

Solution of the crystallographic phase problem by iterated projections

Veit Elser

Department of Physics, Cornell University, Ithaca, NY 14853-2501, USA. Correspondence e-mail: ve10@cornell.edu

An algorithm for determining crystal structures from diffraction data is described which does not rely on the usual reciprocal-space formulations of atomicity. The new algorithm implements atomicity constraints in real space, as well as intensity constraints in reciprocal space, by projections that restore each constraint with the minimal modification of the scattering density. To recover the true density, the two projections are combined into a single operation, the *difference map*, which is iterated until the magnitude of the density modification becomes acceptably small. The resulting density, when acted upon by a single additional operation, is by construction a density that satisfies both intensity and atomicity constraints. Numerical experiments have yielded solutions for atomic resolution X-ray data sets with over 400 non-hydrogen atoms, as well as for neutron data, where positivity of the density cannot be invoked.

© 2003 International Union of Crystallography
Printed in Great Britain – all rights reserved

1. Introduction

This year marks the semi-centennial of the realization by crystallographers that diffraction intensities possess sufficient information to reconstruct an atomistic structure (Sayre, 2002). The simple fact that the scattering arises from a known number of nearly point-like entities, while clearly not as intricate in content as the body of collected intensities itself, constitutes a significant piece of information as well. The first important steps in utilizing atomicity in structure determination were taken by Sayre (1952) in his celebrated equation, and later by Hauptman & Karle (1953) in their probabilistic analysis of structure factors. These pioneering methods were based in reciprocal space and established a framework for all subsequent work on the crystallographic phase problem.

What is remarkable in the fifty-year history of direct methods is that crystallographers have been slow to exploit the efficiency of the fast Fourier transform (FFT) to impose atomicity in real space, the space where it is most naturally expressed. In the first application of the FFT to phase extension and refinement, Barrett & Zwick (1971) viewed the transformation to real space primarily as an expedient for carrying out the convolution in Sayre's equation (and related tangent formula). These authors also made the observation that, with easy access to the real-space density, a variety of phase refinements is available that are outside the scope of the reciprocal-space approach. Density modification, as these real-space-based methods became known, fell short of atomicity in that only global characteristics of the density, such as non-negativity, were addressed. True atomicity constraints in real space first appeared in the *SnB* program (Miller *et al.*, 1994; Weeks & Miller, 1999), and were later adopted by *SHELXD* (Sheldrick, 1997, 1998). However, even these

programs continue to rely heavily on the traditional reciprocal-space approach to atomicity

The aim of the work detailed below was to develop a practical phase-determination algorithm for crystal structures which imposes atomicity entirely within real space. A key component of the algorithm is an iterative operation (difference map) that was discovered by deconstructing the most successful algorithm (hybrid input–output) for the phase problem in optics (Fienup, 1982) and re-expressing it in terms having wider applicability (Elser, 2003a). Experiments with the *atom_retriever* implementation of the new algorithm on a variety of test structures demonstrate both its robustness and its speed. The flexibility of the new approach, with respect to the kinds of constraints that can be imposed in real space, raises hopes of an *ab initio* solver not limited by atomic resolution data.

2. Constraints and projections

The choice of the algorithm's fundamental variables is largely motivated by the mathematical structure of the iterative step (§3). In particular, the object that is iterated should have the property that it can be added, in the sense of a linear vector space, to other objects, and that there is a natural expression for the distance between objects. The set of structure factors, with magnitudes fixed and phases variable, are not good candidates in this respect. A better choice of object, and the one we adopt, is the real-space scattering density sampled on a finite regular grid. The relationship between real-space sampling and reciprocal-space sampling on the reciprocal lattice is quite direct, as illustrated by the two-dimensional example in Fig. 1. Shown on the left (Fig. 1a) is the actual

Table 1

Test structures solved by *atom_retriever*.

Structure	Reference	Data	Reflections	<i>N</i>	Group	<i>d</i> _{min}	Completeness	<i>b</i>	Grid	β	Iterations	(cos $\Delta\varphi$)
Montebrasite	(a)	Neutron	1024	20	<i>P1</i> †	0.66	85%	2.5	22 × 24 × 22	0.7	20	0.76
TEX nitramine	(b)	X-ray	1911	36	<i>P1</i> †	0.76	91%	6.6	20 × 22 × 24	0.7	15	0.77
Amphiphilic pyrrole	(c)	X-ray	4238	48	<i>P1</i> †	0.81	100%	11.0	14 × 26 × 52	0.7	25	0.76
Punctaporonin D	(d)	X-ray	3764	54	<i>P1</i>	0.85	100%	10.4	22 × 26 × 30	0.7	75	0.72
α -helix and 3_{10} / α -helix	(e)	X-ray	7489	148	<i>P1</i>	0.96	100%	7.5	22 × 36 × 42	0.7	1100	0.71
Triphenylphosphine	(f)	X-ray	9982	152	<i>P1</i> †	0.84	100%	2.9	28 × 36 × 42	0.7	25	0.79
α -helical bundle	(g)	X-ray	23681	479	<i>P1</i>	0.90	86%	10.1	40 × 46 × 60	0.7	450000	0.75

References: (a) Groat *et al.* (2002); (b) Karaghiosoff *et al.* (2002); (c) Ramos Silva *et al.* (2002); (d) Poyser *et al.* (1986); (e) Karle *et al.* (1989); (f) Kooijman *et al.* (1998); (g) Privé *et al.* (1999). † The actual structure has symmetry *P*1̄, but was treated as *P1* with twice as many atoms.

scattering density within one unit cell. The structure factors of the corresponding crystal decay with scattering angle so that only a limited range about the origin in reciprocal space is measured. By padding with zeros at the corners, the measured structure factors can be fit into a finite rectangular grid as shown in the middle figure (Fig. 1*b*). Given phases for the structure factors on the bounded reciprocal-space grid, the discrete Fourier transform of the resulting complex structure factors then gives the discretely sampled real-space density shown on the right (Fig. 1*c*). Conversely, given a scattering density on the real-space grid (Fig. 1*c*), the inverse Fourier transform gives the complex structure factors on the bounded reciprocal-space grid, although with magnitudes not necessarily matching the measurements (Fig. 1*b*).

