The charge flipping algorithm

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

This paper summarizes the current state of charge flipping, a recently developed algorithm of *ab initio* structure determination. Its operation is based on the perturbation of large plateaus of low electron density but not directly on atomicity. Such a working principle radically differs from that of classical direct methods and offers complementary applications. The list of successful structure-solution cases includes periodic and aperiodic crystals using single-crystal and powder diffraction data measured with X-ray and neutron radiation. Apart from counting applications, the paper mainly deals with algorithmic issues: it describes and compares new variants of the iteration scheme, helps to identify and improve solutions, discusses the required data and the use of known information. Finally, it tries to foretell the future of such an alternative among well established direct methods.

1. Introduction

Charge flipping (CF) is a deceptively simple structure-determination algorithm that solves the phase problem with much weaker assumptions than classical direct methods. As such, its conceptual importance and potential in crystallographic teaching was clear from the start (Oszlányi & Sütő, 2004, 2005), while, considering the power and widespread use of existing direct methods, it seemed less likely that charge flipping could make a practical impact. We were fortunately wrong in this respect. Just three years after the first publication, and thanks to the quick and creative response of the crystallographic community, numerous applications of CF have already emerged, user programs are being created and there is still progress in improving the basic algorithm. It now seems timely to critically summarize the current state of this quickly developing field.

The paper deals with all aspects of the structure-solution process. It does this in a tutorial style with a particular emphasis on algorithmic issues. §2 gives a brief introduction to the phase problem and dual-space approaches for its solution. In §3, this is followed by a precise description of the basic charge flipping algorithm and its special properties. Then in §4 we introduce some figures of merit that were found useful to identify the convergence in the otherwise unconditional iteration process. In §5, we discuss different ways of choosing the only parameter of the algorithm, and show a clean-up procedure that works after the *ab initio* solution and before a true refinement. §§6 and 7 introduce significant improvements of the basic algorithm first in reciprocal space and then in real space – some of these are published here for the first time. §8 discusses the correct treatment of unobserved data and sharpening of density maps with observed data. This is also the place where the results of all previous improvements are compiled in a table for a carefully selected example. §9 explains that the *ab initio* algorithm can also be utilized in those situations when some preliminary structural information is available. §10 is perhaps the most awaited part of the paper that counts successful applications and lists various user programs. In the concluding section, limitations and future prospects of the charge flipping method are discussed.

2. The phase problem and dual-space approaches

In the simplest case, the electron density of a crystal is faithfully represented in reciprocal space by a complete set of complex structure factors: $F(\mathbf{h}) = |F(\mathbf{h})| \exp[i\varphi(\mathbf{h})]$. If these could be directly measured, the act of structure solution would be nothing but an inverse Fourier transform. Unfortunately, a simple diffraction experiment provides information only on the magnitudes of structure factors but cannot determine the phases. This leads to the famous phase problem of crystallography, with additional complications of limited resolution, missing reflections, unreliable intensities *etc.* The main difficulty of this problem is that any phase set in the high-dimensional space is compatible with the measured data, and only additional information on the electron density can select the correct solution. Mathematical techniques that do this job are called direct methods.

Classical direct methods utilize the constraints of positivity, atomicity and chemical composition that lead to statistical phase relations of structure factors (Hauptman & Karle, 1953; Karle & Hauptman, 1956; Woolfson, 1987; Giacovazzo, 1998). While classical direct methods, in their original form, work entirely in reciprocal space, there are other variants that switch back and forth between real and reciprocal spaces. The advantage of these dual-space methods is that they offer remarkable freedom of imposing constraints in each of the
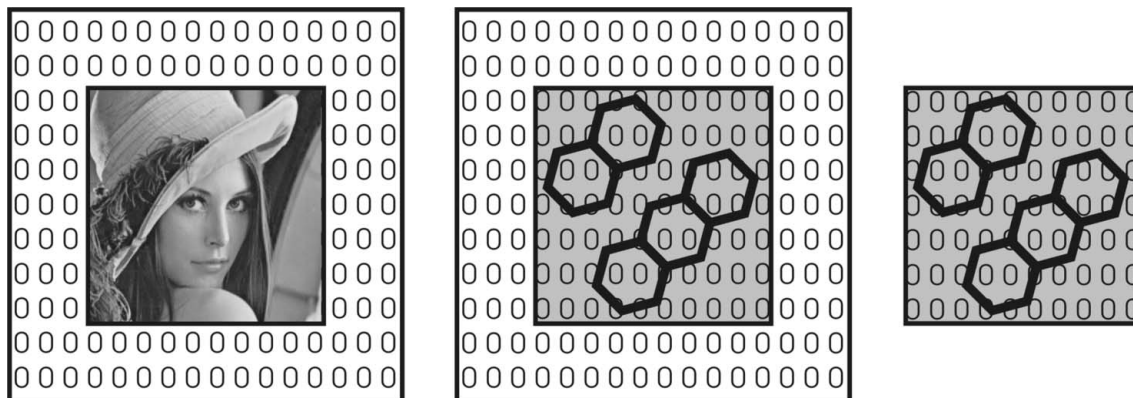
Figure 1
Left: The standard Lena image as a non-periodic object of a square support, surrounded by a region of zero density. Middle: A much simpler object where zero density is also present inside the confining square. Right: The same molecular object without the surrounding region of known zero density. This corresponds to the unit cell of a periodic crystal.

spaces. The disadvantage is that this cannot be done simultaneously: satisfying the constraint in one space may often violate the constraint in the other. The usual solution is some iterative scheme that is computationally expensive and was not practical before the widespread use of the fast Fourier transform (FFT) (Cooley & Tukey, 1965; Barrett & Zwick, 1971; Frigo & Johnson, 2005). Therefore, dual-space procedures were initially used for improving partial solutions like phase correction (Hoppe & Gassmann, 1968) and various density-modification methods in protein crystallography (Wang, 1985; Abrahams & Leslie, 1996; Zhang *et al.*, 2001) – their *ab initio* application is more recent. Today we can be sure that all current programs of structure determination contain smaller or larger sections of various dual-space procedures and differ significantly from the pure classical scheme (Miller *et al.*, 1993; Sheldrick, 1998; Burla *et al.*, 2005; Yao *et al.*, 2006).

It is also fair to pay tribute to another field. The first *ab initio* dual-space method that works without statistical phase relations comes from optics, where the Gerchberg–Saxton–Fienup algorithm (Gerchberg & Saxton, 1972; Fienup, 1982) is the standard phase-retrieval tool of non-periodic objects. This method requires that (i) preliminary information on the size and shape of the support of the object is available, and (ii) the support is surrounded by a known region of zero scattering density. If the surrounding zero region is large enough, the continuous Fourier transform can be measured at a sufficiently fine sampling to fulfil the Nyquist criterion, and, if the shape of the support is known with sufficient accuracy, an iterative scheme can reconstruct any general object. Note that this is also the mathematical background of recent progress in single-particle imaging (Sayre, 2002). In contrast to non-periodic objects, the scattering density of a periodic crystal is always undersampled (Sayre, 1952) and, therefore, the above scheme is not applicable. Fortunately, crystal structures are very special objects, which occupy only a small fraction of space with high density values and leave a large fraction of space for nearly zero density (see Fig. 1). Given high-resolution data, we may attempt to find and use these zero plateaus. Charge flipping is a method that follows this plan.

3. The basic charge flipping algorithm

The simplest Fourier cycle (shown in Fig. 2) comprises four steps: (i) a real-space modification of the electron density, (ii) a Fourier transform to reciprocal space, (iii) a reciprocal-space modification of calculated structure factors and (iv) an inverse Fourier transform back to real space. The use of the FFT requires that the continuous electron density is also sampled at regular grid points. With observed data within a resolution sphere of radius $|\mathbf{h}|_{max} = 1/d_{min}$, the necessary grid spacing is $\Delta r \le d_{min}/2$.

