Acta Crystallographica Section A
Foundations of Crystallography
ISSN 0108-7673

Received 21 July 2004
Accepted 29 October 2004

Ab initio structure solution by charge flipping. II. Use of weak reflections

Gábor Oszlányi* and András Sütő

Research Institute for Solid State Physics and Optics of the Hungarian Academy of Sciences, POB 49, H-1525 Budapest, Hungary. Correspondence e-mail: go@szfki.hu

The original charge flipping algorithm [Oszlányi & Sütő (2004). *Acta Cryst.* A**60**, 34–141] is an amazingly simple structure solution method which works *ab initio* on high-resolution X-ray diffraction data. In this paper, a new version of the algorithm is presented that complements the phase exploration in reciprocal space. Instead of prescribing observed moduli of all structure factors, weak reflections are treated separately. For these reflections, calculated moduli are accepted unchanged and calculated phases are shifted by a constant $\Delta\varphi = \pi/2$. This means that the observed data of weak reflections are *not used* in the iteration, except for the knowledge that they are indeed weak. The improvement is drastic, in some cases the success rate is increased by a factor of ten, in other cases a previously unsolvable structure becomes solvable by the modified algorithm.

1. Introduction

Recently, we presented an amazingly simple *ab initio* structure solution algorithm, which was called *charge flipping* (Oszlányi & Sütő, 2004). The method is based on the fact that the unit cell is mostly empty. The electron density of a crystal has extended regions of zeros, which appear as small oscillations given high-resolution X-ray diffraction data. Therefore, we can limit the algorithm to find only the high-density atomic regions, and accept that the only function of low density is to effectively explore the phase space. In this scheme, the low density will bear no relationship to the true charge density. This is not a high price to pay because, once a starting structural model is available, standard refinement programs can safely complete and refine the electron density. The charge flipping algorithm works iteratively in the manner of Fourier recycling. In real space, the modification is the sign reversal of electron density below a threshold while, in reciprocal space, the observed structure-factor moduli are combined with the calculated phases. The method is truly *ab initio*, the knowledge of chemical composition or even atom types is not required. Symmetry information is not used either, all structures are allowed to float in space group $P1$.

In crystallography, all ingredients of the charge flipping algorithm have been present for quite some time. Solvent flipping was used as a density-modification method in protein crystallography (Abrahams & Leslie, 1996; Abrahams, 1997), an electron-density threshold was applied by the low-density elimination algorithm (Shiono & Woolfson, 1992), and the space group $P1$ was utilized to relieve symmetry constraints (Pavelčík, 1994; Sheldrick & Gould, 1995; Burla *et al.*, 2000). It is their mix that makes the algorithm work. However, we did not arrive at the algorithm along this path. The charge flipping algorithm has its roots in phase retrieval of non-periodic objects (Gerchberg & Saxton, 1972; Fienup, 1982; Millane, 1990), we found it as a modification of Fienup's hybrid input–output algorithm (Fienup, 1982). This was a natural direction, since non-periodic objects with a surrounding support always supply an abundance of zeros. We first handled the object and its support in the same way, then succeeded with less-and-less oversampling and finally reached the Bragg sampling of crystals. Diffraction data of crystals are always undersampled (Sayre, 1952), there are fewer zeros, no known support, a situation that means a greater challenge to the algorithm. We continue to work on the crystallographic phase problem, and anticipate that some developments here will find application in the reconstruction of non-periodic objects.

The charge flipping algorithm is still in its infancy. However, two research groups have already shown that it works for experimental data (Wu *et al.*, 2004; Palatinus, 2004), most surprisingly in the case of modulated crystals (Palatinus, 2004). This is an important result that proves that the algorithm is not based on atomicity. The abundance of low-density values helps to develop it, but the two notions are not identical. Without the use of atomicity, the new method offers special applications, but in normal cases it has lower efficiency than more elaborate dual space direct methods (Miller *et al.*, 1993; Sheldrick, 1998).

