszlányi, G. & Sütő, A. (2005). *Acta Cryst.* A**61**, 147–152.
Oszlányi, G. & Sütő, A. (2008). *Acta Cryst.* A**64**, 123–134.
Palatinus, L. & Chapuis, G. (2007). *J. Appl. Cryst.* **40**, 786–790.
Smolin, Y. I. (1970). *Sov. Phys. Crystallogr.* **15**, 36–37.
Stadelmann, P. (2004). *JEMSEMS*, Java version. CIME-EPFL, Lausanne, Switzerland.
Vincent, R. & Midgley, P. A. (1994). *Ultramicroscopy*, **53**, 271–282.
Wu, J., Weierstall, U. & Spence, J. C. H. (2005). *Nature Mater.* **4**, 912–916.
Zhou, Z. & Harris, K. D. M. (2008). *J. Phys. Chem. A*, **112**, 4863–4868.