References

AL HADDAD, M. & BECKER, P. J. (1988). *Acta Cryst.* A**44**, 262-270.

BECKER, P. J. & AL HADDAD, M. (1989*a*). *Acta Cryst.* A**45**. Submitted.

BECKER, P. J. & AL HADDAD, M. (1989*b*). *Acta Cryst.* A**45**. Submitted.

FELLER, W. (1970). *An Introduction to Probability Theory and its Applications.* New York: John Wiley.

KATO, N. (1973). *Z. Naturforsch. Teil A*, **28**, 604-609.

KATO, N. (1980). *Acta Cryst.* A**36**, 763-778.

TAKAGI, S. (1962). *Acta Cryst.* **15**, 1311-1312.

TAKAGI, S. (1969). *J. Phys. Soc. Jpn*, **26**, 1239-1253.

VAN KAMPEN, N. G. (1981). *Stochastic Processes in Physics and Chemistry.* Amsterdam: North-Holland.

ZIMAN, J. M. (1979). *Models of Disorder. The Theoretical Physics of Homogeneously Disordered Systems.* Cambridge Univ. Press.

Acta Cryst. (1989). A**45**, 337-342

A Simulated Annealing Approach to the Search Problem of Protein Crystallography

BY S. SUBBIAH* AND S. C. HARRISON†

Committee on Higher Degrees in Biophysics and Howard Hughes Medical Institute, 7 Divinty Avenue, Harvard University, Cambridge, MA 02138, *USA*

(*Received* 26 *July* 1988; *accepted* 20 *December* 1988)

Abstract

With the growing availability of computer power it has become routine to perform exhaustive multi-dimensional searches in protein crystallography. Specifically, in cases where homologous or partially homologous structures are available, the initial interpretation of poor electron density maps is done by performing computer-intensive rotational and translational searches in real space. Often such calculations of the best fit between structure and map cannot even be attempted owing to the vast computing effort involved (years of MicroVAX II time). Here, the combinatorial optimization method, simulated annealing, is shown to reduce substantially the computing effort involved and also to permit computations that are beyond the reach of current algorithms. This is illustrated with practical examples involving the structure determinations of the human histocompatibility antigen HLA-A2 and an influenza virus hemagglutinin-sialic acid complex.

1. Introduction

In cases where poor and not-readily interpretable electron density (e.d.) maps of protein structures are available, it has become common to perform real-space rotation and translation searches using homologous or partially homologous protein structures as search objects against the available map (Huber, 1965; Colman & Webster, 1985; Reynolds, Remington, Weaver, Fisher, Anderson, Ammon & Matthews, 1985). These exhaustive searches in general require large amounts of computer time and often such searches cannot be attempted for this reason. However, when attempted they typically require on the order of hundreds of MicroVAX II c.p.u. hours (Colman & Webster, 1985; Reynolds *et al.*, 1985) (*NB* a Cray I is on the order of 100 times faster than a MicroVAX II). Even in these cases it may be necessary to approximate the search object to the α-carbon backbone alone. Such exhaustive calculations ascertain the best-fit six-dimensional orientation of the search object in the map. However, depending on the appropriateness of the structural homologue and the 'noise' level of the map, the optimal solution may not necessarily correspond to the true solution of the crystallographic problem. Nonetheless, the correct solution is always found to be among the better solutions to the search procedure. In short, the problem is that of using the least-approximate search object and to obtain a list of good solutions to the search problem while minimizing the amount of computer effort.

The complex optimization technique of *simulated annealing* (also known as the Metropolis algorithm) is shown in this paper to have these desired qualities (Metropolis, Rosenbluth, Rosenbluth, Teller & Teller, 1953). This method, which is a variant of the Monte Carlo method, has been applied in recent years with much success to large optimization problems ranging from spin-glass theory in solid-state physics to the classic travelling salesman problem of computer science (Van Hemmen & Morgenstern, 1983; Kirkpatrick, Gelatt & Vecchi, 1983). It has also recently been found to be an effective tool in the

* Committee on Higher Degrees in Biophysics.

