calculated using a similarity measure based on common lines with the already-angled set serving as a reference.

Some of the disadvantages of the angular reconstitution were addressed in the common-lines-based method for determining orientations for $N > 3$ particle projections simultaneously (Penczek *et al.*, 1996). In this method, the problem is formulated in terms of minimization of the variance of the 3D structure, as expressed in terms of common-lines discrepancy between N projections. In a sense, the design of the method is the exact opposite of the 'standard' common-lines approach: instead of trying to the determine the Eulerian angles (rotation matrices $\mathbf{R}_i$) based on angles α_{ij} of common lines in the planes of the projections, one assumes that rotation matrices $\mathbf{R}_i$ are known, finds the set of angles α_{ij} of common lines and computes the overall discrepancy along these lines. For a pair of projections i and j, the in-plane angles of common lines are found by solving the system of equations

$$\mathbf{R}_i\mathbf{n}_{ij} = \mathbf{R}_j\mathbf{n}_{ji} \tag{2.5.7.14}$$

for α_{ij} and α_{ji}. The discrepancy minimized in the method is the variance of the 3D structure that, by analogy to the 2D case (2.5.7.10) and (2.5.7.11) is

$$L_{\rm cl}(\{F\}, \{\mathbf{R}\}) = \sum_{m=0}^{M}\sum_{k=1}^{N} \left\| F_k(u_m, \alpha; \mathbf{R}_k) - \langle F\rangle_k \right\|^2 u_m^2 \Delta u \Delta\Omega_{kl},$$
$$\langle F\rangle_k = \frac{1}{N-1}\sum_{l=1, l\neq k}^{N} F_l(u_m, \alpha; \mathbf{R}_l), \tag{2.5.7.15}$$

where $L_{\rm cl}$ is written in Fourier 3D polar coordinates and $F_k(u_m, \alpha; \mathbf{R}_k)$ is the Fourier transform of the kth projection in 2D polar coordinates (u_m, α) with the orientation in 3D Fourier space given by the rotation matrix $\mathbf{R}_k$. All Fourier planes F_k are considered to have zero thickness, so all discrepancies are calculated only along common lines and the 'partial average' $\langle F\rangle_k$ is in fact an arrangement of $N-1$ Fourier planes in 3D space. An approximation to $u_m^2\Delta u\Delta\Omega_{kl}$ is calculated by equating the values of $\Delta\Omega_{kl}$ to the areas of the Voronoi diagram cells constructed on a unit sphere for points of intersection of common lines with this unit sphere (see Section 2.5.6.6). Generally the method performs very well, particularly if the projection images cover 3D angular space evenly.

Some macromolecules, particularly those that have an elongated barrel-like shape, will have a strongly preferential orientation with respect to the direction of electron beam showing only what are often called 'side views', *i.e.*, projections perpendicular to rotation along one axis corresponding to single-axis tilt data-collection geometry. These orthoaxial projections form single-axis tilt reconstruction geometry. In this case, Fourier transforms of all projections share only one common line, the line coinciding with the rotation axis, and clearly the common-lines-based method is not applicable to the *ab initio* structure determination. To cope with this situation, a method termed *Sidewinder* was developed (Pullan *et al.*, 2006). It is based on the observation that a Fourier transform of a finite object with a diameter D can be considered to have a nonzero thickness $1/D$ (Fig. 2.5.6.5). Thus, if the angle between two central sections of the 3D Fourier transform of this object, as derived from 2D Fourier transforms of its projections, is not too large, then these two sections will share information in Fourier space that is proportional to the amount of overlap of the two 'slabs' in Fourier space. Using this observation, the general idea employed in *Sidewinder* is to calculate pairwise cross-correlation coefficients (CCCs) between class averages of side views and to use this information to deduce the values of the azimuthal Eulerian angles using the Monte Carlo minimization method (Fishman, 1995).