The charge flipping algorithm is a special variant of the above scheme. To make the name-giving real-space modification plausible, Fig. 3 shows the electron-density samples (pixels) sorted in ascending order for two structures calculated at different resolutions. These plots illustrate that at high resolution a small number of pixels with large values carry the majority of structural information. The information content of the many pixels scattered around zero is much less, it is enough to know that their values are confined in a narrow band. If we combine the constraint of positivity with a weak perturbation of these nearly zero plateaus, we might get an algorithm that simultaneously explores the phase space and decreases its effective dimensionality. The choice of CF is to reverse the sign of small electron density. This introduces small discontinuities while keeping the standard deviation of the whole density distribution practically constant.
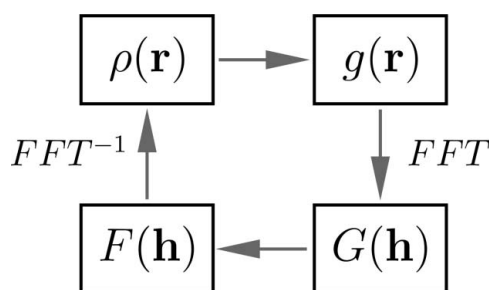

Figure 2
Steps of the simplest Fourier cycle.

Here we recall the precise steps of initialization + one iteration cycle:

0. The algorithm is initialized by selecting a random phase set $\{\varphi(\mathbf{h})\}$ which satisfies Friedel's law. The structure factors are created as $F(\mathbf{h}) = F_{obs}(\mathbf{h}) \exp[i\varphi(\mathbf{h})]$, where $F_{obs}(\mathbf{h})$ are the observed moduli. Unobserved moduli and $F(0)$ are set to zero. The starting electron density $\rho(\mathbf{r})$ is obtained by an inverse FFT.

1. Given a positive threshold δ, the electron density $\rho(\mathbf{r})$ is divided into two parts, $\rho = \rho_1 + \rho_2$ with $\rho_1 = \rho$ if $\rho \geq \delta$ and $\rho_2 = \rho$ if $\rho < \delta$. 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 the FFT. Thus we arrive at a parallelepiped-shaped volume that contains both observed reflections within the resolution sphere and unobserved reflections outside it.

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

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

Evolution of a dynamical system is a good analogue of the above process, the electron density (or phase set) can be considered as a point of high-dimensional space that is driven by the perturbations and constraints along a complicated path. A successful run has three characteristic parts: an initial transient, a long stagnation period before and the stable state after the convergence. The transient usually takes only ~ 10 cycles, and is needed to reach a subspace of zero plateaus

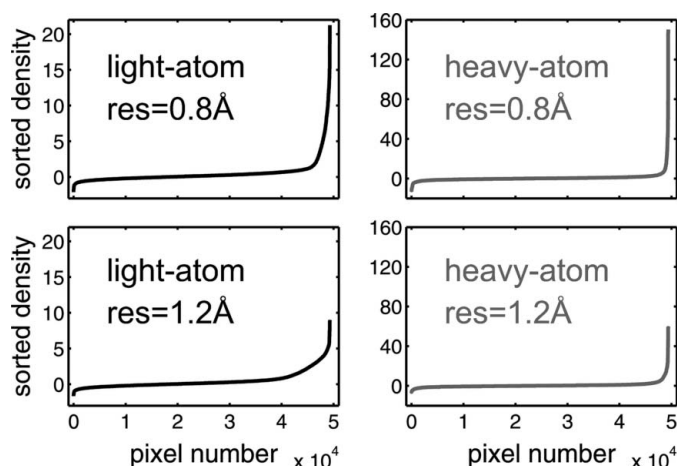**Figure 3**
Electron-density samples (in e Å^{-3} units) calculated on a 0.4 Å grid and sorted in ascending order. Left and right columns correspond to a light-atom structure with composition $C_{160}N_2O_{10}$ and to a heavy-atom structure with composition $C_{48}O_4P_4S_4Cl_{32}Nb_8$ (Irngartinger *et al.*, 1999; Stumpf *et al.*, 1999). Top and bottom rows come from ideal F_{obs} data at $d_{min} = 0.8$ and 1.2 Å resolutions. It can be clearly seen that large electron density is concentrated in a small fraction of space – this behaviour is stronger with higher-resolution data and when heavy atoms are present.

with positive peaks. In the stagnation period, a more detailed exploration of this subspace follows. Although the CF algorithm is truly deterministic, it is also chaotic, the actual path is extremely sensitive to small changes of the starting point or of δ. When the path reaches the region of convergence, the solution occurs quickly, often 10–100 cycles are sufficient for completion. After this, the solution is remarkably stable, the same algorithm that led to the solution is unable to kick it out from this state.

There are several attractive properties of the CF algorithm. The first is its extreme simplicity and easy implementation. The second is its truly *ab initio* character. The electron density is represented on a grid but the concept of atoms is not acknowledged during the solution. Therefore, CF makes no use of atom types, chemical composition or even the total charge of the unit cell. There is no need for utilizing symmetries either, all structures may be allowed to float freely in the space group $P1$. Electron density can continuously evolve up to the solution, where simple peak picking identifies atoms – the chemical composition and space-group symmetry can be determined afterwards. The principle of such an algorithm differs a great deal from that of both classical direct methods and global optimization, it works without statistical phase relations or any cost function. Therefore, charge flipping is complementary to other direct methods, it can work well in troublesome situations such as unknown chemical composition, ambiguous space groups, the presence of pseudosymmetry or disorder.

4. Figures of merit

The iteration process is unconditional and could continue forever. The practical question is when to stop it, preferably in an automatic way and after recognizing that the solution has been found. The choice of some good figures of merit (FOM) is discussed below.

In our first paper, the convergence was indicated by a sharp drop in the following quantities: $F(0)$ corresponding to the total density, the usual R factor, and the average phase change. These are examples of three main types of FOM that characterize (i) the electron density itself, (ii) the fit to the observed data, and (iii) the unobserved phases. Here we introduce some variants of (i) and (ii) and suggest all of them are followed as a vector. This can be useful when we do not have high-resolution F_{obs} data, and the usual drop in the R factor becomes too small or even absent.

So here is a list of some useful FOM's:

$$R = \sum \left| F_{obs} / \sum F_{obs} - F_{calc} / \sum F_{calc} \right|,$$
$$\text{where } F_{calc} = |G|;$$
$$\text{CC} = \langle xy \rangle / (\langle x^2 \rangle \langle y^2 \rangle)^{1/2},$$
$$\text{where } x = F_{obs} - \langle F_{obs} \rangle \text{ and } y = F_{calc} - \langle F_{calc} \rangle;$$
$$F(0) = \text{total charge};$$
$$\sum \rho^3 = \text{peakiness};$$
$$R_H = \sum |\text{sort}(\rho) - \text{sort}(\rho_{ref})| / \sum |\text{sort}(\rho_{ref})|.$$

feature articles

The R factor and the correlation coefficient (CC) characterize the fit of calculated and observed moduli, and with the usual definition both are insensitive to the correct scaling of moduli. At convergence, the R factor drops, and the correlation coefficient increases. Usually, both show enough contrast but, with low-resolution data, normalized structure factors (E's) or neutron diffraction data, it can happen that one or both become insensitive to the solution.

