

***FOX*, 'free objects for crystallography': a modular approach to *ab initio* structure determination from powder diffraction**

Vincent Favre-Nicolin and Radovan Černý

Copyright © International Union of Crystallography

Author(s) of this paper may load this reprint on their own web site provided that this cover page is retained. Republication of this article or its storage in electronic databases or the like is not permitted without prior permission in writing from the IUCr.

FOX, 'free objects for crystallography': a modular approach to *ab initio* structure determination from powder diffraction

Vincent Favre-Nicolin*† and Radovan Černý

Laboratoire de Cristallographie, Université de Genève, Switzerland. Correspondence e-mail: vincent.favre-nicolin@ujf-grenoble.fr

A new program has been developed for *ab initio* crystal structure determination from powder diffraction data (X-ray and neutron). It uses global-optimization algorithms to solve the structure by performing trials in direct space. It is a modular program, capable of using several criteria for evaluating each trial configuration (*e.g.* multi-pattern). It is also modular in the description of the crystal content, with the possibility of describing building blocks in the sample, such as polyhedra or molecules, and with automatic adaptive handling of special positions and sharing of identical atoms between neighbouring building blocks. It can therefore find the correct structure without any assumption about the connectivity of the building blocks and is suitable for any kind of material. Several optimization algorithms (simulated annealing, parallel tempering) are available, with the possibility of choosing the convergence criterion as a combination of available cost functions. This program is freely available for Linux and Windows platforms; it is also fully 'open source', which, combined with an object-oriented design and a complete developer documentation, ensures its future evolution.

© 2002 International Union of Crystallography
Printed in Great Britain – all rights reserved

1. Introduction

Powder diffraction using X-rays and neutrons plays a major role in the structure determination of new materials which are not available in the form of single crystals. *Ab initio* structure determination from powder diffraction data is more difficult than single-crystal diffraction analysis, because of the projection of the three-dimensional diffraction pattern on a single dimension. During the past 12 years, the number of structure determinations using powder diffraction increased to around 600 (Le Bail, <http://www.cristal.org/iniref.html>). The methods used can be divided into two groups, as follows.

(a) Reciprocal-space methods: so-called direct methods. Reciprocal-space methods require the extraction of structure-factor amplitudes from the powder pattern, which is complicated by a strong overlap of diffraction peaks; this is especially important because low-intensity reflections play an important role in direct methods (Casarano *et al.*, 1984, 1991). Significant improvement of the structure-factor amplitude extraction from a powder pattern has been achieved recently using the high angular resolution attainable at synchrotron-based powder diffractometers, as well as by employing texture (Lasocha & Schenk, 1997; Černý, 1998; Wessels *et al.*, 1999), anisotropic thermal expansion (Shankland, David & Sivia, 1997), fast iterative Patterson squaring (Estermann & Gramlich, 1993) and Bayesian approaches (Sivia & David, 1994).

However, in some cases (texture and thermal expansion) additional experimental equipment and more time for collecting additional powder patterns are needed. Furthermore, for high-symmetry closely packed compounds (inter-metallics), the analysis of *E* maps can be very difficult and ambiguous.

(b) Direct-space methods: locating building units (molecules, polyhedra, atoms) in the cell and comparison of calculated and observed diffraction patterns. Direct-space methods do not require the extraction of structure factors or any additional experimental equipment. Intense work has been done during the past 10 years in the development of this method for molecular crystals. In general terms, it is a global-optimization problem of a great complexity, in which the agreement between the observed and calculated diffraction patterns is maximized. Several optimization algorithms have been used: Monte Carlo (MC) search (Harris *et al.*, 1994), MC search with simulated annealing (SA) (Andreev *et al.*, 1997; Newsam *et al.*, 1992) and genetic algorithm (GA) (Shankland *et al.*, 1997; Kariuki *et al.*, 1997). Several computer programs using direct-space methods have appeared recently (for a full review see <http://www.cristal.org/iniref.html>). Mostly, they use MC search with SA [*e.g.* *POWDER SOLVE* (Engel *et al.*, 1999), *ESPOIR* (Le Bail, 2001), *TOPAS* (Bruker AXS, 2000), *ENDEAVOUR* (Putz *et al.*, 1999), *DASH (ex-DRUID)* (David *et al.*, 1998), *GA (GAP)* (Shankland *et al.*, 1997), *GAPSS* (Kariuki *et al.*, 1997)]. Some of them combine the diffraction data with a crystal potential-energy minimization.

† Current address: DRFMC/SP2M/IRS, 17 rue des Martyrs, 38054 Grenoble Cedex 9, France

They are mostly oriented towards molecular compounds, and have not very often been applied to non-molecular (inorganic) crystals. Little effort has been spent in the parameterization of the crystal structure (description of the structure with building blocks) for non-molecular (inorganic) crystals. Greater progress has been made only for framework (zeolite) crystals, using either SA [*ZEFSAII* (Falcioni & Deem, 1999)] or a combined search in the direct space (expected structure topologies) and reciprocal space (diffraction data and direct methods with Fourier recycling) [*FOCUS* (Grosse-Kunstleve *et al.*, 1997)]. A similar approach is currently being developed within the program *EXPO* (Altomare *et al.*, 2002). Generally, using available programs it is difficult to extend the validation criteria for a specific problem.

To address these issues we have used a modular approach for *FOX*, 'free objects for crystallography', as follows.

(i) The crystal structure can be described using any combination of isolated atoms, molecules or polyhedra. The correct structure can be found without any assumption on the actual connectivity between these building blocks, nor any *a priori* knowledge about atoms on special positions.

(ii) Several optimization algorithms can be used, with an easy path from one to another, each being easily upgradable. Each optimized object (crystal structure, powder pattern) can define its cost function (CF) and any combination of CFs can be used as a criterion.

(iii) It is possible to use jointly all available data sets: powder pattern and single-crystal, neutron and X-ray.

(iv) Extending the program (parametrization of the crystal structure, algorithms, criterion for validating the structure) is easy.