In the present paper, we focus on improving the charge flipping algorithm. In the original version, low-density regions of real space were modified to explore the structure-factor phases. Respecting the original philosophy, this exploration can also be done in reciprocal space. We can check what contribution is added to the electron density by *e.g.* the

doi:10.1107/S0108767304027746

weakest 20% of structure factors. Not surprisingly, this contribution is small. Therefore, a different and complementary exploration of the phase space can be done by the real and reciprocal half cycles of the iteration process. In the following, we show that the charge flipping algorithm is more efficient if weak structure factors are treated separately, their calculated moduli are allowed to change freely and their phases are shifted by a constant $\Delta\varphi = \pi/2$. The improvement is drastic, in some cases the success rate is increased by a factor of ten, in other cases, a previously unsolvable structure becomes solvable by the modified charge flipping algorithm.

2. Easy- and hard-to-solve structures

A large pool of dissimilar structures is a precondition for algorithm development. Since our first paper (Oszlányi & Sütő, 2004), we have solved more than a thousand crystal structures using high-resolution synthetic data, and learned that the number of atoms is not the only parameter that determines the difficulty of *ab initio* structure solution. There are remarkably large structures that are easy to solve, and some relatively small ones that require a lot of effort. We need a few difficult examples, matching the current status of the algorithm. Too easy or impossible to solve structures are useless for algorithm development. We praise simulation, this allowed us a fair comparison of algorithm variants by providing the same quality diffraction data for different structures. Throughout this work, we used scattering factors of real atoms at zero temperature and calculated single-wavelength data extending to a resolution of 0.8 Å. In all cases, the tolerance to noise and missing data was extensively checked. Here we summarize our general experience on how the charge flipping algorithm performs with structures of different types.

The most important distinction is between centrosymmetric and non-centrosymmetric structures. Diffraction data of real scatterers have inversion symmetry, and the structure cannot be distinguished from its enantiomer. Therefore, it is very much easier to solve centrosymmetric than non-centrosymmetric structures. The good news for the practical application of the charge flipping algorithm is that 75% of the CSD database (Allen, 2002) is centrosymmetric, so these space groups will not cause difficulty. The bad news is that structures of biological importance are almost entirely non-centrosymmetric. For now we choose non-centrosymmetric examples with space group *P*1 and try to reach the limits here. Note that the method works for all other space groups, but these must be treated as *P*1.

Another simple distinction is between all-light- and heavy-atom structures. Normally, heavy-atom structures are easier to solve. However, atoms that are too heavy may cause problems, they can mask light atoms or prevent structure solution in extreme cases. As charge flipping is also efficient as a *structure completion method*, it is always an option to start the iteration process with a single heavy atom in real space. In *P*1, the coordinates of this atom are not needed. All-light-atom structures are more difficult to solve, especially with more than 100 atoms in the unit cell. So here we select large light-atom structures for further algorithm development.

The third type of classification is according to some structural motif. For simplicity, we call this random *versus* regular structures. Traditional direct methods are based on structure-factor statistics. They work well if the atomic positions can be considered random but break down when a strong structural motif, *e.g.* parallel sheets of planar molecules, occurs. The charge flipping algorithm behaves just the opposite way, it clearly prefers structural motifs: lines, sheets, anything that can be translated into a similar object. In this work, easy-to-solve examples always correspond to regular structures, while random-looking structures are often hard to solve. We believe that randomness is less of a problem than some 'bad' symmetries, whether crystallographic or non-crystallographic, which create traps for the algorithm.

A final observation is that structures with a larger volume per atom are easier to solve. Large intermolecular voids are an advantage, they act as a natural support for the structure even if their position is not known in advance. Fragmented structures with many small molecules are examples of this behaviour. We also made tests by artificially expanding structures. These clearly show that either larger zero regions between unchanged fragments or larger interatomic distances make the charge flipping algorithm more efficient, a situation similar to oversampling.

We have briefly summarized a lot of qualitative experience that helped to select the example structures for the present study. By now, if we see a structure we can guess how difficult its solution will be. It would be useful if the same guess could be done based on observed data alone. We suggest that for a class of equal-size structures a plot of structure-factor moduli $F_{\rm obs}$ sorted in ascending order can serve this purpose. Fig. 1 shows such a plot for a selection of light-atom structures in the space group *P*1. These curves approximately characterize the effective dimensionality of the search space. By this notion, we mean the number of large structure factors determining the