It is remarkable that convergence is indicated by some global properties of the density map *without* considering the data. Of these, $F(0)$ is our current favourite. It works well even when the R factor and CC fail, and its only annoying property is that it continuously creeps downward before the significant drop at convergence. This is in contrast with the R factor or CC which quickly become constant but, even so, it is always worth following $F(0)$. The other useful property of the map is its 'peakiness'. This quantity originates from Cochran's work (Cochran, 1952) and was actively used in the phasing process by Stanley (1979, 1986). Its passive use as a FOM is suggested here. Convergence is indicated by its sharp increase and a characteristic maximum can be observed before the solution slightly degrades. While peakiness is an interesting quantity related to triplets, its utility seems to be limited to high-resolution data. The third property of the density map requires a reference histogram (Lunin, 1988; Zhang & Main, 1990), so this is a slight departure from the pure *ab initio* approach.
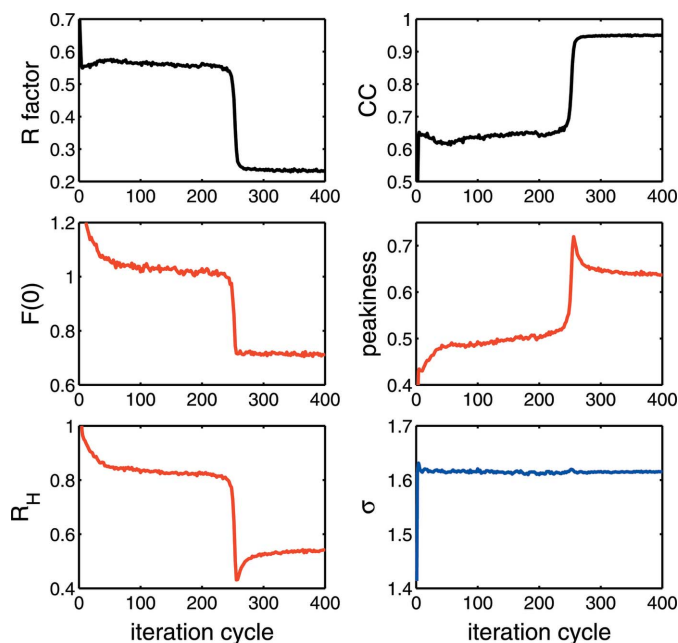

Figure 4
A successful run of the charge flipping algorithm for a typical organic structure using high-resolution data. Subplots show some global quantities (calculated in the g state) as a function of the iteration cycle. First row: the R factor and the correlation coefficient. Second row: the total charge and the peakiness (both normalized by their ideal values). Third row: the match of the histogram and the standard deviation of the electron-density map (e Å^{-3} units). The first five quantities are good figures of merit that clearly indicate the convergence, while σ with its nearly constant value is important for selecting the algorithm's threshold parameter.

(Even so, a sufficiently good reference histogram can be calculated after creating an artificial structure with similar composition and randomly positioned atoms at a minimum distance from each other.) Again, the histogram could be used in an active or passive way. While the phasing power of a histogram is weak, the match of calculated and reference histograms R_H is a very good new FOM. Convergence is indicated by its sudden drop and a characteristic minimum can be observed before the solution settles. A particularly useful property of R_H is that it shows enough contrast also at lower resolutions.

5. Selection of δ and finishing a solution

The basic CF algorithm has a single parameter: the threshold δ. Its optimal choice is a delicate problem and is important because it determines (i) the computational cost of a solution and (ii) the quality of the resulting map.

The computational cost of a solution at a given parameter value (or algorithm variant) can be measured only on the basis of statistics. For this we need a number of runs (N) that start from different random phase sets and are allowed to continue for a maximum number of cycles (M). Thus, the length of the ith run is $x_i \leq M$. If convergence can be detected by some of the previous FOM's, then the iteration can be stopped at $x_i < M$. With $n > 0$ successful runs, we calculate

$$\bar{x}_S = \frac{\sum_{i=1}^{N} x_i}{n} \leq NM$$

which expresses directly the average number of cycles spent for a single solution.

Once the efficiency of a given δ parameter can be measured, we are ready to search for its optimum value. Because peaks of a density map are very much dependent on the grid size, absolute scale and thermal parameters, the choice of δ is a delicate problem, also influenced by the chemical composition and the type of the structure. In the simplest case, δ is kept constant for the complete run and the fraction of pixels that are flipped ($\phi = N_{\rho < \delta}/N_{\text{grid}}$) may vary in each cycle. In practice, ϕ varies strongly only in the initial transient, it is nearly constant during the stagnation period, and increases slightly at the point of convergence. A suggested variant was to change the roles of δ and ϕ, keeping ϕ constant and allowing δ to vary (Wu *et al.*, 2004). Technically, this also required the sorting of pixels in each cycle which is a time-consuming step, but sorting can be replaced by a much faster multipass binning process (Coelho, 2007a). The main advantage of constant ϕ is that its choice is not dependent on the correct scaling of observed and calculated structure factors. It also yields a converged density of slightly better quality because the fraction of pixels that are flipped does not increase at the point of convergence, as with a constant δ. Otherwise, to find the optimal δ or to find the optimal ϕ takes just as much effort and, once found, the cost of a single solution is also the same.

Here we suggest a new parametrization of δ that is more physical than previous approaches. The threshold is now expressed as $\delta = k\sigma$, where k is a fixed number and

$\sigma = (\langle\rho^2\rangle - \langle\rho\rangle^2)^{1/2}$ is the standard deviation of the electron-density map. σ expresses the spread of density values in a single number and is constant at a given resolution due to Parseval's theorem. The standard deviation seems to be a good basis for selecting δ because (i) it makes k independent of the correct scaling of structure factors and (ii) it puts a margin of reliability on the final density map. When σ is recalculated in the g state of each iteration cycle (see Fig. 4), it is still nearly constant throughout the process and is remarkably insensitive to the choice of k. This behaviour was carefully checked for a large number of crystal structures, including those which contain many heavy atoms and a few light atoms in the same structure. In all cases, the optimal value of the new parameter was found to be in a narrow range $k = 1.0$–1.2, and the cost of a solution was the same as if it were calculated with the optimal constant δ or constant ϕ. The positive correlation between δ and σ is intuitively obvious. Higher-resolution data or the presence of heavy atoms increase the positive electron-density values while zero or nearly zero values are always present. The result is a larger spread of electron density that also requires a larger δ to achieve the same perturbation. The cause of the approximate coincidence of σ and the optimal δ is less clear and an explanation should consider the particular shape of the density distribution. We do not say that the problem of selecting the optimal parameter *a priori* is completely resolved. Small structures are readily solved with any k in the 1.0–1.2 range, but with increasing size and complexity the

useful range quickly becomes narrower. One symptomatic treatment, suggested by Coelho (2007b), is to periodically ramp the parameter within the expected limits. This works well but to solve protein-sized structures we still need a real theory of selecting δ.

The quality of the solution is also related to the threshold parameter. The fastest convergence usually requires a relatively large δ and the perturbation of many pixels in the low-density region. These pixels contribute significantly to the structure factors and their flipping is the cause of a limit cycle behaviour after convergence. However, these pixels should not be considered as part of the solution, and can be deleted after the convergence. Thus the correct 'clean-up' procedure of the solution is to switch from charge flipping to ten cycles of low-density elimination as was first introduced in the program *SUPERFLIP* (Palatinus & Chapuis, 2007). Our experience shows that even a single cycle suffices and better quality histograms are obtained if we do not decrease δ. The recipe is: keep or increase $\delta = k\sigma$ used for the solution, set $\rho < \delta$ to zero and complete the Fourier cycle by prescribing F_{obs}. For demonstration, we selected an organometallic compound (Florke & Haupt, 1990) (376 non-H atoms per 7394 Å^3, space group $P2_1/n$) that in addition to light atoms (C, O) contains both heavy (P, Cl) and very heavy atoms (I, Re) in the same structure. Previously, this was a typical problem case for the CF algorithm because the large standard deviation of the electron-density map required a large threshold parameter for the solution, and thus could not provide the position of the
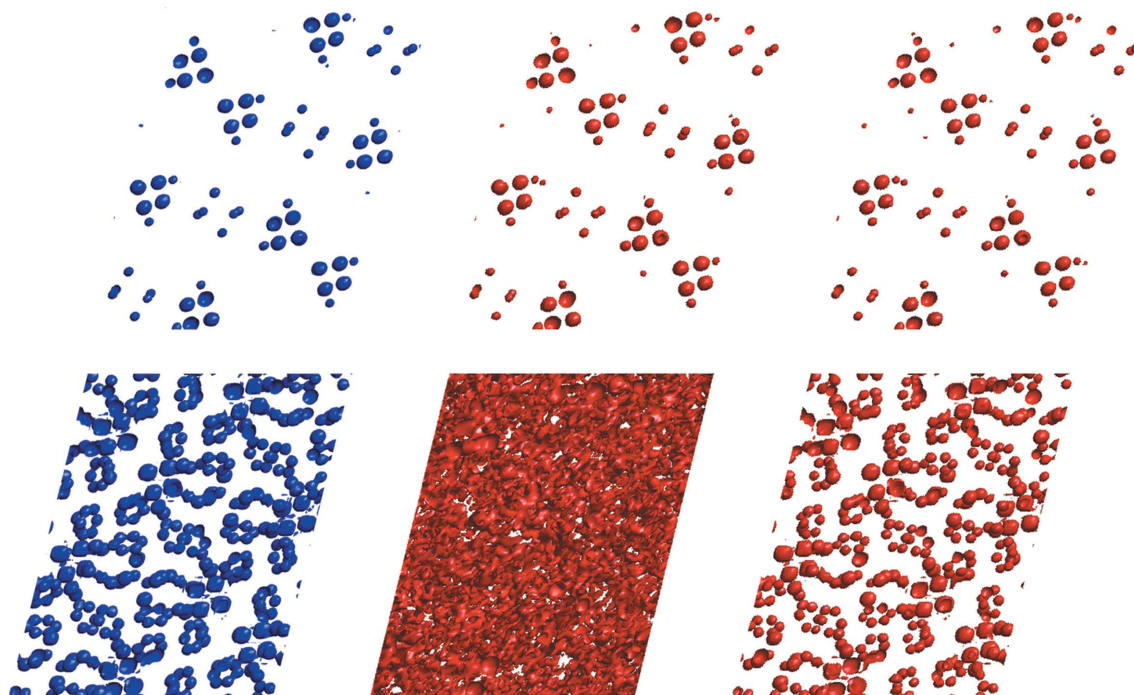