2.1. Intensity constraints

A valid density in real space must first of all have the property that the inverse Fourier transform gives the measured structure-factor *magnitudes*. If not, one can seek the minimal density modification that brings the actual magnitudes into agreement with the measured ones. Using the symbol ρ to represent the vector of densities on the real-space grid, the mathematical operation which accomplishes this is the projection $\Pi_F(\rho)$. The projected density is uniquely defined by the properties that its Fourier transform has the correct (given) magnitudes and the distance $\|\Pi_F(\rho) - \rho\|$ is minimized. It is convenient to use the Euclidean distance since it is preserved by Fourier transformation:

$$\|\rho\|^2 = \sum_{\mathbf{r}} \rho_{\mathbf{r}}^2 = \sum_{\mathbf{q}} |\tilde{\rho}_{\mathbf{q}}|^2 = \|\tilde{\rho}\|^2. \quad (1)$$




Figure 1

Discrete sampling of a scattering density in two dimensions. (a) Density within one unit cell. (b) Structure-factor magnitudes for the density in (a), sampled over a range with $h^2 + k^2 < 10^2$. (c) Discrete sampling of (a) obtained by applying the discrete Fourier transform to the *complex* structure factors [(b) combined with phases].

In (1), the indices \mathbf{r} and \mathbf{q} denote grid points in real space and reciprocal space, respectively, and the complex structure factors $\tilde{\rho}_{\mathbf{q}}$ are related to the real-space density by

$$\tilde{\rho}_{\mathbf{q}} = [\mathcal{F}(\rho)]_{\mathbf{q}} = (1/M^{1/2}) \sum_{\mathbf{r}} \exp(2\pi i \mathbf{q} \cdot \mathbf{r}) \rho_{\mathbf{r}}, \quad (2)$$

where *M* is the total number of grid points (in real or reciprocal space). Using the unit cell's fractional coordinates (*x*, *y*, *z*) to label \mathbf{r} and the Miller indices (*h*, *k*, *l*) for \mathbf{q} , we have $\mathbf{q} \cdot \mathbf{r} = hx + ky + lz$. The invariance of the distance (1) makes it possible to achieve the distance-minimizing property of the projection Π_F very easily in reciprocal space. Specifically, for every complex structure factor $\tilde{\rho}_{\mathbf{q}}$, we wish to find the nearest point in the complex plane lying on a circle corresponding to the measured magnitude $F_{\mathbf{q}}$. The required projection is therefore accomplished by simply rescaling the magnitude of the structure factor:

$$\left[\tilde{\Pi}_F(\tilde{\rho}) \right]_{\mathbf{q}} = \frac{F_{\mathbf{q}}}{|\tilde{\rho}_{\mathbf{q}}|} \tilde{\rho}_{\mathbf{q}}. \quad (3)$$

A vanishing denominator in (3) is not a problem since it represents a set of measure zero if one neglects extinctions (where $F_{\mathbf{q}}$ also vanishes). The projection in real space is expressed symbolically as

$$\Pi_F = \mathcal{F}^{-1} \tilde{\Pi}_F \mathcal{F}. \quad (4)$$

When the Fourier transforms are implemented with the FFT, the computational cost of projecting a density on the intensity constraints grows as *M* log *M*.

A practical algorithm using the Fourier intensity projection Π_F must address the fact that not all the structure factors on the reciprocal-space grid will be measured. In addition to $F_{\mathbf{0}}$, measurements near $\mathbf{q} = \mathbf{0}$ will frequently be absent or very unreliable, particularly for large-unit-cell crystals. At the other extreme, structure factors in the corners of the reciprocal-space grid (Fig. 1*b*) will be absent because data are normally collected within a spherical domain about the origin. The absent structure factors can be treated in a uniform way by applying *bound* constraints rather than value constraints. A bound $|\tilde{\rho}_{\mathbf{q}}| < F_{\mathbf{q}}^B$ in reciprocal space is geometrically a disc, and the projection that restores the bound constraint either leaves the structure factor unchanged, when it is inside the disc, or moves it to the nearest point on the circumference, when it lies outside. Using \mathcal{D} to denote the set of grid points \mathbf{q} for which

measured values $F_{\mathbf{q}}$ exist, and assuming bounds $F_{\mathbf{q}}^B$ can be found for all others, equation (3) should be replaced by

$$\left[\tilde{\Pi}_F(\tilde{\rho}) \right]_{\mathbf{q}} = \begin{cases} (F_{\mathbf{q}}/|\tilde{\rho}_{\mathbf{q}}|)\tilde{\rho}_{\mathbf{q}} & \text{for } \mathbf{q} \in \mathcal{D} \\ (F_{\mathbf{q}}^B/|\tilde{\rho}_{\mathbf{q}}|)\tilde{\rho}_{\mathbf{q}} & \text{for } \mathbf{q} \notin \mathcal{D} \text{ and } |\tilde{\rho}_{\mathbf{q}}| > F_{\mathbf{q}}^B \\ \tilde{\rho}_{\mathbf{q}} & \text{for } \mathbf{q} \notin \mathcal{D} \text{ and } |\tilde{\rho}_{\mathbf{q}}| < F_{\mathbf{q}}^B. \end{cases} \quad (5)$$

At large \mathbf{q} , one can obtain reliable bounds by extrapolating the measured structure factors on a Wilson plot (see §4.2). Near the origin, where usually only few structure factors are absent, an infinite bound is usually adequate and avoids a more difficult estimation problem. An example of a reciprocal-space grid, showing structure-factor values and bounds taken from data for a 148-atom peptide structure (Table 1, reference *e*), is shown in Fig. 2.