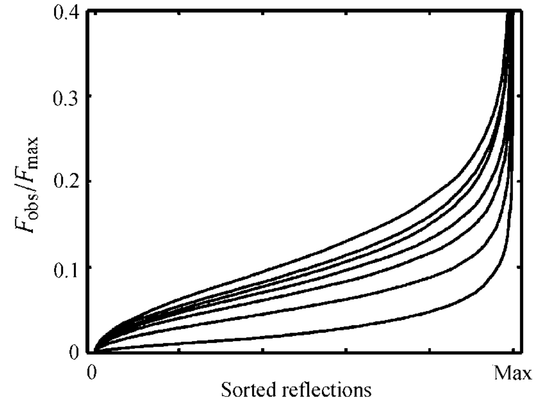

Figure 1
Observed data for a selection of *P*1 structures with 125–150 light atoms in the unit cell. $F_{\rm obs}$ is normalized by its maximum value and sorted in ascending order. Curves from bottom to top characterize a transition from easy- to hard-to-solve structures.

main characteristics of the structure. The smaller this number, the easier the problem of structure solution.

3. Use of weak reflections

For algorithm development, we selected a synthetic peptide with 132 non-H atoms in the unit cell (Karle *et al.*, 1987). This $P1$ test structure carries most properties that cause difficulty for the charge flipping algorithm: it is relatively large, consists only of light atoms, it is non-centrosymmetric and has a continuous spiral backbone. It also corresponds to the 'worst' curve in Fig. 1. Although the original charge flipping algorithm can solve this structure, the success rate is low: 10% or less. For our tests, we always checked 200 random starting phase sets running the algorithm for a maximum of 20000 iteration cycles. We know that success rates may depend on the exact protocol used for structure solution, so we considered the recent suggestion that their use should be abandoned (Elser, 2003). This is a thought-provoking opinion but would require an algorithm that works for any structure and any phase set – unfortunately not the case in practice.

Here we repeat the steps of the original iteration process:

0. A random starting phase set $\varphi(\mathbf{h})$ is generated that satisfies Friedel's law. The structure factors $F(\mathbf{h})$ are created as $F_{\rm obs}\exp(i\varphi)$, where $F_{\rm obs}(\mathbf{h})$ are the observed moduli. The starting electron density $\rho(\mathbf{r})$ is obtained by an inverse fast Fourier transform (FFT).

1. Given a positive threshold δ, the electron density $\rho(\mathbf{r})$ is divided in two parts, $\rho = \rho_1 + \rho_2$ with $\rho_1(\mathbf{r}) = \rho(\mathbf{r})$ if $\rho(\mathbf{r}) > \delta$ and $\rho_2(\mathbf{r}) = \rho(\mathbf{r})$ if $\rho(\mathbf{r}) \le \delta$. The electron density $g(\mathbf{r})$ is generated as $g = \rho_1 - \rho_2$ by flipping the low-density region.

2. Temporary structure factors $G(\mathbf{h})$ are calculated by a FFT.

3. Structure factors $F(\mathbf{h})$ are constructed by using the phases of $G(\mathbf{h})$ and replacing the moduli by $F_{\rm obs}(\mathbf{h})$. $F(0) = G(0)$ is accepted unchanged and $F(\mathbf{h})$ outside the resolution sphere are reset to zero.

4. Finally, the $F(\mathbf{h})$ amplitudes are inverted to obtain the new electron density $\rho(\mathbf{r})$. The next iteration cycle is started from step 1.

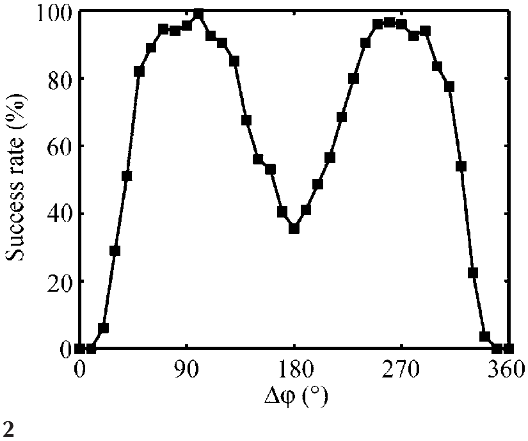

Figure 2
Success rate of the modified charge flipping algorithm as a function of the $\Delta\varphi$ angle. The test structure and the structure-solution protocol are described in the text.

The iteration process is unconditional. Several basic quantities like the R factor, total charge and phase change can be used for monitoring convergence, which is indicated by an abrupt drop in all of them, as shown by Fig. 2 of the paper by Oszlányi & Sütő (2004).