Figure 5
Electron-density map of the unit cell for an organometallic example described in the text. Columns from left to right: the reference structure (blue), the *ab initio* solution by charge flipping using the threshold parameter $\delta = 1.0\sigma$ (red), the final state after one step of the clean-up procedure using the parameter $\delta = 1.5\sigma$ (also red). Top and bottom rows are isosurface plots with levels 1.0σ and 0.2σ. In all cases, heavy atoms can be identified already above the level of one standard deviation. Light atoms appear only after the clean-up procedure and only at the 0.2σ level.

light atoms. The clean-up procedure solves this problem, convincing illustrations are shown in Fig. 5.

The efficiency of the above one-cycle clean-up can be given a simple explanation. Let us follow the strongest reflections $F(\mathbf{h})$ after convergence sets in. Then the Fourier components F_1 and F_2 corresponding to ρ_1 and ρ_2 become orthogonal to each other in the complex plane, see Fig. 4 of Oszlányi & Sütő (2004). From this instant on, a continued iteration enters a two-cycle switching between $F_1 + F_2$ and $F_1 - F_2$. By cutting off ρ_2, we keep only F_1 which, after adjusting its modulus to F_{obs}, is our best approximation to the ideal F.

6. Improvements in reciprocal space

A simple Fourier cycle is susceptible to stagnation and we must be prepared to break it in some way. This requires a fine balance of constraints and weak perturbations, both of which can be implemented in real or reciprocal space. In the basic CF algorithm, the prescription of F_{obs} data is a pure modulus constraint, while the charge flipping step is a mix of the positivity constraint and of a low density perturbation. A large number of algorithm variants can be constructed by adding (or removing) constraints and by introducing new perturbations. Naturally, we want only those variants that increase the efficiency of the algorithm.

The first improvement in reciprocal space was found accidentally. To check the algorithm's tolerance to missing data, we replaced weak reflections by zeros and observed that convergence actually became faster (Oszlányi & Sütő, 2004). The fraction of reflections considered weak was $\alpha = 10–50\%$ of all data within the resolution sphere. The speed-up was significant (a factor of 2–4 for simple structures), and the final R factor increased without an obvious degradation in the quality of the reconstructed electron density. This first try (called 'weak = 0' for short) suggested that weak reflections can be utilized in a similar way to low electron density, their perturbation improves the exploration of phase space and decreases the dimensionality of the problem. Subsequently, we have found many variants of reciprocal-space perturbations. A particularly successful one is casually called the $\pi/2$ version of CF (Oszlányi & Sütő, 2005). This algorithm has three parameters: the threshold δ, the fraction of reflections α that are considered weak, and a phase shift $\Delta\varphi$. In step 3 of the iteration cycle, weak reflections are treated in a special way: their calculated phases are shifted by the constant $\Delta\varphi$ and their calculated moduli are accepted unchanged. (Beware, phases of Friedel pairs must be shifted by $\pm\Delta\varphi$!) An exhaustive search for the best three parameters is not realistic. It is usually sufficient to try $\alpha = 20$ or 40%, check a narrow range above $\Delta\varphi = \pi/2$ and search for the usual δ. The speed-up with good parameters can be quite significant, for large structures the number of iteration cycles spent for a solution often decreases by a factor of 50–100. The range of useful δ also becomes broader, and there is an increased chance that a solution is found after a long stagnation period. Note that the perturbation of weak reflections assumes that there is enough
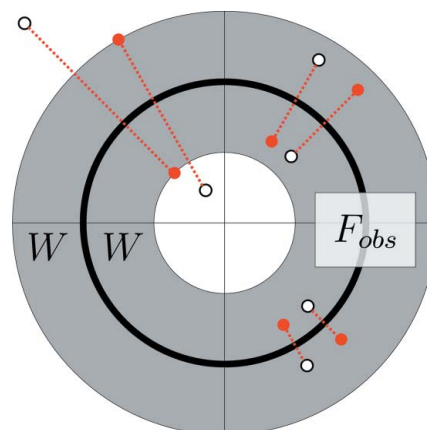

Figure 6
$F_o + \Delta F$ relaxation of the modulus constraint in the charge flipping algorithm. A selected structure factor before and after this step $[G(\mathbf{h}) \rightarrow F(\mathbf{h})]$ is plotted by empty and red symbols in the complex plane for a few possible cases. The measured $F_{obs}(\mathbf{h})$ defines the radius of the (bold) circle that is the centre of the mirror operation and the parameter W defines a (shaded) ring that limits the change of the modulus.

data, so the utility of these algorithms decreases with lower resolution.

There is one more improvement in reciprocal space that is worth special attention. It is essentially an $F_o + \Delta F$ Fourier synthesis (Main, 1979) that relaxes the modulus constraint in step 3 of the iteration cycle. The operation $F(\mathbf{h}) = [2F_{obs}(\mathbf{h}) - |G(\mathbf{h})|] \exp[i\varphi_G(\mathbf{h})]$ gives the mirror image of each calculated structure factor relative to its projection on the circle of radius F_{obs}, and is a clear-cut case of positive feedback with a known reference point. For *ab initio* work, this step was first used as part of a very different scheme (Drendel *et al.*, 1995), and its reinvestigation was suggested to us by Thomas Weber. There are several advantages of the simplest $F_o + \Delta F$ version: (i) it does not discard any reflections, (ii) it does not add any new parameter to basic CF, and (iii) it converges at least as fast as the $\pi/2$ version. However, for larger structures, it is better to limit the reciprocal-space perturbation by a new parameter W. As shown in Fig. 6, the $F_o + \Delta F$ step is applied only when $G(\mathbf{h})$ falls within the ring $F_{obs}(\mathbf{h}) - W < |G(\mathbf{h})| < F_{obs}(\mathbf{h}) + W$, otherwise the mirror operation is limited by the opposite edge of the ring. This is made plausible by our observation that a solution with optimal δ always leads to a uniform misfit of calculated and observed structure-factor moduli. The optimal value of the normalized parameter $w = W/\max(F_{obs})$ is usually larger than 0.25.

7. Improvements in real space

Charge flipping is a local perturbation of low density that does not act at points where the value is above the threshold. Unfortunately, it can happen that at the beginning of the process high positive densities emerge incoherently, and it is practically impossible to remove them by flipping only $\rho < \delta$. These are the runs that end without success and just increase

the statistical cost of a single solution. As shown previously, various reciprocal-space perturbations are able to break this stagnation. The reason is that they act non-locally in real space, also in points of large (and possibly false) density. It is instructive to see whether the same result can be achieved by a modification that acts purely in real space.