2.2. Atomicity constraints

Atomicity can be imposed as a support constraint, where the support S of the density is the subset of real-space grid points \mathbf{r} such that $\rho_{\mathbf{r}} \neq 0$. However, in contrast to phase retrieval with nonperiodic objects, where S is known or can be bounded, in crystallography one only knows that S has atomic characteristics. The simplest definition of an *atomic support* is the union of a known number of compact subsets of grid points, each representing one atom, and having arbitrary locations within the unit cell. Given a particular atomic support S , the projection Π_S of an arbitrary density ρ to a minimally modified density having support S is simply

$$\left[\Pi_S(\rho) \right]_{\mathbf{r}} = \begin{cases} \rho_{\mathbf{r}} & \text{for } \mathbf{r} \in S \\ 0 & \text{otherwise.} \end{cases} \quad (6)$$

If \mathcal{A} denotes the collection of all atomic supports (differing in atom locations), then atomicity projection is defined by

$$\Pi_A(\rho) = \Pi_{S'}(\rho), \quad (7)$$

where $S' \in \mathcal{A}$ is the atomic support that minimizes $\|\Pi_{S'}(\rho) - \rho\|$ over all $S \in \mathcal{A}$. While (6) can be computed quickly, an exhaustive search over all atomic supports in \mathcal{A} to find (7) may be prohibitive. We therefore adopt a heuristic (described below) that quickly finds an atomic support that is usually optimal.

In describing the precise projection operation, we distinguish two cases: atomicity projection for positive atoms (A_+), and atomicity projection for atoms of arbitrary sign (A). The




Figure 2 Structure factors in the plane $h = 0$ for a 148-atom peptide structure (Table 1, reference *e*): (a) measured values; (b) bounds.

former is used with X-ray diffraction data, the latter with neutron diffraction data when atomic species with both signs of scattering length are present. In both cases, we assume the number of atoms per unit cell N is known. In the case of X-ray diffraction, this will usually not include H atoms.

The first step in computing the projection $\Pi_{A_+}(\rho)$ is to sort the density values on the real-space grid. Then, beginning with the largest density, grid points with the property of being a local maximum are identified. A local maximum is defined by having a larger density value than any of its 26 neighboring grid points. Each time a local maximum is found, the 27 density values (maximum + neighbors) are copied, after *positivity projection*, onto a real-space output grid that was initially set to zero. Positivity projection, given by

$$\left[\Pi_+(\rho) \right]_{\mathbf{r}} = \begin{cases} \rho_{\mathbf{r}} & \text{if } \rho_{\mathbf{r}} > 0 \\ 0 & \text{otherwise,} \end{cases} \quad (8)$$

is the minimal modification that restores the positivity of atoms. The search through the sorted densities terminates when N local maxima have been identified and copied (positively) into the output grid. A graphical example of Π_{A_+} in two dimensions (and eight neighbors) is shown in Fig. 3(b).

Two modifications are required to compute the projection to atoms of arbitrary sign, $\Pi_A(\rho)$. To identify large peaks of arbitrary sign, the densities on the real-space grid are sorted by absolute value. Then, densities in the sorted list are identified as local maxima or minima, depending on their sign. Positivity projection is still applied to local maxima, whereas local minima are subjected to its counterpart: negativity projection. Fig. 3(c) shows the action of Π_A . The computational cost of both types of atomicity projection is dominated by the sort of M densities on the real-space grid. Using the quicksort algorithm, this cost grows as $M \log M$, proportional to the cost of Fourier intensity projection (§2.1).

The atomicity projections described require data with sufficient resolution. From the conventional definition

$$\text{resolution} = d_{\min} = 2\pi/Q_{\max}, \quad (9)$$

where $Q = (2\pi/\lambda)2 \sin \theta$ is the magnitude of the physical scattering wavevector, one can obtain in physical units the real-space grid spacing. The relationship between Q and the vector of Miller indices \mathbf{q} is given by

$$(Q/2\pi)^2 = \mathbf{q} \cdot \mathbf{M} \cdot \mathbf{q} \rightarrow (h/a)^2 + (k/b)^2 + (l/c)^2, \quad (10)$$




Figure 3 Atomicity projections of a two-dimensional density. (a) Density ρ within one unit cell. (b) $\Pi_{A_+}(\rho)$, the density closest to (a) having five positive atoms. (c) $\Pi_A(\rho)$, same as (b) but with the signs of the atoms unspecified.

where the matrix \mathbf{M} is a metric constructed from the unit-cell parameters, and the last expression gives the explicit form for an orthorhombic cell with dimensions a , b and c . For the ranges on the Miller indices to be consistent with a Q_{\max} , (10) shows in particular that

$$|h| < a(Q_{\max}/2\pi) = a/d_{\min}. \quad (11)$$

The number of reciprocal-space grid points for the index h is therefore $2(a/d_{\min})$, and since the real-space grid has the same number of points and has physical dimension a , the grid spacing is $d_{\min}/2$.