For structures that have reasonably high symmetry and for those for which it is possible to collect high-quality EM data, it is sometimes possible to determine the initial structure using the 3D projection alignment method, which will be described in the next section. However, the approach is extremely computationally intensive and it is virtually impossible to try the method repeatedly to verify that the approach converges to more-or-less the same 3D structure, as is recommended for other *ab initio* methods described in this section. When the method is successful it is quite powerful, as an intermediate resolution structure can be obtained without going through intermediate and quite laborious steps of analysis of the data. A word of caution is warranted: with the direct method, unless there is external evidence that the obtained structure is correct, it is possible to obtain a self-consistent but entirely incorrect model of the molecule.

In the absence of reliable objective measures of the correctness of the structure, one can apply common sense in order to spot definitely improbable 3D maps. Given the mass of the complex it is possible to calculate the corresponding volume, and thus the threshold at which the map should be examined (Section 2.5.7.11). If at this threshold the mass density is discontinuous or there are pieces of mass surrounding the structures, the map is most likely to be incorrect. Similarly, strong directional artifacts appearing as streaks permeating the structure indicate that either the collected projection images are dominated by one or two views of the structure or that the angular assignment is incorrect. In addition, the 3D map should be centred in the window box; although the centring is not strictly speaking a mathematical requirement for a successful reconstruction, all single-particle structure-determination software packages take advantage of the fact that for centred objects orientation searches are easier to perform. So, if the map is not centred it is a clear indication of the failure of the procedure. Finally, for symmetric structures there should be no large pieces of mass on the symmetry axes.

2.5.7.9. Refinement of a 3D structure

Given an initial low-resolution model of the 3D structure and the data set of 2D projection images of the complex that have Fourier-space information extending beyond the resolution of the model, it is possible to refine the structure such that the full extent of the resolution information in the data will be utilized. In some cases, it is also possible to use as an initial structure in the refinement procedure a structure of a homologous protein, thus avoiding the process of *ab initio* structure determination altogether. The goal of the refinement is to find such orientation parameters for each of the particle projections for which (2.5.7.6) is minimized. There exist various implementations of the structure-refinement strategy and they can be roughly divided into those that perform exhaustive searches for all five orientation parameters (two translations and three Eulerian angles per 2D projection image) and those that perform local searches, usually by employing gradient information. Finally, the strategies may differ in how the correction for the CTF is implemented.

The original 3D projection-matching strategy (Penczek *et al.*, 1994) is based on the observation that given an ideal structure f and the necessary parameters of the CTF and image-formation model, it is straightforward to find five orientation parameters for each projection image. One begins with the determination of the sufficient angular step: assuming the structure is properly sampled at the Nyquist frequency and has a real-space radius of r voxels, the angular step is given by

$$\delta\theta \cong \arctan(1/r). \tag{2.5.7.16}$$

Next, keeping in mind that projection directions are parametrized by two Eulerian angles (φ, θ), one generates a set of projection directions quasi-uniformly distributed over half a unit

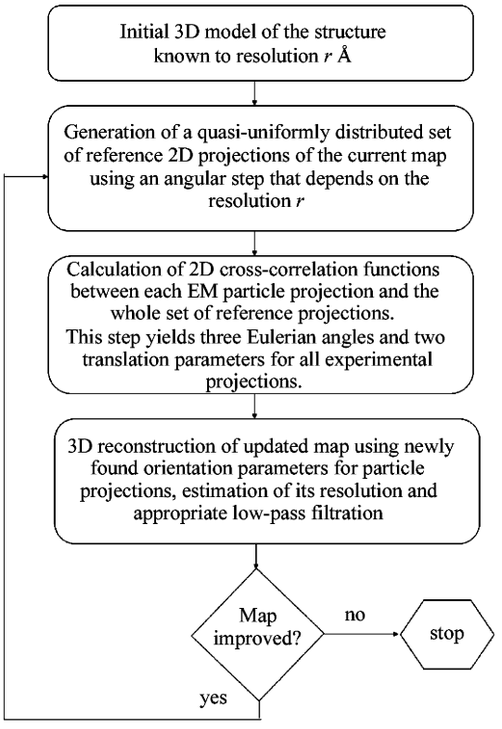

Fig. 2.5.7.4. Schematic of the 3D projection-alignment procedure.