2. The global-optimization process: algorithms and cost functions

The algorithms currently implemented in *FOX* use a so-called reverse Monte Carlo approach (McGreevy & Pusztai, 1988) (see Fig. 1): starting from a random configuration, the free parameters in the structure are varied randomly. Configurations are compared using a CF, which is characteristic of how good the structure is, and which is defined using either *a priori* knowledge about the compound (energy model, bond distances, *etc.*), and/or from experimental data. The CF is a strictly positive real number, smaller values representing better configurations. Thus, in the Markov process, which ensures the sampling of the parameter space, the new generated structure is kept according to the widely used Metropolis algorithm (Metropolis *et al.*, 1953): (a) systematically if the CF is better (smaller) than the preceding configuration, and (b) with probability $\exp(-\Delta CF/T)$ if the new configuration has a higher cost. The relative probability of two independent trial configurations follows Boltzmann law: $P_1/P_2 = \exp[(CF_2 - CF_1)/T]$, with T being the temperature of the distribution. It is important to understand that the overall list of generated trial configurations indeed follows a Boltzmann-type law, but one that is weighted by the density of states (the number of available configurations at each cost), so

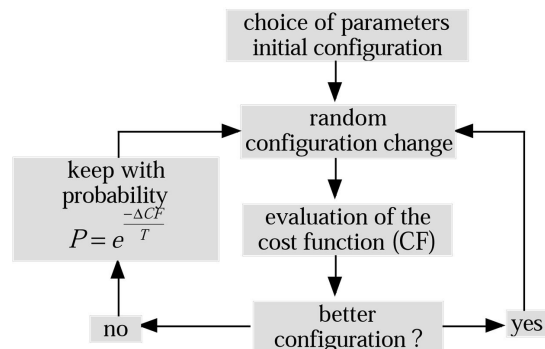


Figure 1

Workflow of a Monte Carlo type algorithm. The parameter space is explored by making random configuration changes from a known parametrization. The cost function (CF), the temperature T (and its variation during the optimization) and the random moves (amplitude of displacements, biased moves) must be carefully chosen to ensure an efficient convergence of the algorithm.

that the actual number of configurations near the global minimum is very small.

The efficiency of such an algorithm depends crucially on the optimization, the nature of the random moves, and the choice of an adequate CF, which we will detail now.

2.1. Simulated annealing versus parallel tempering

In order to make the MC process converge, it is necessary to lower the refinement temperature, so as to favour better configurations. This can be done using SA, in which the temperature is slowly decreased after each trial configuration. But the hypersurface representing the CF as a function of all parameters is extremely complex and presents numerous false minima. Thus the drawback of SA is that if the temperature decreases prematurely, the algorithm can be trapped in a local minimum, with a negligible probability to overcome the energy barrier to reach the global minimum.

In order to cope with this effect, another closely related algorithm available is parallel tempering (PT) (Falcioni & Deem, 1999): instead of using a single chain of configurations with a decreasing temperature, a small number of parallel optimizations are made, each with a different temperature. Periodically (in *FOX*, after performing ten trials in each parallel optimization), the algorithm tests an exchange of configurations between parallel optimization processes, with the same acceptance rules as when evaluating successive trials. This has two advantages: (i) the user does not need to evaluate how many trial configurations will be necessary to reach the global minimum, since the algorithm is invariant with the trial number; and (ii) at all times during the optimization there is the possibility to reach any configuration, however high the associated cost is, by using the higher-temperature process.

In our experience, SA can be significantly faster for very simple structures [less than ten independent parameters or degrees of freedom (DOF)], but is much less efficient for complex structures, where the difficulty in choosing the correct temperature regime and the number of trials often results in a

low success rate. This can be overcome by systematically performing several successive SA runs, but parallel tempering still has a higher success rate and is simpler for the user. We have also found that, in PT, using a larger number of parallel optimizations (30 rather than 5 or 10) leads to a greater success rate of the algorithm, without penalty in the speed of the convergence; this indicates a better sampling of the parameter space.

2.2. Adaptive displacement amplitudes versus temperature values

Besides the temperature regime, it is important to choose an adequate amplitude for the random displacements of all parameters: at high temperatures, it is better to allow relatively large modifications between each trial, so that the less-favourable parts of the hypersurface are not over-sampled. For the best configurations which are located in narrow minima, it is necessary to use smaller steps to sample finely towards the global minimum.

In *FOX*, we use random displacements of all parameters following the base law: $p_{\text{new}} = p_{\text{old}} + \Delta p$, with $|\Delta p| < (0.01 \times M)$ for translational parameters, and $|\Delta p| < (3.6^\circ \times M)$ for angular parameters, where M is a multiplier which is adjusted depending on the location on the hypersurface. The value for M can either be a direct function of the temperature, or dynamically adjusted; this is done by imposing an average acceptance rate for trials. We have found that choosing an acceptance rate between 10 and 30% for trial configurations worked well: it ensures both that not too many trials are accepted (which would mean that the algorithm runs too slowly and wastes computing time), nor too few (which would prevent the adequate sampling of the hypersurface). Generally, we let the maximum amplitude multiplier vary from 8 (higher costs) to 1/8 (lower costs).

One remaining difficult choice for the algorithm is the range of temperatures to use, for it is closely related to the properties of the hypersurface (or equivalently to the CF used): maximum and minimum values, large or narrow valleys towards the

minimum, etc. Once a CF has been chosen, a few trial optimizations would quickly indicate the correct values, but in the prospect of a modular algorithm with a set of user-chosen cost functions, a more general approach is needed.

To obtain this we reversed the amplitude versus temperature problem: contrary to the temperature values, the maximum and minimum displacement amplitude values with which we want to sample the parameter space are universal, independent of the CF used. Thus it is possible to choose a fixed law following (say) an exponential decrease of displacement amplitude values, from 1 Å down to 0.01 Å for translations, and from 36° down to 0.36° for angles. The temperature values are then adjusted so that the same acceptance rate between 10 and 30% is reached. The temperature values are thus naturally tuned depending on the narrowness of the hypersurface. In our experience, this 'smart' temperature tuning has given as good results (*i.e.* same success rate in the same average number of trials) as when using an exponential decrease of the temperature with hand-optimized maximum and minimum values, independently from the choice of CF.

2.3. Cost functions

The definition of the cost associated with each trial structure is essential in a global-optimization algorithm. In *FOX*, we have not imposed any CF on the user, but we have preferred to allow each crystallographic entity (crystal structure, diffraction data, etc.) to define its own CF; each of these CFs should approximately scale between 0 (excellent matching) and 1 (completely wrong), so that all CFs more or less scale like a crystallographic R factor. Moreover, it is easy to add new CFs to all objects. The user can choose any combination of these CFs, and the weighted sum of the CFs is used as an overall CF in the algorithm.