The only change of the above iteration scheme involves step 3. Before using the algorithm, reflections are sorted according to their observed moduli. Structure factors are marked as weak and strong and will be treated separately in the iteration process. Strong reflections are used as before, accepting their calculated phases and replacing their moduli by the observed data. Weak reflections are treated in a new way, their calculated phases are shifted by $\Delta\varphi$ and *their calculated moduli are accepted unchanged*. This means that the observed data of weak reflections are *not used* in the iteration, except for the knowledge that they are indeed weak.

Obviously, this was not our first thought for a new iteration process. We checked many algorithm variants, some are listed in the order of increasing efficiency: (i) all observed reflections are used, (ii) weak reflections are set to zero, (iii) observed moduli of weak reflections are used and their phases are shifted, and (iv) the calculated moduli of weak reflections are accepted and their phases are shifted. The last algorithm variant turned out to be clearly the best. The hint to treat weak reflections separately appeared in our previous paper but in the context of tolerance to missing reflections.

In addition to the electron-density threshold δ, there are two new parameters of the modified algorithm: the number of weak reflections and the shift of calculated phases. In an exhaustive study, all three parameters of the algorithm should be simultaneously optimized. Our experience shows that up to 50% of all observed reflections can be treated as weak. For now, we fix this ratio at 20%, which works well in most cases. The choice of $\Delta\varphi$ and the original δ parameter is more critical. Fortunately, the two parameters are not strongly correlated, a few one-dimensional slices of the parameter space are sufficient.

Fig. 2 shows the success rate as a function of $\Delta\varphi$, while the δ parameter is kept constant at its optimum value. This is an astonishing plot. It shows that even 20% weak reflections play a crucial role; their use can make or break structure solution. At $\Delta\varphi = 0$, the success rate is zero. At this point, weak reflections change freely, which seems to be the worst thing we can do. Note that this differs from the original algorithm. Then, in the $0 < \Delta\varphi < \pi/2$ range, the success rate increases and reaches its maximum slightly above $\pi/2$. At larger phase shifts up to $\Delta\varphi = \pi$, the success rate decreases significantly but does not reach zero again. The plot is then mirror symmetric, positive and negative phase shifts have the same effect. From now on, the modified algorithm is defined as the original charge flipping algorithm completed by a constant $\pi/2$ phase shift of weak reflections.

Fig. 3 shows a comparison of the original and modified algorithms as a function of δ. The modified algorithm is better in two obvious ways: (i) the maximum success rate is higher by a factor of ten, (ii) the practical range of the δ parameter is also wider. The two properties are independent because the

Table 1
Test structures.

Columns: CSD/PDB code and original reference, number of non-H atoms, success rate of the original and modified algorithms (%)

Code and reference	No. of atoms	Original	Modified
`gofmod` (*a*)	188	100	100
`xeqtui` (*b*)	264	100	100
`hegxum` (*c*)	209	0	75
`agazud` (*d*)	219	0	51
`1a7y` (*e*)	314	0	24

(*a*) Biradha *et al.* (1998); (*b*) Witt *et al.* (2000); (*c*) Rontoyianni *et al.* (1994); (*d*) Alexander *et al.* (2002); (*e*) Schaefer *et al.* (1998)

success rates do not scale. A wider δ range is also true for easy structures, where both algorithm variants work for all starting phase sets. For difficult structures, success rates differ so much that there is no point in comparing the average number of iteration cycles per solution. For easy structures, the speed is increased by a factor of up to five. We only add that the modified algorithm is superior in one more way: for difficult structures, it is worth continuing a given run for longer. This is in contrast with the original algorithm, which either works relatively fast or seems to get stuck due to the insufficient perturbation of phases.

To prove further the efficiency of the modified algorithm, we also present examples that the original charge flipping algorithm could not solve. Table 1 shows a selection of large all-light-atom structures in the space group *P*1 and a comparison of success rates for the original and modified algorithms. It contains the two extremes: easy- and hard-to-solve structures. For easy-to-solve structures, the success rate is 100% for both algorithm variants, and the average number of iterations leading to convergence is only a few hundred. Although the size of hard-to-solve structures is similar, they show a completely different behaviour. For these structures, the maximum number of allowed iteration cycles was increased to 100000 and the parameter δ was optimized. Even so, the original algorithm fails completely, while the modified algorithm still works with a fairly high success rate. Although *P*1 is a rather special space group (Xu *et al.*, 2000), we emphasize that improved efficiency of the modified algorithm also holds for other space groups.