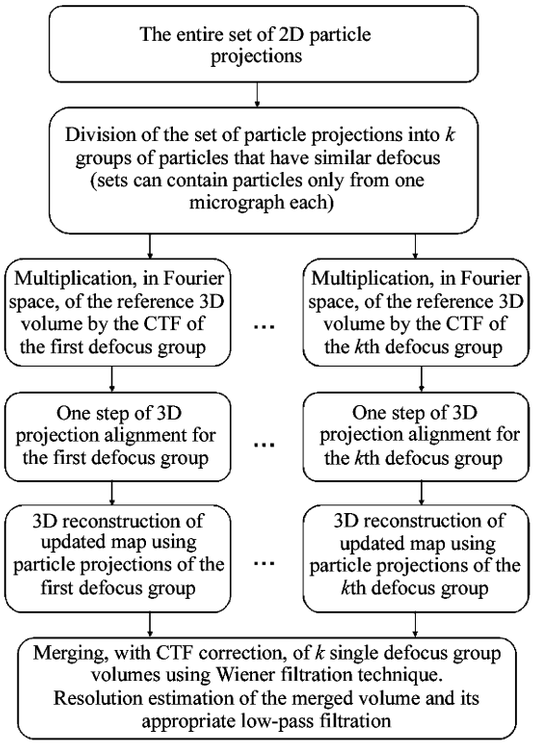

Fig. 2.5.7.5. Schematic of 3D projection alignment with CTF correction performed on the level of 3D maps reconstructed from projection images sorted into groups that share similar defocus settings.

sphere (or, in the case of a symmetric structure, over an asymmetric subunit) by taking fixed steps along the altitude or tilt angle θ and a number of samples azimuthally in proportion to $\sin\theta$ (Penczek *et al.*, 1994). So, for a chosen constant increment $\delta\theta$ and given θ angle the increment of the φ angle varies according to

$$\Delta\varphi = \delta\theta/|\sin\theta|. \qquad (2.5.7.17)$$

If all three Eulerian angles are to be sampled, as is necessary in some applications, then ψ is sampled uniformly in steps of $\delta\theta$.

In order to find the orientation parameters of projection images, one step of projection matching is performed. The reference structure is projected in all directions given by (2.5.7.17), yielding a set of reference images. Next, for each projection image, 2D cross-correlation functions with all reference images are calculated using one of the methods described in Section 2.5.7.6 and the overall maximum yields the translation, the in-plane rotation angle, the number of the most similar reference image (thus the remaining two Eulerian angles) and information about whether the image should be mirrored. Given this, a new 3D structure can be calculated using a 3D reconstruction algorithm (see Section 2.5.6). This simple protocol constitutes the core of 3D projection alignment (Fig. 2.5.7.4).

In a simple implementation of the 3D projection-matching procedure, all projection data are assembled into *defocus groups*, *i.e.*, groups of projection images that have similar defocus settings (Frank *et al.*, 2000). During refinement, for each defocus group the reference volume is multiplied by the CTF with the appropriate defocus value, one step of projection matching is performed and a refined structure is reconstructed for this group (Fig. 2.5.7.5). In addition, the within-group resolution is estimated using the Fourier shell correlation (FSC) approach (2.5.7.19) applied to two volumes calculated from two subsets of projection images randomly split into halves. After all defocus groups have been processed, the individual refined volumes are merged in Fourier space with a CTF correction using Wiener-filter methodology (Penczek *et al.*, 1997),

$$F_{\text{merged}} = \frac{\sum_k \text{CTF}_k \text{SSNR}_k F_k}{\sum_k \text{CTF}_k^2 \text{SSNR}_k + 1}, \qquad (2.5.7.18)$$

where SSNR_k is the spectral signal-to-noise ratio estimated for each defocus group using (2.5.7.22). Subsequently, the resolution of the merged volume is estimated by merging the half-volumes into two half-merged volumes using (2.5.7.18) and comparing them using (2.5.7.19). Next, the merged volume is filtered using (2.5.7.25) and the structure is centred so that its centre of mass is placed at the centre of the volume in which it is embedded.