According to our projection heuristic, a pair of local extrema (of the same sign) can never be neighbors on the grid, but must be separated by at least two grid spacings in one of the three dimensions. Supposing for simplicity that the unit cell is nearly orthorhombic, the strongest bound on the relative displacement of a pair of atom centers arises when the latter is along the body diagonal of the grid. In that case, the displacement must exceed $(2, 2, 2)$ in grid units, otherwise the corresponding local extrema might be neighbors on the grid. The minimum separation of atomic centers must therefore satisfy $r_{\min} > 2 \times 3^{1/2}$ in grid units, or $r_{\min} > 3^{1/2}d_{\min}$. Since for organic structures $r_{\min} \approx 1.4 \text{ \AA}$ (neglecting H atoms), this statement implies $d_{\min} < 0.81 \text{ \AA}$. On the other hand, this bound is derived from the worst-case placement (relative to the grid) of two atoms in a structure of many atoms. Atom pairs displaced along a grid axis, for example, yield the more generous bound $d_{\min} < r_{\min}$. It is therefore not surprising that this form of atomicity projection has succeeded in solving organic structures at resolutions exceeding 0.81 \AA (Table 1).

3. The difference map

Given two sets of constraints on the density, implemented respectively by projections Π_1 and Π_2 , the difference map is an iterative procedure for obtaining a solution density ρ_{sol} that satisfies both constraints, specifically:

$$\Pi_1(\rho_{\text{sol}}) = \rho_{\text{sol}} = \Pi_2(\rho_{\text{sol}}). \quad (12)$$

Although our main interest is the projections $\Pi_1 = \Pi_A$ (or Π_{A+}) and $\Pi_2 = \Pi_F$, we begin with a review of the solution method for a general pair of projections (Elser, 2003a).

3.1. Fixed points and solutions

Starting with an arbitrary initial density $\rho(0)$, a sequence of iterates $\rho(n) = D^n(\rho(0))$ is generated by repeated application of the difference map:

$$D: \rho \mapsto \rho + \beta(\Pi_1 \circ f_2 - \Pi_2 \circ f_1)(\rho). \quad (13)$$

Each of the projections in (13) is composed with a map

$$f_i = (1 + \gamma_i)\Pi_i - \gamma_i \quad (i = 1, 2); \quad (14)$$

$\beta \neq 0$, γ_1 and γ_2 are real parameters. The difference map has two key properties, the first being that a solution, as defined by (12), exists if and only if the map has a fixed point, $\rho^* = D(\rho^*)$.

To see this, note that the difference in (13) vanishes at a fixed point, hence,

$$\Pi_1 \circ f_2(\rho^*) = \Pi_2 \circ f_1(\rho^*) = \rho'. \quad (15)$$

Applying either of the projections to (15) and using the property $\Pi_i \circ \Pi_i = \Pi_i$, we obtain

$$\Pi_1(\rho') = \rho' = \Pi_2(\rho'), \quad (16)$$

thus identifying ρ' with ρ_{sol} . Conversely, if ρ_{sol} exists, the set of fixed points is nonempty since it is easily verified that $D(\rho_{\text{sol}}) = \rho_{\text{sol}}$.

The fixed-point property of the difference map makes no reference to the detailed forms of the maps f_i . These maps are key to the second property: the attractive nature of the fixed points. In an iterative solution method, a fixed point is useless unless it is attractive; moreover, the greater the basin of attraction, the more effective is the method. Replacing the f_i by identity maps, for example, creates unstable (repulsive) directions in a fixed point's local behavior (Elser, 2003a), effectively reducing to zero the probability of arriving at the fixed point. The chosen form (14) of the maps f_i is the simplest, involving just the projections, that for suitable values of the parameters γ_i renders the fixed points of the difference map attractive.

Fixed points of the difference map should not be confused with the solution. The former are not unique, comprising in fact a submanifold in the space of densities; formally, the submanifold of fixed points is given by the intersection of inverse images:

$$(\Pi_1 \circ f_2)^{-1}(\rho_{\text{sol}}) \cap (\Pi_2 \circ f_1)^{-1}(\rho_{\text{sol}}). \quad (17)$$

Because of the strongly mixing property (see §3.4), it makes no sense to expend effort on the construction of the initial iterate, $\rho(0)$. All results presented here constructed $\rho(0)$ from the output of a simple random-number generator. During the course of iterating the difference map, the convergence to a fixed point is assessed by the norm of the difference

$$\varepsilon_n = \|(\Pi_1 \circ f_2 - \Pi_2 \circ f_1)(\rho(n))\|. \quad (18)$$

When ε_n becomes acceptably small, ρ_{sol} is obtained, as in (15), by applying $\Pi_1 \circ f_2$ (or $\Pi_2 \circ f_1$) to the estimate $\rho^* \approx \rho(n)$.

3.2. Parameter values

For a particular β , interpreted as a step size, the values of γ_1 and γ_2 are selected to optimize the convergence at fixed points. The earliest analysis of the difference map (Elser, 2003a) assumed local orthogonality of the two constraint subspaces and found

$$\gamma_1 = -1/\beta \quad (19)$$

$$\gamma_2 = 1/\beta. \quad (20)$$

Subsequent work (Elser, 2003b) considered an average-case analysis for particular kinds of constraints, including atomicity, and found optimal parameters ($1 \rightarrow S$, $2 \rightarrow F$)

$$\gamma_S = -1/\beta \quad (21)$$

$$\gamma_F = \frac{1 + t_F(1 - \beta)}{\beta}, \quad (22)$$

where $t_F \approx 0.5$ is the fraction of reciprocal-space grid points with known (non-negligible) structure factors, and the projection S (support) corresponds to atomicity (A or $A+$). However, because the typical step size is $\beta \approx 0.7$, the small numerical difference between (20) and (22) has not led to a significant change in performance of the algorithm. All experiments quoted in this work used the simpler expression (20).

Since there is as yet no theory for determining the optimal value of β for any particular application, β remains the single parameter of the algorithm that must be optimized empirically. Average solution times, measured in terms of difference-map iterations, are shown in Fig. 4 as a function of β for a 148-atom peptide structure (Table 1, reference *e*). Systematic trends in optimal β values with data resolution and $M/N = \text{gridsize}/\text{atoms}$ have not been performed, although it appears that optimal values fall in the range shown ($0.4 \leq \beta \leq 0.8$).