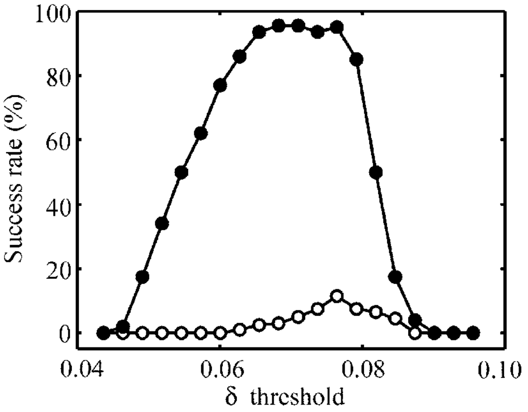

Figure 3
Success rate of the original and modified charge flipping algorithms (empty and solid symbols). The threshold δ is normalized by the maximum of the charge density. The test structure and the structure-solution protocol are described in the text.

Finally, we mention that a single point in Fig. 2 and Fig. 3 corresponds to 200 structure-solution attempts, and in the worst case 4 million iterations. These calculations were greatly facilitated by the use of the `FFTW` software (Frigo & Johnson, 1998), the raw speed up of Fourier transforms was nearly an order of magnitude. For instance, the average time per solution was 37 s for the 132 atom synthetic peptide and 605 s for `agazud`, a difficult example in Table 1. A single parameter set took 0.5–45 h, most time was spent for the cases when the success rate was zero for the original algorithm. The above run times were measured on a PC with a single processor Athlon 3200+.

4. Why does it work?

Any iteration method of *ab initio* structure solution relies on two conditions. The first is the stability of a solution once found. Mathematically, this means that solutions are stable fixed points of the iteration. The second is the absence of metastable solutions. In other words, the algorithm either avoids or escapes traps during the iteration process.

The original charge flipping algorithm certainly has the first property. After convergence, the resulting density proves to be quite robust against deterministic or random perturbations. This holds for the electron density above the threshold, low density may and will change. The stability is remarkable in spite of the fact that the algorithm has a continuous set of fixed points (the structure and its translates) and is chaotic in the sense of a hypersensitive dependence on the initial conditions. Indeed, we checked that it suffices to change just *the last bit of a single initial phase* to provoke an unpredictable shift of the resolved structure. The duality of sensitive dependence on initial conditions and stability after convergence is the main characteristic of *ab initio* structure solution, which is utilized by all multisolution strategies.

The second condition – absence of metastable solutions – is more critical. We have seen that the original algorithm can get stuck in hard to solve cases. The charge flipping algorithm is based on a real-space division of electron density in two parts, and does not act in points where the value is above the threshold. If at the beginning of the procedure high positive densities emerge incoherently, it can be very difficult or practically impossible to remove them by flipping only small densities and keeping the reciprocal-space part of the algorithm unchanged. One can imagine many ways to remedy this situation and it is possible that a suitable modification of the algorithm purely in real space would do equally well. Our choice was guided by three principles: (i) reciprocal-space perturbations may be efficient because they act non-locally in real space, in particular also in points of high densities; (ii) large observed data should be kept intact; (iii) the algorithm should preferably remain deterministic.

Based on the above principles and in analogy with the real-space transformation, it was natural to make a similar division in reciprocal space, determined by the magnitude of the observed Fourier moduli. The particular way we have finally chosen to perturb weak reflections is the result of a series of trials. In retrospect, one can understand why a uniform $\pi/2$ phase shift of weak reflections gives the best result. Let H_1 and H_2 be the sets of $\mathbf{h}$ vectors corresponding to strong and weak reflections, respectively. If $G(\mathbf{h})$ are the temporary structure factors obtained in step 2 of the original iteration cycle, then, according to the modified algorithm, $F(\mathbf{h}) = G(\mathbf{h})\exp(i\Delta\varphi)$ if $\mathbf{h}$ is a weak reflection and $F(\mathbf{h}) = F_{\rm obs}(\mathbf{h})G(\mathbf{h})/|G(\mathbf{h})| = F_{\rm obs}(\mathbf{h})\exp[i\varphi(\mathbf{h})]$ otherwise. This yields the new charge density