The 3D projection-matching approach works very well during the initial stages of the refinement as it constitutes a very efficient approach to an exhaustive search for orientation parameters of all projection data. Once the orientation parameters are known to a degree of accuracy, it is straightforward to modify the procedure such that only subsets of reference projections are generated at a time and projection images are compared only with reference projections within a specified angular distance from their angular direction established during previous iteration. This modification speeds up the procedure significantly and makes it possible to refine structures to very high resolution by using a very small angular step $\delta\theta$. Another possible modification is to introduce an additional step of 2D alignment of the projection data that share the same angular direction (Ludtke *et al.*, 1999). The advantage is that this can correct possible errors of alignment to the projection of a limited-resolution reference structure and also, to an extent, reduces the danger of bias from artifacts in the reference structure. Finally, it is also possible to incorporate into the refinement strategy a correction for the envelope function of the microscope (Ludtke *et al.*, 1999). The 3D projection-matching strategy is widely popular and most EM software packages have implementations of various versions of basic strategies, as outlined above (Frank *et al.*, 1996; Ludtke *et al.*, 1999; Hohn *et al.*, 2007).

A possible improvement over the 3D projection-matching procedure can be achieved by working in transformed spaces in which the distinction between orientation search and 3D reconstruction is removed: (1) spherical harmonics (Provencher & Vogel, 1988; Vogel & Provencher, 1988), which have found applications exclusively in the determination of icosahedral structures (Yin *et al.*, 2001, 2003); (2) Radon transform (Radermacher, 1994), with selected applications in the determination of asymmetric particles (Ruiz *et al.*, 2003); or (3) Fourier transform, implemented in the *FREALIGN* package (Grigorieff, 2007). In *FREALIGN*, the transformation between the arbitrarily oriented Fourier 2D central section and the 3D Fourier Euclidean grid is implemented using trilinear interpolation that includes *ad hoc* correction for the CTF effects. In high-resolution structure-refinement mode, the program uses a gradient-based Powell optimization algorithm (Powell, 1973), thus overcoming the main deficiency of 3D projection-matching algorithms.

A unified approach to direct minimization of (2.5.7.6) was proposed by Yang *et al.* (2005) and is implemented in the *SPARX* package as the YNP method (Hohn *et al.*, 2007). The premise of the YNP method is that the orientation parameters are approximately known (thus the initial 3D map) and both the orientation parameters and the density map are updated simultaneously in a gradient-based optimization scheme. In the YNP method, the derivatives with respect to the density distribution are calculated analytically and the derivatives with respect to orientation parameters are calculated using finite difference approximations. The YNP method is very efficient and its major advantage is that it avoids many problems associated with approximate solutions inherent in methods that work in transform spaces. The projection/backprojection operations are carried out rapidly using linear interpolation, which due to sufficient oversampling of the data does not have a significant adverse impact on the solution. Moreover, because the density map f is updated simultaneously with the orientation parameters, the computationally demanding separate step of 3D reconstruction is eliminated.

2.5.7.10. Resolution estimation and analysis of errors in single-particle reconstruction

The development of resolution measures in EM was greatly influenced by earlier work in X-ray crystallography. In EM, the problem is somewhat more difficult as, unlike in crystallography, both the amplitude and the phase information in the data are affected by alignment procedures (which we consider distant analogues of phase-extension methods in crystallography). Therefore, resolution measures in EM reflect the self-consistency of the results; however, as the data are subject to alignment, there is a significant risk of introducing artifacts resulting from the alignment of the noise component in the data. Ultimately, these artifacts will unduly 'improve' the resolution of the map.