We have tested many variants of the basic algorithm and found that it is much more difficult to find an improvement in real space than in reciprocal space. For the low-density part, we cannot find anything nearly as good as charge flipping. Low-density elimination (LDE) (Shiono & Woolfson, 1992) was considered as a particularly attractive alternative because it was the first method to utilize a small positive threshold δ and was already found useful to complete a CF solution. To make sure that we really probe the *ab initio* phasing power of its name-giving real-space part, we used high-resolution data ($d_{\min} = 0.8$ Å) and kept the simplest modulus constraint. It was found that LDE is extremely susceptible to stagnation, requires a δ parameter about twice as large as CF, and solves only very simple structures. Later, the same authors dropped the original version and opted for continuous functions (Refaat & Woolfson, 1993; Foadi *et al.*, 2000; Matsugaki & Shiono, 2001) that delete $\rho < 0$ and amplify $\rho > 0$. Then there is no threshold and the modification must be considered as a whole. While we do not question the utility of these procedures within their native schemes, we still find that their phasing power is weak – at least as part of the simplest Fourier cycle. In this case, creating small discontinuities of low density seems to be a more efficient way for the exploration of phase space.

After many attempts, we have discovered a purely real-space modification that gives a significant improvement of the basic CF algorithm. This is called 'flip-mem' because the improvement comes at the expense of requiring a short-term memory of the iteration process. In step 1 of the $(n + 1)$th cycle, the usual sign reversal is performed for $\rho^n < \delta$. For $\rho^n > \delta$, the density modification is calculated as $g^{n+1} = \rho^n + \beta(\rho^n - \rho^{n-1})$, where the useful range of the β parameter is 0.5–1.0. Thus, g^{n+1} is outside the interval formed by ρ^{n-1} and ρ^n, and always on the side of ρ^n. This is another example of a positive feedback, different from both simple perturbations and from the negative feedback used by Fienup's algorithm. As we shall see in the table, this purely real space modification performs comparably well to the reciprocal-space $\pi/2$ and $F_o + \Delta F$ variants. The advantage of CF with real-space memory is that it yields better quality histograms.

The last real-space modification listed here does not yield a speed-up for structure solution using X-ray data. Instead, it removes the positivity constraint of charge flipping. In step 1 of the iteration cycle, the $g = -\rho$ sign reversal is applied only for $-\delta < \rho < +\delta$, and leaves both large positive and large negative values unchanged. This is called the *band flipping* version of the algorithm which helps to reconstruct negative scattering densities. Its utility in structure solution of neutron diffraction data is discussed in detail in our recent publication (Oszlányi & Sütő, 2007).

8. Data resolution and sharpening

Charge flipping is based on the physical assumption that the unit cell is mostly empty, *i.e.* large electron density is concentrated in a small fraction of the space. To expose this property, the data must extend to sufficiently high resolution and the thermal vibrations must be small. In a favourable case, complete and accurate $F_{\mathrm{obs}}(\mathbf{h})$ are available up to $d_{\min} = 0.8$ Å, the global thermal parameter can be corrected to $B \approx 0$, and the electron density is represented on a grid with $\Delta r = 0.2$–0.4 Å spacing. While the efficiency of the CF algorithm could really benefit from data of ultra-high resolution, this is seldom available. Instead, we must often work with significantly fewer data; decreasing the resolution to 1.0, 1.2 or 1.5 Å means that only a 51, 30 or 15% fraction of the standard $d_{\min} = 0.8$ Å reflection sphere is measured. At any resolution, (i) we must decide how to handle unobserved reflections, (ii) we must choose between natural and normalized structure factors, and (iii) it may become necessary to replace missing data by other pieces of known information. Points (i) and (ii) are discussed here, point (iii) in the next section.

Let us recall how the basic algorithm implements the modulus constraint. It replaces $|G(\mathbf{h})|$ by $F_{\mathrm{obs}}(\mathbf{h})$ for observed reflections, it resets $F(\mathbf{h}) = 0$ for unobserved reflections outside the resolution sphere, and accepts $F(0) = G(0)$ as calculated.

$F(0)$ is the sum of the total scattering density and determines the zero level of the flip. As such, its selection is more of a question of the algorithm than of data. $F(0)$ is usually initialized at zero, quickly reaches stagnation and drops at the convergence. The main role of freely changing $F(0)$ is not to find the true value of the total charge but to find the middle of the band that can be flipped without significantly changing the distribution of pixels. While we still consider unconstrained $F(0)$ the most useful variant, there is no clear consensus among developers. For example, Palatinus first fixed it at zero for the complete iteration process (Palatinus, 2004), later used it only for refinement, but his first version is preferred by Coelho when using the parameter ϕ and phase constraints (Coelho, 2007b). The choice of handling unobserved data outside the resolution sphere is more straightforward. With a proper match of resolution and grid spacing (reciprocal-space sphere contained by a parallelepiped), the number of unobserved reflections is large, at least as many as the number of observed reflections. These cannot be allowed to change freely because the iteration would lock in a false state without enough constraints. Prescribing upper limits and/or adding a thin unobserved shell is a possible choice, but the least problematic is to limit the search space and reset all $d < d_{\min}$ reflections to zero in each iteration cycle. There are many other types of missing data within the resolution sphere and the maximum-entropy optimized Patterson map (Palatinus *et al.*, 2007) gives a good estimate for these moduli. Extending data to higher than experimental resolution is a different problem that badly needs additional structural information.

Considering that most direct methods operate with normalized structure factors (E's instead of F's), it is a natural

feature articles

Table 1
The average number of iterations required for a single solution using various modifications of the charge flipping algorithm.

Columns: resolution and type of data. Rows: algorithm variants and the number of iteration cycles required for a single solution at the given data. The test structure and algorithm variants are described in the text. Parameters: α is the fraction of reflections considered weak, β is the amplification factor of the positive feedback and w is the relative width of the reciprocal ring where $F_{obs} + \Delta F$ is applied.

	F at 0.8 Å	E at 0.8 Å	F at 1.0 Å	E at 1.0 Å
Basic charge flipping	$>8 \times 10^6$	32000	$>8 \times 10^6$	600000
Weak = 0, α = 0.20	8×10^6	11000	$>8 \times 10^6$	210000
Weak = 0, α = 0.40	4×10^6	1700	$>8 \times 10^6$	32000
Weak = 0, α = 0.60	3×10^6	500	8×10^6	5400
$\Delta\varphi = \pi/2$, α = 0.20	190000	1500	1×10^6	10000
$\Delta\varphi = \pi/2$, α = 0.40	140000	1300	790000	2800
$\Delta\varphi = \pi/2$, α = 0.60	210000	4400	8×10^6	2000
$F_{obs} + \Delta F$, w = 0.25	59000	750	190000	3400
$F_{obs} + \Delta F$, w = 0.50	42000	1600	120000	900
$F_{obs} + \Delta F$, $w = \infty$	48000	2600	130000	650
flip-mem, β = 0.6	1×10^6	1400	8×10^6	10000
flip-mem, β = 0.8	220000	2500	4×10^6	3400
flip-mem, β = 1.0	40000	9700	460000	3100

question whether their use leads to faster convergence also with the charge flipping algorithm. The answer is yes and we can easily see why. Let us assume that the structure contains only a single type of atom with scattering factor $f(\mathbf{h})$. If structure-factor moduli are replaced by $F_{obs}(\mathbf{h})/f(\mathbf{h})$ then we get the point-atom representation of the same structure. Peaks of the electron-density map become higher and concentrated into a smaller number of pixels – this is called sharpening, and means a smaller search space for the algorithm. Unfortunately, artificial negative density – which always comes from resolution cut-off – is also enhanced with sharpening. At each charge flipping cycle, large negative samples must change their sign, which is unnecessary and can be harmful. The two effects of using E's are always coupled. Many tests indicate that with data at high resolution the positive effect dominates and the only harm is the smaller drop of the R factor at convergence.