$$\rho(\mathbf{r}) = \frac{1}{V}\Bigg\{F(0) + 2\sum_{\mathbf{h}\in H_1} F_{\rm obs}(\mathbf{h})\cos[2\pi\mathbf{hr} + \varphi(\mathbf{h})] + 2\sum_{\mathbf{h}\in H_2} |F(\mathbf{h})|\cos[2\pi\mathbf{hr} + \varphi(\mathbf{h}) + \Delta\varphi]\Bigg\}.$$

When $\Delta\varphi = \pi/2$, each partial wave of a weak reflection is replaced by an orthogonal one. This is the largest perturbation we can make. As the density below the threshold is alternating in sign, the algorithm converges to a limit two-cycle and only the phases of every second iteration cycle are comparable. After two cycles, the phase shift of weak reflections adds up to π, which is the largest value modulo 2π. This argument is valid as long as high densities are only weakly perturbed. Using large δ and many weak reflections, the optimal choice of $\Delta\varphi$ is somewhat larger than $\pi/2$.

It is important to let the moduli of weak reflections change freely. But, in addition, non-zero phase shifts of weak reflections are necessary to stabilize the freely changing moduli. After convergence, low density in real space and small structure factors in reciprocal space will not contain structural information. Their role is to explore the phase space and broaden the region that can be hit to obtain the approximately correct structure. This is better achieved by using fewer constraints. If structure-factor moduli of weak reflections are constrained by the observed values, the success rate goes down by a factor of three, though its $\Delta\varphi$ dependence remains unchanged.

The modified charge flipping algorithm is still deterministic. Apart from aesthetic reasons, we had other arguments against a partial randomization of the algorithm. During the work on our first paper (Oszlányi & Sütő, 2004), we learned that randomization of low electron densities is a disastrous perturbation that prevents convergence. Nevertheless, we have tested many variants of random perturbation in reciprocal space. In contrast to what happens in real space, practically any random change of weak reflections improves the convergence of the original algorithm. This shows that, as long as large observed data are left unchanged, a non-local perturbation of the electron density is a useful exploration of phase space. However, none of the random algorithm variants comes near to the efficiency of the modified algorithm, which is absolutely the best and deserves to be treated separately.

5. Conclusions

In this paper, we presented an efficient modification of the charge flipping algorithm used for *ab initio* structure solution of high-resolution X-ray diffraction data. In the original version, low-density regions of real space were used to explore the structure-factor phases (Oszlányi & Sütő, 2004). A complementary exploration of the phase space can also be done in the reciprocal half of the iteration cycles. For this, weak and strong reflections are treated separately: while strong reflections are used as before, weak reflections are treated in a new way. Their calculated moduli are allowed to change freely, and their phases are shifted by a constant $\Delta\varphi = \pi/2$. The $\pi/2$ shift can be understood as the maximum perturbation of reflections that add only a small contribution to high electron densities. Paradoxically, it is better *not to use* observed moduli of weak reflections: in the search for a solution, they create only unwanted constraints.

This method of structure solution is only approximate, both the low density in real space and small structure factors in reciprocal space bear no relation to the true structure. Their role is to develop high electron densities, which can be considered as a (very complete) starting structural model. Our view of the charge flipping algorithm is that of a dynamical system. In the very high dimensional space of phases, a deterministic rule drives the system around. Once an approximate solution is hit, it is stable against perturbations including those that have driven the system so far. This suggests that given favourable conditions (high resolution and small size) the crystallographic phase problem is benign, at least in comparison with the most general global minimum problem of optimization.

The modified algorithm is still based on the existence of extended regions of zeros. A future improvement may be achieved by some sort of sharpening, though higher-resolution and lower-temperature experimental data are always preferable. The reason is that all three trivial symmetries of diffraction by real objects – enantiomer, shift and negative electron density – can cause traps for an algorithm. While centrosymmetric structures are free from the enantiomer problem, too much sharpening opens an ambiguity for negative density. Most non-centrosymmetric symmetries suffer from special ambiguities, at least when structures are handled in the space group $P1$. Further work with symmetries is required to see how this problem can be solved.