† Committee on Higher Degrees in Biophysics and Howard Hughes Medical Institute.

0108-7673/89/050337-06$03.00

crystallographic refinement of macromolecular structures (Brunger, Kuriyan & Karplus, 1987). However, while the algorithm guarantees the generation of a list of very good solutions in very significantly reduced times, owing to the nature of these problems one cannot be certain that the global optimum has been found.

In the context of the crystallographic search problem this uncertainty has not been found to be a difficulty in practice. In the two examples considered in this paper, the optimum in either case was arrived at in a relatively short time. The first case deals with the most objective positioning of a 21-atom sialic acid model in the receptor binding site of the influenza virus hemagglutinin–sialic acid complex electron density map (Weis, Brown, Cusack, Paulson, Skehel & Wiley, 1988). The six-dimensional exhaustive rigid-body rotational and translational search against a map at 2·9 Å took approximately 6·5 MicroVAX II c.p.u. hours. Of the 67 good solutions in the top 2σ range, most were attained by the simulated annealing protocol in approximately 25-fold less time. The four top 1σ best solutions were all found. The second problem involved the placement of a partially built HLA-A2 model with 87 atoms in an electron density map obtained from a crystal form different from that in which the model was originally built (Bjorkman, Saper, Samraoui, Bennett, Strominger & Wiley, 1987). Here, where a thorough exhaustive search would have been impossible with current computing resources, the annealing protocol proved to be approximately 88 000-fold faster. This example serves to illustrate another property of the simulated annealing protocol: the relative scaling between computer effort and size of the tackled problem is less than linear. Thus this approach to the crystallographic search problem affords the possibility of overcoming more complex problems.

Finally it is worth noting that this approach deals with a stage that naturally precedes the crystallographic refinement stage of any X-ray structure determination. Therefore it is highly complementary to the recent developments in crystallographic refinement that also employ simulated annealing (Brunger *et al.*, 1987).

2. The Metropolis algorithm – simulated annealing

Given a vast solution for a quantiy $E(x)$ that requires minimization (or maximization) over the range of variable x, the simulated annealing protocol proceeds along the following steps.

(a) A random start point x_0 is chosen and $E(x_0)$ is computed.

(b) A single small displacement Δx is made and $E(x_0+\Delta x)$ is evaluated.

(c) $\Delta E = E(x_0+\Delta x) - E(x_0)$ is computed.

(i) If ΔE is negative or equal to 0, the new value of x is accepted, x_0 is replaced by $x_0+\Delta x$ and the process returns to step (b) for another cycle.

(ii) If $\Delta E \geq 0$ a random number r lying between 0 and 1 is generated.

If $r < \exp-(\Delta E/T)$, the new value of x is accepted, x_0 is replaced by $x_0+\Delta x$ and the process returns to step (b) for another cycle.

If $r > \exp-(\Delta E/T)$, the new value of x is rejected and the process returns to step (b) for another cycle with x_0 unaltered.

The cycle outlined above is controlled by the parameter T. The value of T governs the probability of accepting displacements, Δx, that in the short run do not decrease E. At the outset of the annealing procedure, T is set to be very much greater than the typical ΔE's. A large T allows nearly all random displacements to be accepted with little regard to the primary objective of reducing E. However, after a fixed number, I_c, of passes through step (b), the value of T is reduced by some factor - typically $T_{new} = gT_{old}$, where g is some fraction between 0 and 1 like 0·9. Then another I_c cycles are attempted at this new T. These sets of I_c cycles at decreasing values of T are repeated until T becomes much smaller than the ΔE's encountered. Under this condition, the procedure is guaranteed to reach a value of x that is stable over iterations on the order of I_c. In the limit of I_c tending to infinity, this value of x corresponds to the global minimum. Otherwise it corresponds to some minimum, very likely a local minimum.