Some algorithms (*SnB*, *SHELXD*) treat the number of non-H atoms as a parameter, with N deliberately chosen significantly smaller than the best estimate of the actual number to improve performance. Various strategies for selecting atoms from the computed density are compared for the 148-atom peptide structure in the review by Sheldrick *et al.* (2001). Preliminary studies with the difference-map algorithm on the same structure showed only a very weak variation in the average number of iterations when N was varied by $\pm 8\%$. This result indicates that, while N is not a useful parameter, the algorithm can tolerate inevitable uncertainties in the actual numbers of atoms. From a logical viewpoint, choosing N larger than the actual number of atoms is always valid: the atomicity constraint has simply been weakened.

3.3. Convergence with imperfect data

When formulated in terms of constraints, the uniqueness of solutions to the phase problem requires an overconstrained situation. Considered geometrically, the two sets of constraints (say intensity and atomicity) are individually submanifolds in the space of densities with relatively low dimensionality. More




Figure 4

Average number of difference-map iterations required to solve a 148-atom structure (Table 1, reference *e*) for the parameter range $0.4 \leq \beta \leq 0.8$ (over 200 solutions per point).

specifically, in the overconstrained case, the sum of the dimensionalities is less than that of the ambient space (M), such that the intersection of generic submanifolds, of the same dimensions, would be empty (rather than a submanifold of positive dimension). The constraint submanifolds in a well posed phase problem, given perfect data, are nongeneric in the sense that a solution is known to exist or, equivalently, the submanifolds have a nonempty intersection in spite of their low dimensionality. The fine tuning implicit in the intensity data (say) required to achieve an intersection, or true solution, is upset by practically any departures from ideality. Chief among these in the crystallographic phase problem are statistical errors in the intensity measurements and the neglect of hydrogen atoms in the treatment of atomicity. Faced with these realities, one must abandon the hope of finding a solution in the strict sense.

Although the constraint submanifolds are not expected to perfectly intersect with realistic noisy data, we expect them to have a small separation (in the space of densities) in the vicinity of the true density. The convergence estimate ε [see (18)] can be interpreted as the currently achieved distance between constraint submanifolds, and solutions should be identified not by its vanishing but by its value dropping a significant amount. Plots of ε_n as a function of iteration n are contrasted in Fig. 5 for synthetically generated (top) and experimental (bottom) data. Experimental data for the 148-atom peptide structure (Table 1, reference *e*) was used for the imperfect data set, and synthetic data for a 148-equal-atom structure having the same M/N ratio was used to simulate perfect data. To ensure perfect compliance with the atomicity constraint, the Gaussian atoms used to create the perfect data




Figure 5

Evolution of the error estimate ε with difference-map iteration for two 148-atom data sets: (top) synthetic data for ideal $3 \times 3 \times 3$ atoms and no H atoms; (bottom) experimental data (Table 1, reference *e*).

were given supports on $3 \times 3 \times 3$ subsets of the real-space grid; atom centers, given by a random-number generator, avoided the minimum separation $r_{\min} = (12)^{1/2}$ (grid units). The sharp drop in ε displayed by the experimental data, and observed in all difference-map solutions reported here, demonstrates the viability of ε as a solution criterion even in the case of imperfect data.

With imperfect data, it appears that the minimum ε occurs shortly after the initial drop, and is not surpassed subsequently. As a best estimate of a fixed-point density, we therefore take the difference map iterate at the ε minimum; the corresponding solution estimate is found by applying $\Pi_F \circ f_{A+}$ or $(\Pi_F \circ f_A)$. Using this procedure on the known 148-atom peptide structure gave a mean figure of merit $\langle \cos \Delta\varphi \rangle = 0.71$ when averaged over all reflections in the data.

The overconstrained nature of the problem solved by the difference map can be appreciated by closer examination of the solution found with perfect data. Shortly after the sharp drop in ε when the solution is first found, ε decays monotonically to zero. This behavior implies that the structure factors of the density at large angles, which were provided only as bounds to the algorithm, are in fact being extrapolated to their true values by the iterative process (as was confirmed by direct examination of the solution's structure factors).

3.4. The solution process

The problem of finding a point in Euclidean space that satisfies a number of constraints, or showing that no such point exists, is known as a *feasibility problem* in the optimization literature. Theoretical studies have mostly focused on the case of convex constraint subspaces, where monotonicity of convergence can usually be proven for a variety of iterative methods. However, since both sets of constraints in the crystallographic phase problem (Fourier intensity and atomicity) are *nonconvex*, no rigorous results are available. The local analysis of the difference map, quoted above, only establishes the favorable fixed-point characteristics of the map, and provides no estimate on the number of iterations required to enter a fixed point's sphere of influence.

A dynamical systems perspective, combined with empirical data, provides a useful, though nonrigorous, picture for the difference map's mode of operation. Some salient features of the evolution of the density are illustrated in Fig. 6. Relatively rapid changes occur on the scale of very few iterations, and continue in an apparent steady state until quite abruptly the fixed point is encountered. During the long period of rapid changes there is no obvious progress toward the solution and the dynamics is well characterized as chaotic in the *strongly mixing* regime, where iterates settle into a stationary probability distribution very quickly. The basin of attraction of the map's fixed points has some overlap with this probability distribution, the magnitude of which determines the mean number of iterations required to arrive at the solution when averaged over starting points.