The modified algorithm performs so much better than the original version that it can be considered as the new charge flipping algorithm. Whether it will make a practical piece of software remains to be seen. Charge flipping certainly started as a toy but recently it has been successfully applied to experimental data of both normal (Wu *et al.*, 2004) and modulated crystals (Palatinus, 2004). Structure solution of modulated crystals, quasicrystals and non-periodic objects are promising new directions; in all cases, the algorithm can work without the requirement of atomicity. It is also likely that either the real- or reciprocal-space halves of the algorithm will be utilized by other dual-space algorithms; we see no reason why they should work only in this combination.

Note added. Since submission of the manuscript, we have found that unobserved reflections outside the resolution sphere (which were previously reset to zero) can be treated in the same way as observed weak reflections described in §3. Both algorithm versions perform similarly well. However, concurrent special treatment of weak and unobserved reflections is too much, they impair each other's effect.

We thank Lukáš Palatinus for useful discussions. This research was supported by OTKA grants T043494 and T042914 and the work of GO was also funded by a Bolyai János Scholarship.

References

Abrahams, J. P. (1997). *Acta Cryst.* D**53**, 371–376.

Abrahams, J. P. & Leslie, A. W. G. (1996). *Acta Cryst.* D**52**, 30–42.

Alexander, J. M., Clark, J. L., Brett, T. J., & Stezowski, J. J. (2002) *Proc. Natl Acad. Sci. USA*, **99**, 5115–5120.

Allen, F. H. (2002). *Acta Cryst.* B**58**, 380–388.

Biradha, K., Dennis, D., MacKinnon, V. A., Sharma, C. V. K. & Zaworotko, M. J. (1998). *J. Am. Chem. Soc.* **120**, 11894–11903.

Burla, M. C., Carrozzini, B., Cascarano, G. L., Giacovazzo, C. & Polidori, G. (2000). *J. Appl. Cryst.* **33**, 307–311.

Elser, V. (2003). *Acta Cryst.* A**59**, 201–209.

Fienup, J. R. (1982). *Appl. Optics*, **21**, 2758–2769.

Frigo, M. & Johnson, S. G. (1998). *Proc. ICASSP*, **3**, 1381–1384.

Gerchberg, R. W. & Saxton, W. O. (1972). *Optik (Stuttgart)*, **35**, 237–246.

Karle, I. L., Flippen-Anderson, J., Sukumar, M. & Balaram, P. (1987) *Proc. Natl Acad. Sci. USA*, **84**, 5087–5091.

Millane, R. P. (1990). *J. Opt. Soc. Am.* A**7**, 394–411.

Miller, R., DeTitta, G. T., Jones, R., Langs, D. A., Weeks, C. M. & Hauptman, H. A. (1993). *Science*, **259**, 1430–1433.

Oszlányi, G. & Sütő, A. (2004). *Acta Cryst.* A**60**, 134–141.

Palatinus, L. (2004). *Acta Cryst.* A**60**, 604–610.

Pavelčík, F. (1994). *Acta Cryst.* A**50**, 467–474.

Rontoyianni, A., Mavridis, I. M., Hadjoudis, E. & Duisenberg, A. J. M. (1994) *Carbohydr. Res.* **252**, 19–32.

Sayre, D. (1952). *Acta Cryst.* **5**, 843.

Schaefer, M., Sheldrick, G. M., Bahner, I. & Lackner, H. (1998) *Angew. Chem. Int. Ed.* **37**, 2381–2384.

Sheldrick, G. M. (1998). *Direct Methods for Solving Macromolecular Structures*, edited by S. Fortier, pp. 401–411. Dordrecht: Kluwer Academic Publishers.

Sheldrick, G. M. & Gould, R. O. (1995). *Acta Cryst.* B**51**, 423–431.

Shiono, M. & Woolfson, M. M. (1992). *Acta Cryst.* A**48**, 451–456.

Witt, D., Lagona, J., Damkaci, F., Fettinger, J. C. & Isaacs, L. (2000). *Org. Lett.* **2**, 755–758.

Wu, J. S., Spence J. C. H., O'Keeffe, M. & Groy, T. L. (2004). *Acta Cryst.* A**60**, 326–330.

Xu, H., Hauptman, H. A., Weeks, C. M. & Miller, R. (2000). *Acta Cryst.* D**56**, 238–240.