Thus, in a physical sense, T can be viewed as a kind of *temperature*. Hence, the annealing proceeds as successive rounds of equilibration at decreasing temperatures. The primary reason for the success of the scheme is its built-in ability to move in unfavourable directions in the short run. This overcomes the difficulty of being trapped in the nearest local minimum. In practice, it is worthwhile keeping a list of the good solutions encountered along the annealing path. Estimates of the goodness of solutions can be pre-determined by performing a very coarse exhaustive search to obtain approximate values for the mean $\langle E \rangle$ and the standard deviation σ_E of the solution space. This list of solutions can later be minimized independently to their respective nearest minima and these can be included in the final set of good solutions.

The parameters of the annealing protocol are largely determined empirically and are characteristic of the general topology of the particular solution space. However, in theory, there is a lower bound on the value of I_c. The expected distance traversed by a random walk of I_c steps in an l-dimensional grid is of order $(I_c/l)^{1/2}$ (Van Kampen, 1981). Thus, to give the annealing scheme an optimal chance of equilibrating at any temperature, the number of annealing steps, I_c, must be such that $(I_c/l)^{1/2}$ is greater than half the

longest edge of the search space. Further, the value of T_0 can be selected by empirically determining typical ΔE's and arranging for T_0 to be greater than this. g, which governs the rate of cooling, has been shown empirically to give consistently good results in a wide range of problems when lying between 0·9 and 1·0 (Kirkpatrick *et al.*, 1983). All these parameters - T_0, g and the ratio by which $(I_c/l)^{1/2}$ exceeds the longest edge of the search space - are largely determined by the nature of the solution space and consequently by the type of problem that generates this space. Hence, when once determined for a given type of problem, for instance the search problem of crystallography, these parameters are expected to remain relatively constant.

3. The crystallographic search problem

The problem is a general one of placing a rigid molecular structure in an electron density map in the most objective fashion. There are a number of different types of such problems in protein crystallography. One particular instance is the placement of small substrate (or equally well an inhibitor) structures in poor e.d. maps of protein binding sites. Another case is when homologue structures are used in the initial location of new protein structures within the unit cell of poor e.d. maps. The location of a known structure within a poor e.d. map from a different space group is yet another variant.

The size, C, of all such search problems is determined by the following parameters.

(a) The dimensions of the e.d. map to be searched - a, b and c. Often this is no more than the asymmetric unit.

(b) The resolution of the available map, k.

(c) The number of atoms in the search object, n.

Ideally, a complete six-dimensional search would involve an *inner* exhaustive translational search for every rotational setting. At each such configuration of the search object some measure of the quality of fit has to be calculated. Given the limited computer resources available, this measure is often simply the sum of the electron density at each of the n atomic locations, for any given configuration of the search object,

$$E = \sum_{i=1}^{n} (\text{e.d. at atom } i).$$

Each such determination of electron density corresponding to a single atom can be viewed as a basic computational operation (b.c.o.). The size of the crystallographic search problem, C, can be estimated in terms of the number of b.c.o. required.

(1) Using the standard 1/3 resolution criterion for interpolation, the number of translational configurations is

$$T \simeq (abc)(k/3)^3.$$

(2) With reference to Fig. 1, the number of rotational configurations is

$$W \simeq (2\pi/\Delta\theta_1)(2\pi/\Delta\theta_2)(\pi/\Delta\theta_3)$$
$$\simeq \tfrac{1}{2}abc(3\pi/k)^3.$$

(3) Therefore, the size of the crystallographic search problem is

$$C \simeq nTW \simeq \tfrac{1}{2}n\pi^3(abc)^2(3/k)^6.$$

4. Implementation of protocol

The particular implementation of the simulated annealing protocol for the crystallographic search problem is outlined in detail in the accompanying flow diagram (Fig. 2). The one obvious difference is that the quantity $E_{xyz\theta_1\theta_2\theta_3} = \sum_{i=1}^{n}$ (e.d. at atom i) is maximized rather than minimized. Therefore ΔE is assumed to be $\Delta E = -(E_{\text{new}} - E_{\text{old}})$. All rotations in the three Euler coordinates θ_1, θ_2 and θ_3 are done about the origin of the e.d. map prior to translation to the search configuration considered. In all cases, a very coarse search ($C \simeq n \times 10^6$) was conducted beforehand to estimate both the mean, $\langle E \rangle$, and the standard deviation, σ_E, of the solution space. Prior experience with exhaustive searches suggested $3\sigma_E$ as a reasonable lower cutoff for the assessment of the goodness of any candidate solution. Accordingly all such solutions encountered along the annealing path were recorded.