Taking this interpretation as a hypothesis, it can be tested by compiling a distribution of solution times for a given problem

instance, as measured by the number of iterations before the sharp drop in ε occurs (Fig. 5). If the strongly mixing property holds, then iterates are effectively subject to a probability of arriving at the attractive basin of a fixed point that is constant in time, and hence solution times will have an exponential distribution. An experiment comprising 3800 trials for the 148-atom peptide structure (all with $\beta = 0.7$) showed exactly this distribution. A solution was found in each trial, the longest requiring 7760 iterations, and the mean for all trials was 1100 iterations. The distribution of iterations, normalized relative to the mean, is plotted in Fig. 7 and compared with the exponential distribution. Overall the agreement is very good: the slight deviation at small iterations can be explained by a combination of the fixed (but small) number of iterations required to converge, first to the stationary probability distribution, then, after arriving at the attractive basin, to the fixed point.

The observed distribution of solution times greatly simplifies the solution protocol and eliminates yet another potential parameter: the bound on the number of iterations. Iteration bounds, typically some multiple of the number of atoms N , are imposed by *SnB* and *SHELXD* and results are quoted in terms of 'success rates'. Given an exponential distribution of solution times (Fig. 7), such a bound for the difference-map algorithm is arbitrary since it will have no effect on the number of solutions found per total iterations performed on all trials. Expressed in more direct terms: the performance of the algorithm is practically unaffected by random restarts, and hence there is no degradation of performance when iterations




Figure 6
Detailed behavior of the densities in the plane $x = 0$ for the solution shown at the top of Fig. 5: (a), (b) consecutive densities $\rho(100)$ and $\rho(101)$, (c) fixed-point density $\rho^* \approx \rho(3000)$, (d) solution density $\rho_{\text{sol}} = \Pi_F \circ f_{A+}(\rho^*)$.

are allowed to continue indefinitely. It is possible, however, that this conclusion will have to be modified if further experimentation with the difference map, say with small β , finds a nonexponential distribution of solution times.

4. Studies of test structures

4.1. The *atom_retriever* computer program

A preliminary implementation of the difference-map algorithm for crystallographic applications exists as the C-language program *atom_retriever*. The software in its current form is best characterized as a library of general-purpose subroutines that manipulate data in the uniform format of discretely sampled densities. At the lowest level are subroutines for performing a variety of projections, Fourier intensity and atomicity projection being of primary interest to crystallographers. The next level of subroutines combine the chosen projections into the difference map. Finally, at the highest level is a collection of drivers and translators that provide options for monitoring and terminating iterations, as well as converting input structure-factor data files into the rectangular arrays used by the algorithm. With this degree of transparency, it is hoped that users will experiment with innovations at the level of the projections, which are at the heart of the algorithm's success.

4.2. Space groups and structure-factor input

The primary input to *atom_retriever* is the rectangular array, containing structure-factor values and bounds (Fig. 2), and used by the Fourier intensity projection subroutine. Software for applying symmetry elements to structure-factor data in the construction of these arrays is still being developed, limiting applications to structures with triclinic (*P1* and $\bar{P}1$) space groups. In the future, the complete set of space groups will be implemented by preparing the initial density $\rho(0)$ with a projection that recognizes in reciprocal space the phase relationships between structure factors for the specified group; no changes are necessary in the *atom_retriever* program, since, apart from numerical rounding effects, the difference map preserves the symmetry of the density. Since symmetry projection software is also under development, the $\bar{P}1$ struc-

tures studied to date were treated as *P1*, with twice the true number of symmetry-inequivalent atoms.

The truncation of the individual atomic supports during atomicity projection to a $3 \times 3 \times 3$ array of grid points was shown by Elser (2003a) to be optimal when the corresponding Gaussian atom has a mean square displacement $\langle u_x^2 \rangle \approx 0.55$ in grid units, or

$$B/8\pi^2 = \langle u_x^2 \rangle \approx 0.14d_{\min}^2, \quad (23)$$

where B is an effective isotropic temperature factor. Defining a dimensionless temperature factor by

$$b = B/d_{\min}^2, \quad (24)$$

one finds that typical X-ray and neutron data sets (see Table 1) satisfy $b < b_{\text{opt}} \approx 11$. This means that $3 \times 3 \times 3$ is usually a generous support, perhaps even sufficient to accommodate hydrogen neighbors.

An effective temperature factor B , which combines the effects of atomic size, thermal vibration and certain kinds of static disorder, is estimated from the data by making a linear least-squares fit of pairs $\{Q^2, \log |F_Q|^2\}$ to the form

$$\log |F_Q|^2 = A - \frac{1}{2}B(Q/2\pi)^2, \quad (25)$$

for measurements in a restricted range $Q > Q_0$, where typically $2\pi/Q_0 = 1.2 \text{ \AA}$. At large spatial frequencies, the structure factors of centrosymmetric and non-centrosymmetric structures are well modeled as, respectively, real and complex Gaussian random variables, with distributions of intensities at wavevector Q given by (Giacovazzo, 1998)

$$P_Q(I) dI = \begin{cases} \exp(-\frac{1}{2}I/\langle I_Q \rangle) dI / (2\pi \langle I_Q \rangle I)^{1/2} & \text{(centric)} \\ \exp(-I/\langle I_Q \rangle) dI / \langle I_Q \rangle & \text{(acentric)}. \end{cases} \quad (26)$$

Since a least-squares fit applied to (25) gives a formula for the average of $\log I_Q$, the distributions (26) lead to the result

$$A - \frac{1}{2}B(Q/2\pi)^2 = \begin{cases} \log \langle I_Q \rangle - \gamma - \log 2 & \text{(centric)} \\ \log \langle I_Q \rangle - \gamma & \text{(acentric)}, \end{cases} \quad (27)$$

where γ is Euler's constant.