The standard definition of normalized structure factors is $E(\mathbf{h}) = F_{obs}(\mathbf{h})/[\sum_j f_j^2(\mathbf{h})]^{1/2}$. While charge flipping can solve many structures without them, their use can significantly speed up convergence for large structures. Apart from the original definition, we have also experimented with two other variants of sharpening. The first one can be thought of as the $n \rightarrow \infty$ limit of the generalized $E^n(\mathbf{h}) = F_{obs}(\mathbf{h})/[\sum_j f_j^n(\mathbf{h})]^{1/n}$ expression. This corresponds to the scaling $F_{obs}(\mathbf{h})/f_H(\mathbf{h})$, where $f_H(\mathbf{h})$ is the scattering factor of the heaviest atom found in the structure. The second one corresponds to neutron-like scattering factors, where $f_j(\mathbf{h})$ is replaced by the constant $f_j(0)$ for every atom. This is a hypothetical case that requires recalculated synthetic data and cannot be used as a correction for real data. Still, it is instructive to see how it behaves because using $f_j(0)$ means more sharpening than what comes from the standard definition. Tests of several different structures can be

summarized by the following statements: (i) for light-atom structures all three versions of sharpening give similar results; (ii) for structures containing heavy atoms sharpening by $f_H(\mathbf{h})$ gives just as good results as standard E's, while ideal point atoms are much worse; (iii) so, independently of the structure type, knowing the heaviest atom can give sufficient sharpening even if the chemical composition is not available. Dividing by $f_H(\mathbf{h})$ enhances the contrast between heavy and light atoms, by sharpening the former and damping the latter, it decreases the relative contribution of many atoms.

The most attractive property of the use of E's is that all previous improvements of the algorithm are compatible with it. Any improvement obtained with F's can be combined with E's to give an even better performance. A fairly complete set of results is compiled in Table 1 for a carefully selected organic structure (Alexander *et al.*, 2002) [219 non-H atoms per unit cell, space group $P1$, a hard-to-solve example of our earlier paper (Oszlányi & Sütő, 2005)]. This contains the cost of a single *ab initio* solution using all previously discussed (suggested) algorithm variants, with ideal F or E data, and at 0.8 and 1.0 Å resolution. The statistics of 100–400 runs were calculated at each parameter, and each run was continued for a maximum of 20000 cycles. The threshold parameter of different algorithms was carefully optimized and only the best results were included in Table 1. The choice of $\Delta\varphi = \pi/2$ is not optimal, better results were obtained with values higher by 10–20°.

The most remarkable observation is that the efficiency of closely related algorithms differs so much – at least by four orders of magnitude. The other observation is less surprising. The type of data and the resolution strongly and generally affect the efficiency: E's are much better than F's and 0.8 Å resolution is better than 1.0 Å. Otherwise, these algorithm variants can often perform comparably well after some optimization: the cost of a solution may differ only by a factor of 2–5. Probably, it is more important to pick one of the suggested algorithms and spend some time finding a good parameter. At this point, there is no simple recipe for choosing parameters, mainly because the practical measure of success is not the number of solutions per unit time once the best parameter is found but the total time spent on finding a reasonable parameter plus the first solution. Experience in this respect is being accumulated. A good parameter must bring sufficiently strong perturbation to explore the phase space but not too strong, otherwise the iteration will miss the solution. We have also tried more complicated combinations of algorithm variants. It seems that no significant speed-up can be achieved in this way, the sum of useful perturbations is somehow limited by the relatively weak constraints used so far. We have probably reached the point where it is appropriate to add more constraints that may allow stronger perturbations and a further speed-up of convergence. A final word of caution: the speed of finding a solution is not everything. The quality of the electron-density map is just as important and should always be improved by switching back to natural F_{obs}, adding a few iteration cycles of the basic algorithm, plus the final clean-up procedure.

9. Use of known information

Charge flipping is a flexible tool that can be used in different ways and at different stages of the structure-solution process. It either operates in a truly *ab initio* manner without relying on preliminary information or can be applied to complete a partially known structure/phase set, it is able to check the stability of a solution but it can also be adapted to work as an ingredient of other dual-space schemes. In previous publications, we mostly emphasized the *ab initio* nature of the algorithm and gave only hints of its practical utility for structure completion. Here this option is discussed in more detail.

Known structural information might be diverse. To list a few possibilities: atomicity, atom types, chemical composition, histogram of electron density, symmetry elements, full space-group symmetry, connectivity of a molecular structure, common polyhedra of an extended solid, the positions of heavy atoms, the positions of some fragments of the structure, two-dimensional projections of the correct electron density, low-resolution phases, structures of related compounds, statistical phase relations *etc.* and any mix of these. Whether known information is required at all depends on both the complexity of the structure and the amount of available data. With simple structures containing heavy atoms, structure determination may already succeed with powder data of average quality, while for large and complex structures even high-resolution single-crystal data may be insufficient. Even in those cases when required and available, not all types of known structural information can be utilized easily and combined with others. A good part of the crystallographic literature deals with the question of how to use various pieces of known information in an optimal way. Here we pick only two simple constructs that can easily be implemented.

The first one assumes that some other method has already located a fraction of the structure with reasonable accuracy, *e.g.* the position of heavy atoms, an oriented fragment or a typical layer. Then the corresponding electron density can be taken as a good starting point in real space that replaces the usual random-phase initialization of the CF algorithm. Afterwards, the original iteration process can proceed without any modification. This means that no attempt is made to fix the model, it either evaporates or develops into the complete structure within a small number (10–100) of cycles. For structure completion, the threshold parameter δ must be set slightly smaller than what is required for an *ab initio* solution. In the case of failure, multiple runs may also be attempted but then some noise must be added to the starting density so that individual runs differ. As might be expected, with a larger fraction of the correct electron density, the chance of a successful structure completion increases. However, using high-resolution data it can already succeed with as little as 0.5–1.0% of the total density if the known part is carried by heavy atoms, or with 2–4% if carried by light-atom fragments. To give an impression of the efficiency of structure completion by charge flipping: vitamin B12, rubredoxin and lysozyme can be

solved with 0.8–1.2 Å data and starting from the known positions of heavy atoms.

The second construct works in reciprocal space and requires the knowledge of both the modulus and phase for a subset of reflections. These are used not only for initialization but are kept as constraints all through the iteration process. So this second construct is more strict than the first one and can also succeed with a smaller amount of known information. Here we mention an important aspect of algorithm implementation. Our program used for development has a large degree of freedom of deciding how individual reflections are treated. It can assign separate 'how-to-use' codes for each reflection that determines in each cycle whether the reflection is reset to zero, allowed to change freely, its modulus is constrained, both its modulus + phase are constrained, its phase is shifted by $\Delta\varphi$, it is used in $F_o + \Delta F$ mode, it follows phase constraints, or any other option. Reflection groups may also be created and be given the same code. The groups can be created on the basis of the measured magnitude, the resolution shell, symmetry equivalents, available phase information *etc*. The program might even modify its strategy and change the corresponding codes after a number of iteration cycles. A practical case where the simplest strategy works is that of classical phase extension: let us assume that both modulus and phase is known for $d > 3$ Å, only the modulus is known for $d = 1.5$–3 Å and there is no measured data below $d = 1.5$ Å. In other words, 15% of the standard 0.8 Å sphere has known moduli, and only 1.9% has known phases. With reasonable phase error (uniform $\Delta\varphi = [-45°, +45°]$), phase extension from 3.0 to 1.5 Å seems to be an easy exercise for the previous examples of vitamin B12, rubredoxin and lysozyme.

Not all pieces of structural information can be used as simply as in the previous constructs, *i.e.* only at the start or perpetually in each iteration cycle. Often it is better to prescribe constraints (symmetry, atomicity, histogram) only partially and/or only once in every nth iteration cycle, where $n = 10$–50. Also, there exist other pieces of known information (distance, connectivity, polyhedra) that cannot be accommodated in the charge flipping scheme so far. To find the most general and most efficient options of using known information certainly needs further investigation.

10. Applications and user programs

In this section, we briefly discuss the real-life problems that charge flipping is suited for, and the user programs where the algorithm can be put into action. We emphasize that most of the results presented here were achieved by others, a small enthusiastic part of the crystallographic community, who picked up the method early and developed it in a creative way.