2.3.1. R_{wp} and integrated R_{wp} . In powder diffraction, the most often used CF is the weighted profile R factor calculated over the whole powder pattern:

$$R_{\text{wp}} = \left[\frac{\sum_i w_i (y_i^{\text{obs}} - y_i^{\text{calc}})^2}{\sum_i w_i (y_i^{\text{obs}})^2} \right]^{1/2} \quad (1)$$

Even if R_{wp} is the function which takes full advantage of the information available in the powder pattern, it is often the case that due to lattice defects or diffractometer aberrations the pattern has badly defined reflection profiles. For this reason it is interesting to use integrated R and R_w factors: these are calculated by summing the total intensity scattered around each peak position, in the range $2\theta_0 - f \cdot \text{FWHM}$ to $2\theta_0 + f \cdot \text{FWHM}$, with $f = 1$ for a Gaussian shape and $f = 2$ for a Lorentzian (FWHM is the full width at half-maximum): this range generally allows one to include more than 90% of the diffracted intensity for all reflections (see Fig. 2). Should two reflections overlap, the integration segments are divided between the two reflections: the summation is always made on the full calculated and observed profiles, so that no assumption is made about the actual values of the Bragg intensities, contrary to methods that require the extraction of intensities.




Figure 2 Integrated R factors (iR_{wp}) can be used for global optimization that is insensitive to the reflection profiles, without having to extract individual reflection intensities. Full observed and calculated profiles are calculated; the values compared in the weighted R factor are the integrated intensities calculated around each expected reflection position. Integration areas are separated so as not to take any point into account more than once.

This allows one to overcome most of the difficulties in describing accurately the reflection profile (especially for diffraction data with multiple phases), and the small amount of information lost still allows one to find the correct structure solution, even if such a method would not be suitable for a refinement. Note that it is also possible to use reflection intensities (e.g. pseudo-single-crystal intensities extracted from the powder pattern), using the weighted R factor as another CF.

2.3.2. Structural cost functions: 'anti-bump'. Any CF can be valuable to find the correct structure, either to find the global minimum, or to disfavour unsound configurations and thus reduce the overall parameter space to be sampled. An energetic evaluation of each configuration can thus be used, with potentials evaluating quantitatively the bonding of each atom with its neighbours. The difficulty with energetic descriptions is that they are highly specific to the kind of material studied, depending on the nature of the chemical bond (ionic, metallic or covalent). A mistake in the energetic description can easily lead to an incorrect global minimum.

For this reason, no energetic description is available yet in *FOX*, and we have preferred to implement a simple anti-bump (AB) CF that adds a penalty when two atoms are closer than a minimum distance. This minimum distance can be input by the user for each pair of atom types. For identical elements, this function also allows the merging of elements (when the distance tends toward zero), so that for identical atom types which completely overlap, the penalty decreases to 0. This CF can be used for any kind of material, and only helps the algorithm to avoid improbable configurations, without affecting the position of the global minimum (at the global minimum all atoms are in a chemically sensible coordination, with the anti-bump CF kept at a null value at and near the global minimum).

3. Modelization of the crystal structure

In order to allow a wide variety of compounds, the crystal structure can be described as any combination of 'scatterer' objects, which can be independent atoms, molecules or polyhedra. The description using the largest building blocks is vital in an *ab initio* structure determination process, since the number of trials required will (roughly) vary exponentially with the DOF (degrees of freedom).

Describing the structure using building blocks consists of using all the *a priori* information about the connectivity of the atoms to reduce the DOF: e.g. a PO_4 group requires six parameters (three translations + three rotations) as a tetrahedron, but 12 parameters if described as individual atoms. The number of DOF can generally be reduced by a factor of two for inorganic structures, and at least by a factor of three for most molecular structures, due to the good *a priori* knowledge of all bond distances and bond angles. It can further be reduced for molecules in the presence of rigid parts (cycles). By taking into account this *a priori* information, not only can the global minimum be found faster, but it also

ensures that the solution will be sound from the connectivity point of view.

3.1. Molecular components

3.1.1. Description from bond lengths, angles and dihedral angles. For molecules, the plane formula (and possibly the absolute configuration) is known before solving the structure. From common knowledge of organic compounds, the bond distances and angles can be determined with a very good precision (less than 0.03 \AA for bond distances and 5° for angles). The only completely free parameters are the dihedral (torsion) angles of non-constrained bonds (single bonds not part of a small cycle), so that the number of parameters to be determined is generally nearly equal (or even smaller if there are several rigid cycles) to the number of atoms (instead of three times the number of atoms). To describe this geometry correctly, we use internal coordinates recorded in a Z matrix (see Fig. 3). In this description the atoms are listed in a 'matrix', with the position of each atom determined from the bond distance, a bond angle and a dihedral angle with respect to three preceding atoms in the matrix.

3.1.2. Optimizing molecule parametrization for global-optimization algorithms. The reduction of the average number of parameters per atom has a price which lies in the interdependence of the parameters: in an optimization with 'free' atoms, the algorithm can (ideally) put the atoms in the correct positions one after the other. Of course, in the case of global optimization from diffraction data, the situation is more complex since all atoms contribute to all points of the data set, but once a large number of trial configurations have been generated, most atoms fall generally near a position where the electronic density is high, and thus near the final correct position. From that point in the optimization, with free atoms or small groups of atoms, the convergence towards the global minimum is generally extremely fast, since the atoms can move more or less independently.




Figure 3

Building blocks (molecules, polyhedra) can be described by storing in (a) a matrix which records the (b) bond distances, angles and (c) dihedral angles describing each atom position relative to the position of three preceding atoms. This leads to a 'natural' parametrization which fully takes advantage of *a priori* knowledge about the sample contents (polyhedra, molecule plane formula).

For molecules the situation is very different since a modification of one torsion angle in the middle of the molecule affects the position of all atoms in the second part of the molecule: practically, this means that it is impossible to find the position of the last atom (in the Z -matrix description) without finding first the other atoms. Thus for a given number of parameters, the number of trial configurations to construct will be much larger than for a modelization with free atoms or building blocks. The 'first stage' of the global optimization (putting all atoms roughly near the high-density zones, at 1–2 Å from the final position) runs at the same speed for a modelization with free atoms and for that with molecules, but the 'second stage' (having atoms within 0.2–0.4 Å of the final position) is longer for flexible molecules. In terms of a hypersurface description of the parameter space, this means that the global minimum is very deep and abrupt, with a limited number of pathways to reach it. In order to minimize this effect, the order of the atoms in the Z matrix (the order in which the torsion angles are taken into account) must be carefully chosen, so that the most flexible parts of the molecules (which will be determined only at the end of the optimization) are moved towards the end of the list.