The probability distributions (26), with $\langle I_Q \rangle$ determined by the least-squares fit and (27), can now be used to determine bounds on the magnitudes of unmeasured structure factors with $Q > Q_0$, the bulk being in the corners of the reciprocal-space grid where $Q > 2\pi/d_{\min}$. If M' is the number of such (symmetry-inequivalent) structure factors, then $1/M'$ is an acceptable probability for an actual intensity to exceed the bound. Equating this with the probability $I > (F_Q^B)^2$ computed using (26) gives

$$F_Q^B = \begin{cases} (2\langle I_Q \rangle)^{1/2} \text{erfc}^{-1}(1/M') & \text{(centric)} \\ (\langle I_Q \rangle \log M')^{1/2} & \text{(acentric)}. \end{cases} \quad (28)$$

There are far fewer unmeasured structure factors with $Q < Q_0$, and the bound $F_Q^B = \infty$ was used for these in all the studies reported here.




Figure 7
Distribution of difference-map solution times (iterations) for the 148-atom peptide structure and $\beta = 0.7$. Solution times on the abscissa are normalized by the mean number of iterations, 1100. The curve shows the exponential distribution with mean unity, predicted by the strongly mixing hypothesis of difference map dynamics.

4.3. Discussion of tests

The results of a selection of tests of the *atom_retriever* program are summarized in Table 1. In each test, the entire available set of experimental structure factors was used, with missing data replaced by bounds in the manner described above; the grid dimensions correspond to $(2|h|_{\max} + 2) \times (2|k|_{\max} + 2) \times (2|l|_{\max} + 2)$. No effort was made to individually optimize performance with respect to β ; the value chosen, $\beta = 0.7$, is near the minimum in the average number of iterations for the 148-atom peptide structure (see Fig. 4). Mean figures of merit $\langle \cos \Delta\varphi \rangle$ for the *P1* structures were determined relative to calculated phases. For the *P1* structures, where calculated phases were not available, an internal figure of merit was determined by treating the structures as *P1* (with twice as many atoms) and taking the nearest set of centric phases as the true phases. When the number of iterations required by the algorithm has a broad distribution (see Fig. 7), the average number is quoted.

Atomicity projection for positive atoms (Π_{A+}) was used in all the X-ray data sets; the neutron data set for the mineral montebasite provided the sole application of the projection to atoms of arbitrary sign, Π_A . Nuclei with negative scattering length, such as Li, show up as light contrast in a field of dark atoms in plots of the scattering density ρ_{sol} . Fig. 8 shows the scattering density in one layer of the montebasite solution. The small number of iterations required to find the solution is typical of few-atom structures, where the error diagnostic (18) decreases almost monotonically with iteration (Fig. 8). When appropriately translated, the scattering density was nearly centrosymmetric, the resulting internal figure of merit being significantly larger than what would be obtained from random phases: $\langle \cos \Delta\varphi \rangle_{\text{rand}} = 2/\pi \approx 0.64$.

Rapid solutions with a near monotonic error decay were also observed for the other few-atom structures: nitramine, pyrrole, punctaporonin and triphenylphosphine. The nitramine and pyrrole structures were selected for their high density and cell aspect ratios, respectively. These characteristics had no noticeable effects on the algorithm's performance. Triphenylphosphine, interestingly, has more atoms (when treated as *P1*) than the 148-atom peptide structure, which requires many more iterations. Data resolution is not the cause of this anomaly, as was confirmed by truncating the




Figure 8
atom_retriever solution for the neutron data set of the mineral montebasite (Table 1); (left) near-monotonic decrease of the error with iteration. (right) scattering density ρ_{sol} in a plane of the structure showing an atom (Li) with negative scattering length (light contrast).

triphenylphosphine data to the same 0.96 Å resolution as the peptide: the solution was again found in only 25 iterations. A more likely cause is the presence of a moderately heavy P atom in each of the eight triphenylphosphine molecules of the structure, in contrast to the near equality of the non-H atoms of the peptide. Arbitrary atomic charges are trivially accommodated by the form of atomicity projection used by *atom_retriever*. A comparably non-specific atomicity was achieved relatively late in the development of the reciprocal-space-based framework (Hauptman, 1976; Rothbauer, 2000).

It is premature to assess the algorithm's prospects in solving large structures from atomic resolution data. The average solution time for the 148-atom peptide was 35 s on a 2 GHz Pentium 4 (single processor), not far behind *SnB2.2* (30 s) and *SHELXD* (11 s). More critical in evaluating performance is the growth in the average number of iterations with structure size. The largest structure attempted, a synthetic α -helical bundle, required nearly half a million iterations per solution (about 30 h). Fig. 9 shows the error estimate for a typical run, together with electron-density contours obtained directly from the discretely sampled density ρ_{sol} . The single noticeably larger peak in the density was identified with the chloride ion in the structure. This synthetic peptide differs from the smaller structures in that the number of ordered non-H atomic sites (N) is not known *a priori* because of the solvent contribution. The results quoted all used $N = 479$, the number of non-H atoms discovered in the original refinement (Privé *et al.*, 1999), but the structure was also solved with N as small as 440.




Figure 9
atom_retriever solution of the synthetic α -helical bundle (Table 1); (top) evolution of the error estimate, (bottom) electron-density contours in a plane containing the chloride ion (top, left of center).

5. Conclusions

The crystallographic phase problem consists of two logically distinct, though technically coupled, parts. Using the terms ‘needles’ for solutions and ‘hay’ for non-solutions, these parts are: (i) distinguishing needles from hay, and (ii) finding the proverbial needle in a haystack. As this work hopefully demonstrates, it is quite straightforward to recognize needles when presented with one: one only requires projections that act trivially (with negligible change) when operating on an object (*i.e.* scattering density) that satisfies all the known constraints (Fourier intensity, atomicity). If the constraints are too weak, as in a low-resolution data set, the phase problem may not be soluble in principle, because needles cannot be distinguished from the hay. On the other hand, even for low-resolution data, it is known (Podjarny *et al.*, 1987) that in certain circumstances needles can still be recognized. A properly phased protein crystal, for example, will often have a well defined solvent region and a characteristic histogram of density values (possibly at multiple length scales) within the body of the molecule, even when individual atoms are not resolved. If a projection (minimal density modification) can be constructed that applies in this more general setting, the first part of the phase problem can be said to be solved.