For periodic single crystals measured with X-radiation, the utility of the CF algorithm was clear at the time of the initial discovery (Oszlányi & Sütő, 2004). The first demonstration using experimental data followed quickly (Wu *et al.*, 2004), proving the initial claim that the CF algorithm is reasonably tolerant to imperfections of data. However, the heavy-atom structures used in the first demonstration led to high noise

levels of the electron-density map and it was necessary to average multiple runs to improve the map and identify light atoms. This is a good general technique but today we know that the clean-up procedure would do an even better job at the end of a single run, and the option of using multiple runs would still be open. Since then, several unknown structures have successfully been solved by charge flipping (Meffre *et al.*, 2007; Richeter *et al.*, 2007; van der Lee & Astier, 2007; Wardell, Low & Glidewell, 2007; Wardell, Wardell *et al.*, 2007) but, as current direct-methods programs perform so well, using the CF method remained mostly a curiosity.

A more challenging example of structure solution using single-crystal data came as an awkward case of pseudosymmetry (Oszlányi *et al.*, 2006): a large non-centrosymmetric structure where six molecules are stacked on top of each other with an approximate translational symmetry of every second molecule. This nearly regular distribution of atoms is rather unfavourable for classical direct methods because statistical phase relations rely on the random distribution of atomic positions. In contrast, the same atomic arrangement is a big advantage for charge flipping because it amplifies a few Fourier components and decreases the search space. Since the first demonstration, some even larger and otherwise problematic structures were solved by the CF algorithm with surprising ease (Evans, 2007). It is now certain that categories of easy- and hard-to-solve structures can be very different for classical direct methods and CF, and the solution of pseudosymmetric structures will be a main application where charge flipping can complement other direct methods.

Whatever its advantages, for periodic crystals charge flipping may be considered 'just another method' of structure determination. The situation for the field of aperiodic crystals is quite different because it cannot utilize three-dimensional atomicity and has long needed a general and comprehensible method of structure determination. It was Palatinus who realized that, when charge flipping develops extended zero plateaus of the electron density and finds the atoms, it does this without relying on the concept of three-dimensional atomicity. Furthermore, the CF algorithm contains no three-dimensional specific part, it can work in any dimension. So he extended the method to higher dimensions and succesfully solved modulated structures directly in superspace (Palatinus, 2004). The first demonstration already included examples that rank among the most complex modulated structures described so far. Since then, several unknown structures have also been solved successfully (Zúñiga *et al.*, 2006; Palatinus *et al.*, 2006). A further important development in superspace applications is the first structure solution of a decagonal quasicrystal with five-dimensional charge flipping (Katrych *et al.*, 2007). While the superspace descriptions of modulated structures and quasicrystals differ, it seems that charge flipping is generally applicable to both structure types.

When single crystals are not available, structure solution may be attempted using powder diffraction data. Here reflection overlap is a special problem that usually requires the repartitioning of measured intensities. Often the choice of the space group is also uncertain, so a method that works in $P1$ has

clear advantages. The first demonstration of powder charge flipping (pCF) came from the Arizona group (Wu *et al.*, 2006), they included a Le Bail-like repartitioning of intensities in each iteration cycle and solved some relatively simple structures. A different repartitioning scheme based on histogram matching was introduced by Baerlocher, McCusker & Palatinus (2007) that significantly increased the complexity of known structures that can be solved by pCF. Later, the same group solved an 864 atom zeolite catalyst structure that had resisted all previous attempts for a decade (Baerlocher, Gramm *et al.*, 2007). For this spectacular result, it was necessary to combine powder data with other pieces of known information: two-dimensional projections of the electron density measured by high-resolution electron microscopy. It is now clear that charge flipping will be a useful method for solving structures from high-resolution powder diffraction data. General improvements of the CF algorithm with low-resolution data and specific improvements of the pCF repartitioning scheme are still required.

Finally, we list some user programs where the CF algorithm is already implemented. These are: *BayMEM*, *SUPERFLIP*, *PLATON*'s *FLIPPER* module, *TOPAS Academic* and Bruker-AXS *TOPAS*. The program *BayMEM* (van Smaalen *et al.*, 2003) was the host of the first implementation used for solving aperiodic structures. Later, Palatinus and Chapuis created a dedicated program *SUPERFLIP* (Palatinus & Chapuis, 2007) that works with arbitrary dimensions of electron density and data. It also includes a powder option, has strong symmetry analysis capabilities and good connections to other superspace programs like *JANA2000* and *BayMEM* (Petříček *et al.*, 2000; van Smaalen *et al.*, 2003). The third implementation is the module named *FLIPPER* in Ton Spek's well known *PLATON* suite (Spek, 2003) that is designed for more automatic work. Its special functionality is that by using *PLATON*'s other tools it gradually determines the space group and atom types as the iteration process progresses. The latest implementation is *TOPAS Academic* (Coelho, 2007*a*) and the corresponding version of *TOPAS* (Bruker-AXS, 2007). Their common kernel is a programmable program written for periodic structures and with many new variants of the algorithm. Perhaps the most significant improvement developed here is the combination of CF with the tangent formula (Coelho, 2007*b*). This new constraint allows stronger perturbations that lead to a remarkable speed-up of convergence. It is true for all the listed user programs that they are being developed quickly and their current status can be checked only in the user manual. Also, their authors are responsive both to error reports and to requests to include further functionality.

11. Concluding remarks

In this paper, we have attempted to survey a still extending horizon of improvements and applications of the charge flipping algorithm. Finally, we would like to stress that, in spite of the spectacular speed-up that can be reached by some algorithm variants, the heart of the method remains the

name-giving real-space transformation. When we write in more than one place about its 'surprising' simplicity, we admit our own lack of understanding of it beyond the level of intuition. Deeper insight might come from a still missing mathematical proof of convergence of the algorithm.

It has been known for a long time that, in the case of sufficiently high resolution data, scattering intensities contain more than enough information to reproduce the electron density or, equivalently, the missing phases. Atomic scattering factors can be described with nine parameters. Adding three coordinates and, in the worst case, six anisotropic thermal parameters, we have 18 parameters per atom. With a few hundred atoms per unit cell, the electron density is then coded by several thousand real parameters, while at 0.8 Å resolution the number of available Fourier moduli is roughly ten times as much. An algebraic extraction of the electron density from nothing but the known scattering intensities is hopeless for systems of this size; the only imaginable way is *via* an iterative algorithm.

Charge flipping appears to be an algorithm that can do this job. It successfully combines a relaxed positivity constraint ($\rho > -\delta$) with a perturbation of low densities (sign change for $\rho < \delta$). Its strength is the decomposition of the electron density in the sum of two parts, ρ_1 and ρ_2. When the corresponding Fourier transform $F_1 + F_2$ is changed into $F_1 - F_2$, the direction of the complex vectors changes considerably, making an efficient exploration of the space of phases possible. It is to be mentioned that charge flipping is not a projection, not even a relaxed over-projection (Stark & Yang, 1998). Its fixed points, the real N-dimensional vectors (N here is the number of pixels) with components either 0 or $\geq \delta$ form a non-convex disconnected manifold. Application of charge flipping to other than fixed points may increase the distance to this manifold
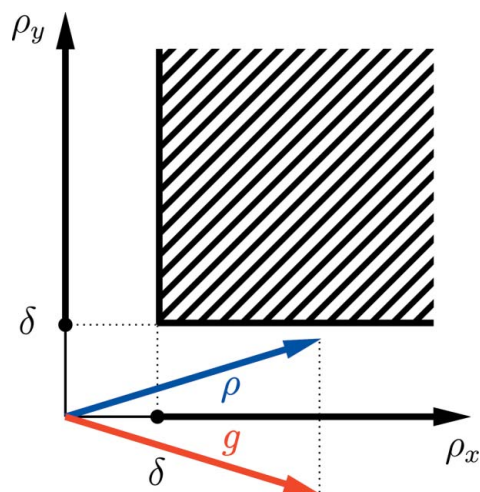
(see Fig. 7) while (relaxed) projections always decrease the distance.