Another way to decrease this effect is not to fix all known parameters (bond distances and angles), but to allow them to vary within small limits to 'create' additional pathways towards the global minimum. An even better way consists in making more intelligent random moves of parameters: normally the modification of one torsion angle displaces all atoms after the torsion angle. It is also possible to modify only (i) the atoms before the torsion angle, but this requires modifying the global orientation and position of the molecule, or (ii) one atom position in the chain, by compensating the move on the torsion angles of neighbouring bonds. The diffi-

culty in both cases is that there is no simple method to minimize the changes so we must use a costly (in terms of processor time) algorithm for the local minimization, and thus these 'smart moves' can only be tried infrequently (2% of moves in the current version of *FOX*). The use of these smart moves in the tests with the cimetidine sample (see §5.2) has decreased the failure rate from 20 to 10%.

We believe that many more optimizations could be obtained for this 'second stage' of the global optimization of molecule structures by having more independent atoms, but this is a delicate task if one does not want to increase the DOF. This could certainly be achieved by changing the modelization of molecules using methods developed for macromolecular compounds (even if the structure determination methods, based on electronic density rather than the optimization with regard to diffraction data, are very different), such as (i) decomposing the molecule in several parts (*e.g.* amino acids) which are connected *a posteriori* in the optimization, or (ii) with a combination of free atoms and bond/bond angle/dihedral angle restraints. This combination of free atoms combined with restraints would also yield a modelization which is invariant with the order of the atoms in the molecule, thus with a better scalability than the modelization based on a simple Z matrix. This modification is being considered for future versions of *FOX*.

3.2. Non-molecular components

3.2.1. Building blocks: polyhedra. Inorganic chemists and crystallographers can very often predict the type of atomic coordination from the formula of their compounds. A wide range of polyhedra are available in *FOX* (tetrahedron, octahedron, cube, prism, square plane, icosahedron), and these are described using Z matrices as for organic molecules, and therefore use a natural description using bond distances, angles and dihedral angles. To describe a crystal structure adequately, it is necessary to take into account the possibility of corner-sharing between polyhedra, which we will describe now along with the handling of special positions.

3.2.2. Handling of special positions and corner-sharing. One important difficulty for the global optimization of inorganic materials is the very often high symmetry, which implies that a number of atoms will be in special positions. Except in simple cases with small unit cells, it is not generally possible to determine *a priori* which atom will lie on a special position. The optimization algorithm must therefore be able to determine when atoms fall on these positions, while also allowing atoms to move away from these positions, without any intervention from the user.

A simple approach is, for each generated configuration, to determine the distance of each atom from a special position (either determining directly the distance to the symmetry plane, axis or centre, or by measuring the distance with its symmetric positions with respect to the symmetry), and if that distance is smaller than (say) 0.5 Å, then the atom is projected on the special position. This algorithm has several shortcomings: (i) it can be complex (ambiguous) to implement




Figure 4

The dynamical occupancy correction (DOC): when several identical atoms approach each other, either because (a) the atom is nearing a special position or (b) there is a formation of a bridge between two building blocks, it is possible (c) to correct the occupancy of all the overlapping atoms, so that the electronic (or nuclear) density used in the scattered intensity correction corresponds to a single atom. This allows a continuous correction for any kind and any number of overlapping atoms.

when one atom is near the intersection of several symmetry elements, (ii) it is not ‘continuous’, since it brutally moves an atom, not allowing the algorithm to see a ‘better’ configuration until it is very close to it, and (iii) the displacement of the atom may break its connectivity (or even be impossible to do) when the material is described in terms of building blocks (polyhedron or molecule).

To cope with this we have implemented a general algorithm using a dynamical occupancy correction (DOC, see Fig. 4): we generate all atom positions in and near the asymmetric unit, then for all unique atoms a neighbour table with all atoms less distant than $d_{\min} = 1 \text{ \AA}$ is generated. The occupancy is then corrected as follows:

$$\text{Occupancy}^{-1} = 1 + \sum_{\text{neighbour}} |d_{\min} - d_i|. \quad (2)$$

From this formula (the sum is over all neighbours of the same atomic type), it is straightforward to see that if n atoms overlap, the dynamical occupancy correction converges to $1/n$. After this correction has been applied, there will be n atoms with occupancy equal or converging to $1/n$. Thus the algorithm for computing structure factors will use the correct electronic (or nuclear) density.

This algorithm is very straightforward to implement and addresses all the shortcomings of the ‘projection’ approach. Moreover, it can correct any type of overlapping with any number of involved atoms, for example when two equivalent atoms are shared between adjacent polyhedra. Furthermore, the exact content of the elementary cell does not need to be known exactly, since excess atoms are simply merged¹ by the program. The only drawback is that it requires the calculation of the distance table, which requires a significant amount of processor time, even compared with the time required for computing structure factors.

4. Diffraction data

Powder pattern data (neutron, synchrotron or laboratory X-ray) can be modeled by a sum of phases: one background phase and any number of crystalline phases associated with declared crystal structures, of known or unknown structure.² For each crystalline phase, the profile can be described using the standard Gaussian, Lorentzian or pseudo-Voigt profiles (Young, 1999). Preferred orientation can also be optimized using the March–Dollase model (Dollase, 1998).

It is also possible to perform the global optimization of a crystal structure from single-crystal data (*i.e.* extracted intensities), provided that all corrections (absorption, Lorentz–polarization) have been performed beforehand (*i.e.* F_{obs}^2).

¹ Note that merged only means that the overlapping atoms look like a single atom in terms of scattering. The algorithm does not remove any atom, so that in a further evolution of the algorithm, the overlapped atoms can move away from each other; the correction is therefore reversible.

² Several unknown phases can be determined on a given powder pattern, provided that the indexing was possible before trying global-optimization methods.

5. Applications

The algorithms and modelizations proposed in *FOX* have been validated both by testing on known structures, and by solving new structures of various complexity. Table 1 gives a summary of the results obtained. The space group and unit-cell information was either taken from the original publication, when available, or determined using the *DICVOL91* program (Boultif & Louer, 1991). A pseudo-Voigt profile shape with a profile width following Caglioti’s function (Caglioti *et al.*, 1958) was used in all cases. The cell and profile parameters were refined on the full powder pattern using the program *FullProf* (Rodriguez-Carvajal, 1993). For all structures, the experimental data were cropped to $1/2d < 0.25 \text{ \AA}^{-1}$ (approximately), which is sufficient for structure solution purposes, since the ratio of DOF to the number of (independent) reflections was greater than 5.