The resolution measures used in EM fall into two categories: measures based on averaging of Fourier transforms of individual images and measures based on comparisons of averages calculated for subsets of the data. In the first group, we have the Q-factor (van Heel & Hollenberg, 1980; Kessel *et al.*, 1985) and the spectral signal-to-noise ratio (SSNR) introduced for the 2D case by Unser and co-workers (Unser *et al.*, 1987), and for the 3D case for a class of reconstruction algorithms data are based on direct Fourier inversion by Penczek (Penczek, 2002). The second group of measures includes the differential phase residual (DPR) (Frank *et al.*, 1981) and the Fourier ring correlation (FRC) (Saxton & Baumeister, 1982). A marked advantage of these measures is that they are equally well applicable to 2D or 3D data. In the latter case, the volumes resulting from 3D reconstruction algorithms take the place of the 2D averages.

The resolution measures used in single-particle reconstruction are designed to evaluate the SSNR in the reconstruction as a function of spatial frequency (Penczek, 2002). The 'resolution' of the reconstruction is reported as a spatial frequency limit beyond which the SSNR drops below a selected level, for example below one.

The FSC is evaluated by taking advantage of the large number of single-particle images: the total data set is randomly split into halves; for each subset a 3D reconstruction is calculated (in two dimensions, a simple average); and two maps f and g are compared in Fourier space,

$$\mathrm{FSC}(f, g; u) = \frac{\sum_{|\|\mathbf{u}_n\|-r|\le\varepsilon}^{n_r} F(\mathbf{u}_n)G^*(\mathbf{u}_n)}{\left\{\left[\sum_{|\|\mathbf{u}_n\|-r|\le\varepsilon}^{n_r} |F(\mathbf{u}_n)|^2\right]\left[\sum_{|\|\mathbf{u}_n\|-r|\le\varepsilon}^{n_r} |G(\mathbf{u}_n)|^2\right]\right\}^{1/2}}. \quad (2.5.7.19)$$

In (2.5.7.19), 2ε is a preselected ring/shell thickness, the $\mathbf{u}_n$ form a uniform grid in Fourier space, $u = \|\mathbf{u}_n\|$ is the magnitude of the spatial frequency and n_r is the number of Fourier voxels in the shell corresponding to frequency u. The FSC yields a 1D curve of correlation coefficients as a function of u. Note that the FSC is insensitive to linear transformations of the densities of the objects. An FSC curve everywhere close to one reflects strong similarity between f and g; an FSC curve with values close to zero indicates the lack of similarity between f and g. Particularly convenient for the interpretation of the results in terms of 'resolution' is the relation between the FSC and the SSNR, which is easily derived by taking the expectation of (2.5.7.19) under the assumption that both f and g are sums of the same signal and different realizations of the noise, which are uncorrelated with the signal and between them (Saxton, 1978):

$$E[\mathrm{FSC}] \cong \frac{\mathrm{SSNR}}{\mathrm{SSNR}+1}. \quad (2.5.7.20)$$

By solving (2.5.7.20) for SSNR we obtain

$$\mathrm{SSNR} = \frac{\mathrm{FSC}}{1-\mathrm{FSC}}, \quad (2.5.7.21)$$

which, taking into account that the FSC was calculated from the data set split into halves, has to be modified to (Unser *et al.*, 1987)

$$\mathrm{SSNR} = 2\left(\frac{\mathrm{FSC}}{1-\mathrm{FSC}}\right). \quad (2.5.7.22)$$

In order to calculate the FSC that corresponds to a given SSNR, one inverts (2.5.7.22) to

$$\mathrm{FSC} = \frac{\mathrm{SSNR}}{\mathrm{SSNR}+2}. \quad (2.5.7.23)$$

Equations (2.5.7.21) and (2.5.7.22) serve as a basis for various 'resolution criteria' used in EM. The often-used 3σ criterion (van Heel, 1987*b*) equates resolution with the point at which the FSC is larger than zero at a 3σ level, where σ is the expected standard deviation of the FSC that has an expected value of zero, in essence finding a frequency for which the SSNR is significantly larger than zero. The 3σ criterion has a distinct disadvantage of reporting the resolution at a frequency at which there is no significant signal, while tempting the user to interpret the detail in the map at this resolution. Moreover, as the FSC approaches zero, its relative error increases, so the curve oscillates widely