different initial velocities and/or temperatures can be taken for further refinement or structure-factor averaging (see below).

The annealing schedule can, in principle, be any function of the simulation step (or 'time' domain). The two most commonly used protocols are linear slow-cooling or constant-temperature followed by quenching. A slight advantage is obtained with slow cooling (Brünger *et al.*, 1990). The duration of the annealing schedule is another parameter. Too short a protocol does not allow sufficient sampling of conformational space. Too long a protocol may waste computer time, since it is more efficient to run multiple trials than one long refinement protocol (unpublished results).

18.2.4.4. *An intuitive explanation of simulated annealing*

The goal of any optimization problem is to find the global minimum of a target function. In the case of crystallographic refinement, one searches for the conformation or conformations of the molecule that best fit the diffraction data and that simultaneously maintain reasonable covalent and non-covalent interactions. Simulated-annealing refinement has a much larger radius of convergence than conjugate-gradient minimization (see below). It must, therefore, be able to find a lower minimum of the target E [equation (18.2.3.1)] than the local minimum found by simply moving along the negative gradient of E.

It is most easy to visualize this property of simulated annealing in the case of a one-dimensional problem, where the goal is to find the global minimum of a function with multiple minima (Fig. 18.2.4.1). An intuitive way to understand a molecular-dynamics simulation is to envisage a ball rolling on this one-dimensional surface. When the ball is far from the global minimum, it gains a certain momentum which allows it to cross barriers of the target function [equation (18.2.4.3)]. Slow-cooling temperature control ensures that the ball will eventually reach the global minimum rather than just bouncing across the surface. The initial temperature must be large enough to overcome smaller barriers, but low enough to ensure that the system will not escape the global minimum if it manages to arrive there.

While temperature itself is a global parameter of the system, temperature fluctuations arise principally from local conformational transitions, for example, from an amino-acid side chain falling into the correct orientation. These local changes tend to lower the value of the target E, thus increasing the kinetic energy, and hence the temperature, of the system. Once the temperature control has removed this excess kinetic energy through 'heat dissipation', the reverse transition is very unlikely, since it would require a localized increase in kinetic energy where the conformational change occurred in the first place (Fig. 18.2.4.1). Temperature control maintains a sufficient amount of kinetic energy to allow local conformational corrections, but does not supply enough to allow escape from the global minimum. This explains the observation that, on average, the agreement with the diffraction data will improve, rather than worsen, with simulated annealing.

18.2.5. Examples

Many examples have shown that simulated-annealing refinement starting from initial models (obtained by standard crystallographic techniques) produces significantly better final models compared to those produced by least-squares or conjugate-gradient minimization (Brünger *et al.*, 1987; Brünger, 1988; Fujinaga *et al.*, 1989; Kuriyan *et al.*, 1989; Rice & Brünger, 1994; Adams *et al.*, 1997). In another realistic test case (Adams *et al.*, 1999), a series of models for the aspartic proteinase penicillopepsin were generated from homologous structures present in the Protein Data Bank. The sequence identity among these structures ranged from 100% to 25%, thus providing a set of models with increasing coordinate error compared to the refined structure of penicillopepsin. These models, after truncation of all residues to alanine, were all used as search models in molecular replacement against the native penicillopepsin diffraction data. In all cases, the correct placement of the model in the penicillopepsin unit cell was found.

Both conjugate-gradient minimization and simulated annealing were carried out in order to compare the performance of the $E^{\rm LSQ}$ least-squares residual [equation (18.2.3.2)], MLF (the maximum-likelihood target using amplitudes) and MLHL (the maximum-likelihood target using amplitudes and experimental phase information). In the latter case, phases from single isomorphous replacement (SIR) were used. A very large number of conjugate-gradient cycles were carried out in order to make the computational requirements equivalent for both minimization and simulated annealing. The conjugate-gradient minimizations were converged, *i.e.* there was no change when further cycles were carried out.

For a given target function, simulated annealing always outperformed minimization (Fig. 18.2.5.1). For a given starting model, the maximum-likelihood targets outperformed the least-squares-residual target for both minimization and simulated annealing, producing models with lower phase errors and higher map correlation coefficients when compared with the published penicillopepsin crystal structure (Fig. 18.2.5.1). This improvement is illustrated in σ_A-weighted electron-density maps obtained from the resulting models (Fig. 18.2.5.2). The incorporation of experimental phase information further improved the refinement significantly despite the ambiguity in the SIR phase probability distributions. Thus, the most efficient refinement will make use of simulated annealing and phase information in the MLHL maximum-likelihood target function.

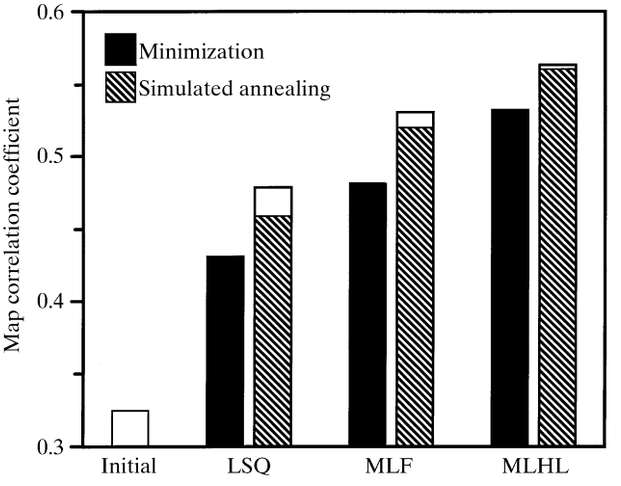

Fig. 18.2.5.1. Simulated annealing produces better models than extensive conjugate-gradient minimization. Map correlation coefficients were computed before and after refinement against the native penicillopepsin diffraction data (Hsu *et al.*, 1977) for the polyalanine model derived from Rhizopuspepsin (Suguna *et al.*, 1987, PDB code 2APR). Correlation coefficients are between σ_A-weighted maps calculated from each model and from the published penicillopepsin structure. The observed penicillopepsin diffraction data were in space group $C2$ with cell dimensions $a = 97.37$, $b = 46.64$, $c = 65.47$ Å and $\beta = 115.4°$. All refinements were carried out using diffraction data from the lowest-resolution limit of 22.0 Å up to 2.0 Å. The MLHL refinements used single isomorphous phases from a $K_3UO_2F_5$ derivative of the penicillopepsin crystal structure, which covered a resolution range of 22.0 Å to 2.8 Å. Simulated-annealing refinements were repeated five times with different initial velocities. The numerical averages of the map correlation coefficients for the five refinements are shown as hashed bars. The best map correlation coefficients from simulated annealing are shown as white bars.

references