The algorithm used for the results shown here was systematically PT. The maximum amplitude of random moves in the parallel worlds of the PT algorithm followed an exponential decrease from 1 \AA and 40° down to 0.01 \AA and 0.5° , and the temperature was dynamically adjusted as explained in §2.2. The success was determined by comparing all atomic positions to the known values for previously known structures, and for new ones the Rietveld refinement was performed. For repetitive tests, a critical value of the cost under which the structure was near enough to the final solution was first determined (practically, 1% above the best CF value achievable during the global optimization), and the results for individual runs were checked afterwards; tests were considered as failed if the number of trials exceeded three times the typical number of trials required for finding the structure solution. For all structures, it was necessary to estimate the number of atoms, polyhedra or molecules present in the asymmetric unit cell and add all these building blocks in the crystal structure. When in doubt (*e.g.* total deuterium content for deuterides), it was possible to exceed the number of independent components, the DOC being able to merge the excess atoms automatically. The bond lengths, angles and dihedral angles (for molecules and polyhedrons) were taken from known values for similar compounds, and for all parameters which were not completely free (not considered as a DOF), the values were never fixed but allowed to vary with tight limits, generally $\pm 0.05 \text{ \AA}$ for bond lengths, and $\pm 5^\circ$ for angles. The starting configuration used random values for all DOF, within limits for the limited parameters. Isotropic displacement parameters were fixed to an estimated value depending on the atomic weight of the element.

5.1. Inorganic structures

The first tests were conducted on PbSO_4 ; the powder data used were from the examples list of the program *ESPOIR* (Le Bail, 2001). For this compound, both neutron and X-ray diffraction patterns were available, so we used it as a test for multi-pattern global optimization. We tested the two possible parameterizations for the structure: (i) using free atoms (Pb + S + 4O), and (ii) using one free atom (Pb) and one tetrahedron

Table 1

 Structures solved using *FOX*.

Notes: (a) New structure. (b) Final number of independent atoms found by *FOX* (M: metal; D: deuterium). (c) The number of atoms in the modelization can be higher than the final number of independent atoms, due to the *a priori* unknown exact composition or the presence of special positions; excess atoms are merged by the dynamical occupancy correction (see §3.2.2). (d) Xl: laboratory X-rays; Xs: synchrotron; N: neutron; Xl+N: joint optimization; Xs,N: successive use of neutron and synchrotron data. (e) Time is given for a 14 GHz (Athlon) computer running Linux, for the median. (f) iR_{wp} : integrated R_{wp} ; R : calculated on extracted reflection intensities; AB: anti-bump. (g) The number of trials required to find the structure with this modelization was widely distributed; see text.

	Space group	Atoms ^b	DOF ^c	Building blocks ^c	Data ^d	Success	Trials × 10 ³ (time) ^e	Cost function ^f
PbSO ₄	<i>Pnma</i>	5	18	6 free atoms	Powder (Xl+N)	100%	<50 (<30 s)	$iR_{wp}(X) + iR_{wp}(N)$
PbSO ₄	<i>Pnma</i>	5	9	Pb + SO ₄	Powder (Xl+N)	100%	<50 (<30 s)	$iR_{wp}(X) + iR_{wp}(N)$
CsOH.H ₂ O ^a	<i>I4₁/amd</i>	2	9	3 free atoms	Powder (Xl)	100%	<50 (<30 s)	iR_{wp}
LiBH ₄ -RT phase ^a	<i>Pnma</i>	5	9	Li + BH ₄	Powder (Xs)	100%	<50 (<30 s)	$iR_{wp} + AB$
LiBH ₄ -HT phase ^a	<i>P6₃mc</i>	4	9	Li + BH ₄	Powder (Xs)	100%	<50 (<30 s)	$iR_{wp} + AB$
NdNi ₄ MgD _{3.6} ^a	<i>Pmn2₁</i>	5 M, 3 D	15, 15	5 M, 5 D free atoms	Powder (Xs,N)	100%	100 (60 s)	$iR_{wp}(X, \text{then } N) + AB$
Zr ₃ NiO _{0.6} D _{6.32} ^a	<i>Cmcm</i>	4 D	15	5 D free atoms	Powder (N)	100%	100 (60 s)	iR_{wp}
LaNi ₂ Mn ₃ D _{5.5} ^a	<i>P6/mmm</i>	5 D	15	5 D free atoms	Powder (N)	100%	100 (60 s)	$iR_{wp} + AB$
ErFe ₂ D _{4.72}	<i>Pmn2₁</i>	10 D	60	20 D free atoms	Powder (N)	100%	500 (300 s)	$iR_{wp} + AB$
Cimetidine	<i>P2₁/c</i>	17	14	C ₁₀ N ₆ S	Powder (Xs)	90%	4000 (2800 s)	iR_{wp}
Cimetidine	<i>P2₁/c</i>	17	14	C ₁₀ N ₆ S	Extracted (Xs)	90%	4500 (1400 s)	R
Potassium tartrate	<i>P2₁2₁2₁</i>	11	14	K + C ₄ O ₆	Powder (Xl)	95%	1200 (720 s)	iR_{wp}
Al ₂ (CH ₃ PO ₃) ₃ ^a	<i>P1</i>	17	24	2Al + 3H ₃ CPO ₃	Powder (Xl)	100%	750 (750 s)	$iR_{wp} + AB$
Al ₂ (CH ₃ PO ₃) ₃ ^a	<i>P1</i>	17	27	AlO ₄ + AlO ₅ + 3H ₃ CP	Powder (Xl)	^g	6500 ^g (6500 s)	$iR_{wp} + AB$

(SO₄). For both modelizations the correct structure was found quickly (less than 50000 trials). The Pb atom position was found first, mostly from the information from the X-ray pattern, and then the other atoms were found. In the correct structure, there are only three independent oxygen atoms (which we did not know *ab initio*); the dynamical occupancy correction correctly reported this. Two oxygen atoms (out of the four initially present) are merged in both modelizations.

The first new structure solved by *FOX* is a new polymorph of CsOH.H₂O (Černý *et al.*, 2002). The structure contains only two independent non-hydrogen atoms (one Cs and one O), both on special positions. They were localized very fast starting from three free atoms (one Cs and two O).