There are some typical ways to improve the basic algorithm. The earliest among them was the special treatment of weak reflections, either by resetting them to zero or by letting them freely evolve and applying a $\pi/2$ phase shift to them. Both introduce a perturbation in the algorithm, associated with a relaxation of the modulus constraint on weak reflections. As we understand it, they accelerate convergence defensively, not by driving the system towards a solution but by preventing it settling down in a trap. This is in contrast with the two other accelerators introduced in this paper, the real-space flip-mem and the reciprocal-space $F_o + \Delta F$. The first one uses a one-step memory while the second works without memory, and both realize a positive feedback. Because $F_o + \Delta F$ concerns the data, it can be considered as a perturbation or relaxation of the modulus constraint. However, it can also be viewed as a pure real-space intervention which replaces ρ^n by $2\rho^n - g^n$ everywhere and prior to charge flipping in the next cycle. As such, it resembles the $\beta = 1$ version of flip-mem, and it is no surprise that the two variants perform similarly, *cf*. Table 1. Note also their difference arising from a different intuition behind them, and that $F_o + \Delta F$ performs distinctly better – it would be hard to tell why.

The above improvements do not affect the *ab initio* character of the CF method. A major improvement with only a slight deviation from being *ab initio* comes from the use of normalized structure factors. Table 1 shows convincingly its superiority compared with the use of unnormalized structure factors. One may ask how far in resolution and system size it is possible to go with the CF algorithm and the use of E's. Based on preliminary computations not reported here, we think that light-atom structures of near-protein-size should be possible to handle up to 1.0–1.2 Å resolution, while for smaller heavy-atom structures the resolution limit might be at 1.5 Å. By application of the tangent formula or other phase relations and constraint of the algorithm with known structural information, these limits can be pushed even farther. We hope to assist and perhaps contribute to these future developments of the charge flipping algorithm.

References

Abrahams, J. P. & Leslie, A. G. W. (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.

Baerlocher, Ch., Gramm, F., Massüger, L., McCusker, L. B., He, Z., Hovmöller, S. & Zou, X. (2007). *Science*, **315**, 1113–1116.

Baerlocher, Ch., McCusker, L. B. & Palatinus, L. (2007). *Z. Kristallogr.* **222**, 47–53.

Barrett, A. N. & Zwick, M. (1971). *Acta Cryst.* A**27**, 6–11.

Bruker-AXS (2007). *TOPAS User's Manual.*

Burla, M. C., Caliandro, R., Camalli, M., Carrozzini, B., Cascarano, G. L., De Caro, L., Giacovazzo, C., Polidori, G. & Spagna, R. (2005). *J. Appl. Cryst.* **38**, 381–388.

Cochran, W. (1952). *Acta Cryst.* **5**, 65–67.

Figure 7
The real space of the two-pixel charge flipping algorithm. The non-convex manifold of fixed points is the union of two half-infinite lines and a quarter-plane. Blue and red arrows correspond to an example of the CF transformation, $(\rho_x, \rho_y) \rightarrow (\rho_x, -\rho_y)$, which increases the distance to the manifold.

feature articles

Coelho, A. A. (2007*a*). *TOPAS Academic 4.1 Technical Reference*, http://members.optusnet.com.au/alancoelho.

Coelho, A. A. (2007*b*). *Acta Cryst.* A**63**, 400–406.

Cooley, J. W. & Tukey, J. W. (1965). *Math. Comput.* **19**, 297–301.

Drendel, W. B., Davé, R. D. & Jain, S. (1995). *Proc. Natl Acad. Sci. USA*, **92**, 547–551.

Evans, J. S. O. (2007). Personal communication.

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

Florke, U. & Haupt, H. J. (1990). *Z. Kristallogr.* **193**, 309.

Foadi, J., Woolfson, M. M., Dodson, E. J., Wilson, K. S., Yao, J. X. & Zheng, C. D. (2000). *Acta Cryst.* D**56**, 1137–1147.

Frigo, M. & Johnson, S. G. (2005). *Proc. IEEE*, **93**, 216–231.

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

Giacovazzo, C. (1998). *Direct Phasing in Crystallography*. Oxford University Press.

Hauptman, H. & Karle, J. (1953). *Am. Crystallogr. Assoc. Monogr. No. 3. Solution of the Phase Problem. I. The Centrosymmetric Crystal*. Michigan: American Crystallographic Association.

Hoppe, W. & Gassmann, J. (1968). *Acta Cryst.* B**24**, 97–107.

Irngartinger, H., Weber, A. & Oeser, T. (1999). *Angew. Chem. Int. Ed. Engl.* **38**, 1279–1281.

Karle, J. & Hauptman, H. (1956). *Acta Cryst.* **9**, 635–651.

Katrych, S., Weber, Th., Kobas, M., Massüger, L., Palatinus, L., Chapuis, G. & Steurer, W. (2007). *J. Alloys Compd.* **428**, 164–172.

Lee, A. van der & Astier, R. (2007). *J. Solid State Chem.* **180**, 1243–1249.

Lunin, V. Yu. (1988). *Acta Cryst.* A**44**, 144–150.

Main, P. (1979). *Acta Cryst.* A**35**, 779–785.

Matsugaki, N. & Shiono, M. (2001). *Acta Cryst.* D**57**, 95–100.

Meffre, A., Barboiu, M. & van der Lee, A. (2007). *Acta Cryst.* E**63**, m255–m257.

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.

Oszlányi, G. & Sütő, A. (2005). *Acta Cryst.* A**61**, 147–152.

Oszlányi, G. & Sütő, A. (2007). *Acta Cryst.* A**63**, 156–163.

Oszlányi, G., Sütő, A., Czugler, M. & Párkányi, L. (2006). *J. Am. Chem. Soc.* **128**, 8392-8393.

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

Palatinus, L. & Chapuis, G. (2007). *J. Appl. Cryst.* **40**, 786–790.

Palatinus, L., Dušek, M., Glaum, R. & El Bali, B. (2006). *Acta Cryst.* B**62**, 556–566.

Palatinus, L., Steurer, W. & Chapuis, G. (2007). *J. Appl. Cryst.* **40**, 456–462.

Petříček, V., Dušek, M. & Palatinus, L. (2000). *JANA2000. The Crystallographic Computing System*. Institute of Physics, Praha, Czech Republic.

Refaat, L. S. & Woolfson, M. M. (1993). *Acta Cryst.* D**49**, 367–371.

Richeter, S., Hadj-Aïssa, A., Taffin, C., van der Lee, A. & Leclercq, D. (2007). *Chem. Commun.* **21**, 2148–2150.

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

Sayre, D. (2002). *Struct. Chem.* **13**, 81–96.

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

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

Smaalen, S. van, Palatinus, L. & Schneider, M. (2003). *Acta Cryst.* A**59**, 459–469.

Spek, A. L. (2003). *J. Appl. Cryst.* **36**, 7–13.

Stanley, E. (1979). *Acta Cryst.* A**35**, 966–970.

Stanley, E. (1986). *Acta Cryst.* A**42**, 297–299.

Stark, H. & Yang, Y. (1998). *Vector Space Projections*. New York: Wiley.

Stumpf, K., Blachnik, R. & Roth, G. (1999). *Z. Kristallogr.* **214**, 251–254.

Wang, B. C. (1985). *Diffraction Methods for Biological Macromolecules*, edited by H. W. Wyckoff, C. H. W. Hirs & S. N. Timasheff, Vol. 115, pp. 90–113. Orlando: Academic Press.

Wardell, J. L., Low, J. N. & Glidewell, C. (2007). *Acta Cryst.* E**63**, o1848–o1850.

Wardell, S. M. S. V., Wardell, J. L., Low, J. N. & Glidewell, C. (2007). *Acta Cryst.* E**63**, o970–o971.

Woolfson, M. M. (1987). *Acta Cryst.* A**43**, 593–612.

Wu, J. S., Leinenweber, K., Spence, J. C. H. & O'Keeffe, M. (2006). *Nature Materials*, **5**, 647–652.

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

Yao, J. X., Dodson, E. J., Wilson, K. S. & Woolfson, M. M. (2006). *Acta Cryst.* D**62**, 901–908.

Zhang, K. Y. J., Cowtan, K. D. & Main, P. (2001). *International Tables for Crystallography*, Vol. F, edited by M. G. Rossmann & E. Arnold, pp. 311–324. Dordrecht: Kluwer Academic Publishers.

Zhang, K. Y. J. & Main, P. (1990). *Acta Cryst.* A**46**, 41–46.

Zúñiga, F. J., Palatinus, L., Cabildo, P., Claramunt, R. M. & Elguero, J. (2006). *Z. Kristallogr.* **221**, 281–287.