The difference map solves the second part of the phase problem by providing a uniform scheme for combining the applicable projections into an algorithm that finds the needle. Its eventual success is practically guaranteed if the corresponding constraints are strong enough to make needles distinct from hay. Perhaps most remarkable of all is the empirical fact that needles can be found in a reasonable time at all. Though the attractive basins of the difference map’s fixed points are, by proper choice of parameters, tuned to be as large as possible (Elser, 2003*a,b*), this local optimization cannot predict the average solution time. With more experience, we anticipate a body of empirical relationships between average solution times and characteristics of the data that can fill this gap. Some progress in this direction was recently achieved (Elser, 2003*c*) in a highly idealized version of the phase problem (Zwick *et al.*, 1996), where the discretely sampled (one-dimensional) density is known to be two-valued (a binary sequence).

atom_retriever was written while the author was a guest of the Center for Experimental and Constructive Mathematics at Simon Fraser University. Encouragement and technical suggestions from the director, Jonathan Borwein, and staff, in particular Rob Ballantyne, are gratefully acknowledged. *atom_retriever* would not have made the transition to a practical piece of software without George Sheldrick’s generosity in sending progressively more challenging data sets along with *SHELXD* benchmarks. I also thank him for his critical reading

of the manuscript. Charles Weeks kindly provided the *SnB* benchmark quoted in §4.3. The free use of experimental data sets is greatly appreciated, with particular thanks to Bryan Chakoumakos (Table 1, reference *a*) and George Sheldrick (Table 1, references *d* and *e*). Nick Elser was instrumental in adapting *atom_retriever* to run on the Cornell Theory Center computing cluster. This research was conducted using the resources of the Cornell Theory Center, which receives funding from Cornell University, New York State, federal agencies, foundations, and corporate partners. The author’s primary source of support is the National Science Foundation under grant ITR-0081775.

References

- Barrett, A. N. & Zwick, M. (1971). *Acta Cryst.* **A27**, 6–11.
- Elser, V. (2003*a*). *J. Opt. Soc. Am.* **20**, 40–55.
- Elser, V. (2003*b*). *J. Phys. A: Math. Nucl. Gen.* **36**, 2995–3007.
- Elser, V. (2003*c*). In preparation.
- Fienup, J. R. (1982). *Appl. Opt.* **21**, 2758–2769.
- Giacovazzo, C. (1998). *Direct Phasing in Crystallography*, pp. 12–17. Oxford University Press.
- Groat, L. A., Chakoumakos, B. C., Hoffman, C. M., Morell, H., Fyfe, C. A. & Schultz, A. J. (2002). *Am. Mineral.* **88**, 195–210.
- Hauptman, H. (1976). *Acta Cryst.* **A32**, 877–882.
- Hauptman, H. & Karle, J. (1953). *Am. Crystallogr. Assoc. Monogr.* No. 3. Ann Arbor, USA: Edwards.
- Karaghiosoff, K., Klapötke, T. M., Michailovski, A. & Holl, G. (2002). *Acta Cryst.* **C58**, o580–o581.
- Karle, I. L., Flippen-Anderson, J. L., Uma, K., Balaram, H. & Balaram, P. (1989). *Proc. Natl Acad. Sci. USA*, **86**, 765–769.
- Kooijman, H., Spek, A. L., van Bommel, K. J. C., Verboom, W. & Reinhoudt, D. N. (1998). *Acta Cryst.* **C54**, 1695–1698.
- Miller, R., Gallo, S. M., Khalak, H. G. & Weeks, C. M. (1994). *J. Appl. Cryst.* **27**, 613–621.
- Podjarny, A. D., Bhat, T. N. & Zwick, M. (1987). *Ann. Rev. Biophys. Biophys. Chem.* **16**, 351–373.
- Poyser, J. P., Edwards, R. L., Anderson, J. R., Hursthouse, M. B., Walker, N. P. C., Sheldrick, G. M. & Whalley, A. J. S. (1986). *J. Antibiot.* **39**, 167–169.
- Privé, G. G., Anderson, D. H., Wesson, L., Cascio, D. & Eisenberg, D. (1999). *Protein Sci.* **8**, 1400–1409.
- Ramos Silva, M., Matos Beja, A., Paixão, J. A., Sobral, A. J. F. N., Lopes, S. H. & Roche Gonsalves, A. M. d’A. (2002). *Acta Cryst.* **C58**, o572–o574.
- Rothbauer, R. (2000). *Z. Kristallogr.* **215**, 157–168.
- Sayre, D. (1952). *Acta Cryst.* **5**, 60–65.
- Sayre, D. (2002). *Struct. Chem.* **13**, 81–96.
- Sheldrick, G. M. (1997). *Recent Advances in Phasing: Proceedings of the CCP4 Study Weekend (DL-Conf-97-001)*, edited by K. S. Wilson, pp. 147–158. CCLRC, Warrington, UK.
- Sheldrick, G. M. (1998). *Direct Methods for Solving Macromolecular Structures*, edited by S. Fortier, pp. 401–411. Dordrecht: Kluwer.
- Sheldrick, G. M., Hauptman, H. A., Weeks, C. M., Miller, R. & Usón, I. (2001). *International Tables for Crystallography*, Vol. F. *Crystallography of Biological Macromolecules*, edited by M. G. Rossmann & E. Arnold, p. 342. Dordrecht: Kluwer.
- Weeks, C. M. & Miller, R. (1999). *J. Appl. Cryst.* **32**, 120–124.
- Zwick, M., Lovell, B. & Marsh, J. (1996). *Int. J. Gen. Sys.* **25**, 47–59.