FOX is particularly efficient in the localization of deuterium atoms (using neutron data) in a metallic matrix of metal hydrides (determined from X-ray data). No preliminary search for available interstices in the matrix is necessary. New structures of metal hydrides determined by *FOX*: LiBH₄, room-temperature (RT) and high-temperature (HT) phases (metal and hydrogen positions) (Soulié *et al.*, 2002), NdNi₄MgD_{3.6} (metal and deuterium positions) (Guénée *et al.*, 2002), Zr₃NiO_{0.6}D_{6.32} (deuterium positions) (Černý *et al.*, 2002), LaNi₂Mn₃D_{5.5} (deuterium positions) (Guénée & Yvon, 2002), ErFe₂D_{4.72} (deuterium positions) (Paul-Boncour *et al.*, 2001).

The simplest of these structures is LiBH₄ (RT and HT), where the hydrogen atoms form a tetrahedron around the boron atom. The optimization was made with one Li and one BH₄ tetrahedron. The structures of both phases were determined only from X-ray data; the hydrogen atoms were found thanks to the combination of (i) their known tetrahedral coordination around the boron, (ii) the fact that there are only light atoms in the structure, and (iii) the use of anti-bump CF.³

³ For all hydride structures, the AB cost function was parametrized with a 'repulsion' between deuterium and metal atoms, but not between two deuterium atoms, to enable the location of partially occupied deuterium sites close to each other.

The metal-atom positions in NdNi₄MgD_{3.6} were determined from laboratory X-ray data. The anti-bump CF was helpful in this case to discard configurations with a very low R_{wp} when the quality of the diffraction pattern was not sufficient. When locating the deuterium positions (NdNi₄MgD_{3.6}, Zr₃NiO_{0.6}D_{6.32}, LaNi₂Mn₃D_{5.5}, ErFe₂D_{4.72}), all metal-atom positions were fixed, and an excess of deuterium atoms (with regards to the probable number of independent deuterium atoms, as expected from absorption results) was introduced randomly in the metal matrix. The deuterium positions were found quickly, with the excess D atoms merged. After removing the excess atoms from the model, the occupancies of the D positions were also optimized, which helped to find a better starting point for the Rietveld refinement (occupancies were generally correct within 10% of the refined values).

For ErFe₂D_{4.72}, the program *FOX* was used to confirm the D positions proposed from the search for available interstitial sites in the metal matrix (Paul-Boncour *et al.*, 2001). There were 20 possible D positions identified, possibly with partial occupancies. We have used 20 independent D atoms (with 100% occupancies) in a first optimization to find all actual positions. This yielded 10 independent positions actually occupied. It should be noted that the large DOF involved was compensated by the fact that all atoms optimized were identical (invariant by permutation), and the possible D positions were limited (through the use of the anti-bump CF) by the positions of the metal atoms. The occupancies of the ten D positions were then optimized. The final results, with eight sites with occupancies in the 70–100% range and two sites with lower occupancies, were in agreement (within 20% for occupancies) with the final results of the Rietveld refinement (except for the two sites with lower occupancies which were overestimated).

5.2. Molecular compounds

We have also conducted tests of previously known molecular compounds. The first test was made on potassium

tartrate $K^+ \cdot (2R,3R)C_4H_3O_4^-$, which was modeled by one free atom and the tartrate molecule using a Z matrix (H atoms were not included); no assumption was made about the absolute or relative configuration of asymmetric C atoms in the model. The DOC was not applied to speed up the optimization (generally, special positions or sharing of atoms between molecules are not present for molecular compounds, so the DOC should be avoided). The correct configuration is generally found after 1.2 million trials on average; the solutions reported correspond in half of the runs to the $(2R,3R)$ absolute configuration, and in the other half to the $(2S,3S)$ absolute configuration.

The other test was made on cimetidine ($C_{10}N_6H_{16}S$), which was used as a test case for *ab initio* structure determination from powder diffraction (Cernik *et al.*, 1991). The molecule was modeled using one Z matrix. The global optimization was tested both using the original powder pattern, and using extracted intensities; in both cases the correct solution was found in about 4 million trial configurations, with a 40% speed increase using extracted reflection intensities as no full powder pattern has to be generated for each trial. It is interesting to compare the results between the cimetidine and potassium tartrate examples: in both cases there are 14 DOF, but in the cimetidine example the number of trials required is about three times larger, which is due to the important flexibility in the middle of the cimetidine structure. This effect is related to the 'rigidity' of the modelization using a Z matrix: finding a 'rough' configuration of the molecule is relatively fast, but then the displacements of individual atoms to their final correct position cannot be done easily because of the direct parametrization from bond, bond angles and dihedral angles (see discussion in §3.1.2).

5.3. Hybrid structure

An interesting example is the structure of a novel aluminium methylphosphonate $[Al_2(CH_3PO_3)_3]$, which was solved using *FOX* (Edgar *et al.*, 2002), and illustrates the possibilities of the modular approach available in *FOX*. Determination of the structure was first attempted using the direct methods software package *SIRPOW* (Altomare *et al.*, 2002), but failed to lead to any partial structure solution, probably due to the difficulty of extracting reflection intensities from the diffraction pattern. For this material, individual 'building blocks' were identified using ^{27}Al , ^{13}C and ^{31}P NMR spectra, which revealed the presence of three crystallographically inequivalent C- PO_3 units as well as two aluminium atoms in tetrahedrally and fivefold coordinated sites, respectively. From this *a priori* knowledge, two models were used to solve the structure: (i) using three H_3C-PO_3 blocks and two Al atoms (24 DOF, not counting the position of H atoms), and (ii) one AlO_4 tetrahedron, one AlO_5 trigonal bipyramid, and three H_3C-P groups (DOF 27, counting only five DOF for the H_3C-P since H atoms positions cannot be found reliably). In both cases the same structure was obtained, and was then refined using *GSAS* (Larson & Von Dreele, 1987). The inclusion of the H atoms and the use of the anti-bump CF was very important in

these optimizations, since it helped differentiate C and O atoms, as well as (indirectly) P and Al, and therefore reduced the depth of local minima.⁴

The second modelization was much less successful during repeated runs (the distribution of the number of required trials to find the solution was wide, ranging from 300000 to 16 million), which is probably due to the difficulty in finding the correct positions of the AlO_4 and AlO_5 groups, whereas in modelization (i), there are two groups of identical building blocks (Al and H_3C-PO_3), and the invariance by permutation reduces significantly the parameter space to be explored. Another explanation for the different success rates and speed of convergence is the fact that the AlO_4 and AlO_5 groups are relatively isotropic scatterers (and therefore hard to orientate), whereas the H_3C-PO_3 are anisotropic and easy to place in the correct position and orientation. The role of the DOF (27 versus 24) is here less important than the nature of the building blocks.

6. The FOX program

6.1. Programming

6.1.1. *ObjCryst++*: object-oriented crystallographic computing library. The *FOX* program relies on the *ObjCryst++* library, which will be presented elsewhere (Favre-Nicolin, 2002; <http://objcryst.sourceforge.net/ObjCryst/>). All crystallographic entities are organized as high-level objects (Crystal, Scatterer, ScatteringPower, ScatteringData, OptimizationObj, etc.), which can independently define their set of parameters. The use of the object-oriented language C++ gives the library and program very good expandability through the use of inheritance and overloading. All objects can create graphical user interfaces (GUIs), which are used for the *FOX* program, and which rely on the cross-platform toolkit wxWindows (<http://www.wxwindows.org>).

6.1.2. Speed optimizations. Global-optimization algorithms require the evaluation of a very large number of trial configurations (from 10000 to infinity), since the search must be ergodic (*i.e.* explore the entire parameter space) to ensure finding the global minimum. Thus in programming *FOX*, emphasis was placed on the speed of calculation, employing storage of all results that can be re-used between successive trials. This is achieved by (i) delegating the computing as much as possible to objects, each of which is designed to avoid useless recalculations, and (ii) optimizing structure-factor calculations by grouping reflection data in vector arrays and using tabulated values for sine and cosine evaluation. The resulting speed achieved (Athlon 1.4 GHz computer running Linux) varies from 500 to 5000 trial structures per second, depending on the complexity of the structure, the space group, the number of reflections and the resolution of the powder pattern. For the cimetidine example, on 150 extracted intensities (single-crystal mode), about 3200 trial configurations per

⁴ With modelization (i) but without H atoms (and still the AB CF), finding the solution requires on average 1.3 million trials, *i.e.* almost twice as many trials than with the H atoms.

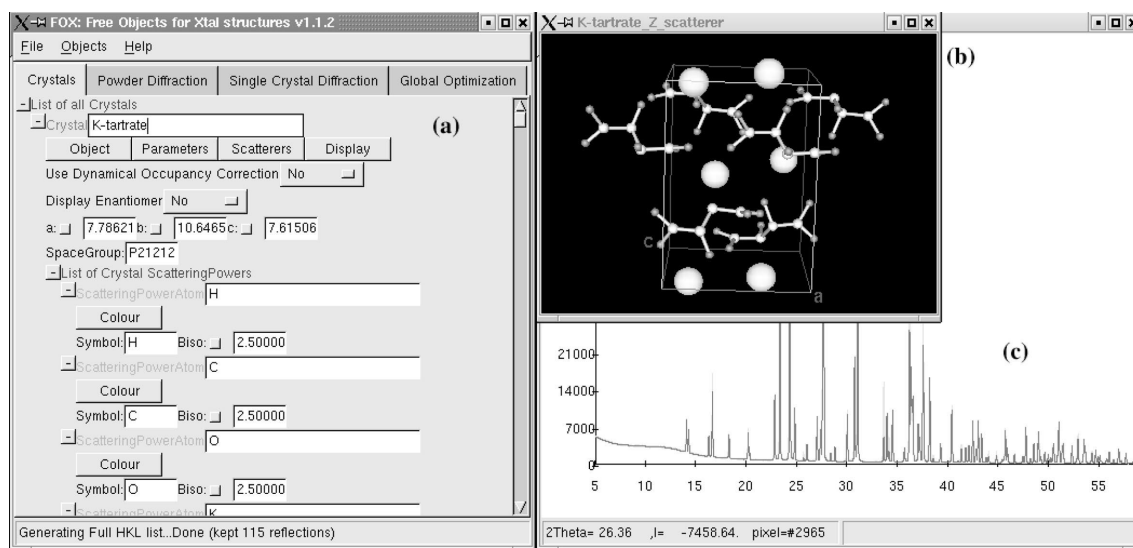


Figure 5

The interface of *FOX* consists of (a) one main window, in which crystal structure(s), diffraction experiment(s) and algorithms can be described. All crystal structures can be displayed in three dimensions in a separate window (b), and all powder diffraction patterns (calculated and observed) can be displayed in a window (c). All windows and (parameter values are updated during the optimization each time a new ‘best’ configuration is reached.

second can be evaluated (11 million per hour). Optimization on the powder pattern with the same number of reflections is approximately 40% slower (6 million per hour).

6.1.3. User interface. The *FOX* program uses a full graphical user interface (Fig. 5), which was written using the cross-platform framework wxWindows [http://www.wxwindows.org]. The main *FOX* window uses four panels: the first describes crystal structure(s) with unit-cell parameters, space-group symbols, a list of scattering powers, and a list of scatterers (either atoms, polyhedra or molecules); each structural parameter can be fixed or limited here. Crystal structures can be shown in three dimensions using a separate window. The second (respectively third) panel describes the powder (respectively single-crystal) diffraction experiments. The last panel describes the optimization algorithms: number of trials, temperature and mutation amplitude regime (SA and PT), with the ability to select objects (crystal structure, diffraction data) to optimize. The CF can be defined as any linear combination of the CF available for the chosen objects; thus if several diffraction data (e.g. X-rays and neutron patterns) are available for the same sample, they can be combined to perform a multi-pattern optimization. During the optimization process, all parameter values and graphics (powder pattern and three-dimensional crystal structure) are automatically updated each time a new best configuration is found. The complete optimization environment can be saved in a text file in extensible markup language (XML, http://www.w3.org/XML/).

6.1.4. Availability. *FOX* and the library *ObjCryst++* can be freely downloaded from http://objcryst.sourceforge.net with mirrors from the CCP14 project. It is currently available for Windows (98 and above) and Linux platforms. It is a free, open-source project (http://www.gnu.org/philosophy/philosophy.html) distributed under the GNU General Public License

(http://www.gnu.org/licenses/gpl.html) and therefore available as source code. It could be recompiled for other operating systems or processors.

6.1.5. Future developments. The use of an object-oriented programming model allows relatively easy expansion of the modelization of crystal structures, diffraction experiments and optimization algorithms. Among possible developments are (i) the improvement of modelization for molecular compounds using restraints, (ii) the development of new cost functions (e.g. energetic cost functions, constraints on the stoichiometry, etc.), (iii) the use of a learned electronic density or envelope (Brenner *et al.*, 2002) to guide the optimization, (iv) the use of new algorithms, and so on.

We acknowledge the following persons for their kind supply of either powder diffraction data or samples: Laure Guénée, Bernard Bertheville, Guillaume Renaudin and Gérald Bernardinelli from the University of Geneva, Switzerland, and Valerie Paul-Boncour from CNRS Thiais, France. Yuri Andreev from the University of St Andrews is acknowledged for a fruitful discussion on global-optimization methods. Ralf Grosse-Kunstleve from Lawrence Berkeley National Laboratory is acknowledged for help in using the *sginfo* and *atominfo* libraries. Lachlan Cranswick is acknowledged for useful suggestions and mirroring the project on the CCP14 Web site. The work was supported by the Swiss National Science Foundation through grant No. 2100-053847.98.

References

- Altomare, A., Giacovazzo, C., Ianigro, M., Moliterni, A. G. G. & Rizzi, R. (2002). *J. Appl. Cryst.* **35**, 21–27.
 Andreev, Y. G., MacGlashan, G. S. & Bruce, P. G. (1997). *Phys. Rev. B Cond. Matter*, **55**, 12011–12017.
 Boulton, A. & Louer, D. (1991). *J. Appl. Cryst.* **24**, 987–993.

- Brenner, S., McCusker, L. B. & Baerlocher, C. (2002). *J. Appl. Cryst.* **35**, 243–252.
- Bruker AXS GmbH (2000). *TOPAS User's Manual*.
- Caglioti, G., Paoletti, A. & Ricci, F. P. (1958). *Nucl. Instrum. Methods*, **3**, 223–226.
- Cascarano, G., Giacovazzo, C., Camalli, M., Spagna, R., Burla, M. C., Nunzi, A. & Polidori, G. (1984). *Acta Cryst.* **A40**, 278–283.
- Cascarano, G., Giacovazzo, C. & Guagliardi, A. (1991). *Acta Cryst.* **A47**, 698–702.
- Cernik, R. J., Cheetham, A. K., Prout, C. K., Watkin, D. J., Wilkinson, A. P. & Willis, B. T. M. (1991). *J. Appl. Cryst.* **24**, 222–226.
- Černý, R. (1998). *Adv X-ray Anal.* **40**, CD-ROM.
- Černý, R., Favre-Nicolin, V. & Bertheville, B. (2002). *Acta Cryst.* **C58**, i31–i32.
- Černý, R., Zavalij, I. Yu. & Koval'chuk, I. V. (2002). *J. Alloys Compd.* Submitted.
- David, W. I. F., Shankland, K. & Shankland, N. (1998). *Chem. Commun.* **8**, 931–932.
- Dollase, W. A. (1998). *J. Appl. Cryst.* **19**, 267–272.
- Edgar, M., Carter, V. J., Tunstall, D. P., Grewal, P., Sawers, L.-J., Favre-Nicolin, V., Cox, P. A., Lightfoot, P. & Wright, P. A. (2002). *Chem. Commun.* Submitted.
- Engel, G. E., Wilke, S., König, O., Harris, K. D. M. & Leusen, F. J. J. (1999). *J. Appl. Cryst.* **32**, 1169.
- Estermann, M. & Gramlich, V. J. (1993). *Appl. Cryst.* **26**, 396–404.
- Falcioni, M. & Deem, M. W. (1999). *J. Chem. Phys.* **110**, 1754–1766.
- Favre-Nicolin, V. (2002). In preparation.
- Grosse-Kunstleve, R. W., McCusker, L. B. & Baerlocher, Ch. (1997). *J. Appl. Cryst.* **30**, 985–995.
- Guénée, L., Favre-Nicolin, V. & Yvon, K. (2002). *J. Alloys Compd.* In the press.
- Guénée, L. & Yvon, K. (2002). *J. Alloys Compd.* Submitted.
- Harris, K. D. M., Tremayne, M., Lightfoot, P. & Bruce, P. G. (1994). *J. Am. Chem. Soc.* **116**, 3543–3547.
- Kariuki, B. M., Serrano-González, H., Johnston, R. L. & Harris, K. D. M. (1997). *Chem. Phys. Lett.* **280**, 189–195.
- Larson, A. & Von Dreele, R. (1987). Los Alamos National Laboratory, Report No. LA-UR-86-748.
- Lasocha, W. & Schenk, H. (1997). *J. Appl. Cryst.* **30**, 561–564.
- Le Bail, A. (2001). *Mater. Sci. Forum*, **378–381**, 65–70. (See also <http://sdpd.univ-lemans.fr/sdpd/espoir/>.)
- McGreevy, R. L. & Pusztai, L. (1988). *Mol. Simul.* **1**, 359–367.
- Metropolis, N., Rosenbluth, A. W., Rosenbluth, M. N., Teller, A. H. & Teller, E. (1953). *J. Chem. Phys.* **21**, 1087.
- Newsam, J. M., Deem, M. W. & Freeman, C. M. (1992). *Accuracy in Powder Diffraction II, NIST Special Publ.* No. 846, pp. 80–91.
- Paul-Boncour, V., Filipek, Y. M., Percheron-Guégan, A., Marchuk, I. & Peilassek, J. (2001). *J. Alloys Compd.* **58**, 317–318.
- Putz, H., Schön, J. C. & Jansen, M. (1999). *J. Appl. Cryst.* **32**, 864–870.
- Rodriguez-Carvajal, (1993). *Physica B*, **192**, 55.
- Shankland, K., David, W. I. F. & Csoka, T. (1997). *Z. Kristallogr.* **212**, 550–552.
- Shankland, K., David, W. I. F. & Sivia, S. (1997). *J. Mater. Chem.* **7**, 569–572.
- Sivia, D. S. & David, W. I. F. (1994). *Acta Cryst.* **A50**, 703–714.
- Soulié, J. P., Renaudin, G., Černý, R. & Yvon, K. (2002). *J. Alloys Compd.* In the press.
- Wessels, T., Baerlocher, Ch., McCusker, L. B. & Creyghton, E. J. (1999). *J. Am. Chem. Soc.* **121**, 6242–6247.
- Young, R. A. (1999). *The Rietveld Method*. International Union of Crystallography Monographs on Crystallography.