This particular implementation of the general annealing protocol described earlier involves some modification in the parameter I_c. At each value of T at most I_1 steps were allowed. Since all steps result in either the new configuration being accepted or rejected, the sum of the number of acceptances, I_a, and the number of rejections, I_r, must equal I_c. When I_a exceeded some pre-set value I_2 (arranged to be less than I_1), T was decreased and the next round of equilibration was initiated. Such equilibration rounds, where I_2 was reached before I_1, were deemed

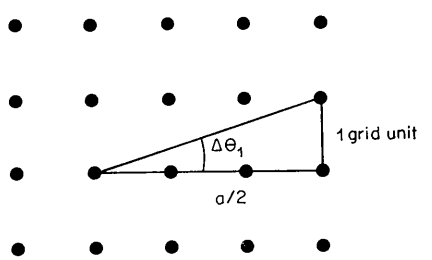

Fig. 1. Given a resolution k and using the standard $\frac{1}{3}$ resolution criterion for the choice of grid spacing of an electron density map, the arc length = $r\theta$ formula can be used to estimate the required angular interval $\Delta\theta_1$ for a real-space rotational search. With a the length of an edge of the map to be searched, $(a/2)\Delta\theta_1 = 1$ grid unit $= k/3$ implies that the number of intervals for a complete search of 2π rad $= 3\pi a/k$.

cooling cycles. Conversely, when I_1 steps were completed without obtaining I_2 acceptances a *freezing* cycle was said to have occurred. Three consecutive cycles of freezing were required for termination of the annealing protocol. The optimal ratio of I_1 to I_2 was empirically determined to be 2. Similarly, the optimal values for g and T_0 were found to be in the range 0·98–0·99 and 2–3σ_E respectively.

Once the annealing was complete, both the converged solution and the list of good solutions (*i.e.* $> 3\sigma_E$) were independently minimized by a standard least-squares minimization procedure to their nearest local minima. These minima were all confirmed to be truly independent minima and duplicate minima were dropped from the list. This pruned list was accepted as the desired collection of best solutions for any given random start.

5. Experiment 1

(*a*) *Problem*

The structure of the bromelain-related influenza virus hemagglutinin protein was solved to 3·0 Å resolution by Wilson, Skehel & Wiley (1981). This glycoprotein, a trimer of 503 × 3 amino acids binds to sialic-acid-containing receptors. The crystal structure of a 21-atom sialic acid–hemagglutinin complex has been recently solved to 2·9 Å (Weis *et al.*, 1988). The placement of the sugar in the binding-site electron density was confirmed by performing a six-dimensional rigid-body search. Here, while the $P4_1$ asymmetric unit is 163·2 × 163·2 × 177·4 Å, the portion of the map that adequately includes the receptor binding site is much smaller – 13·53 × 10·63 × 13·53 Å. Accordingly, the parameters of the problem were: $a = 14$, $b = 11$, $c = 14$, $k = 2{\cdot}9$, $n = 21$ and $C = 6{\cdot}1 \times 10^8$ b.c.o. This search required 6·5 h of c.p.u. time on a MicroVAX II. There were four good solutions in the top σ of the distribution. In the top 2σ range there were 67 independent good solutions. Of these solutions, the top solution happened also to make the most chemical sense. However, in principle this correspondence between the objective best fit and the chemically correct solution need not occur. This solution has since been confirmed by crystallographic refinement and NMR (Weis *et al.*, 1988).

To what extent can simulated annealing reproduce these results? How much faster will simulated annealing achieve this?

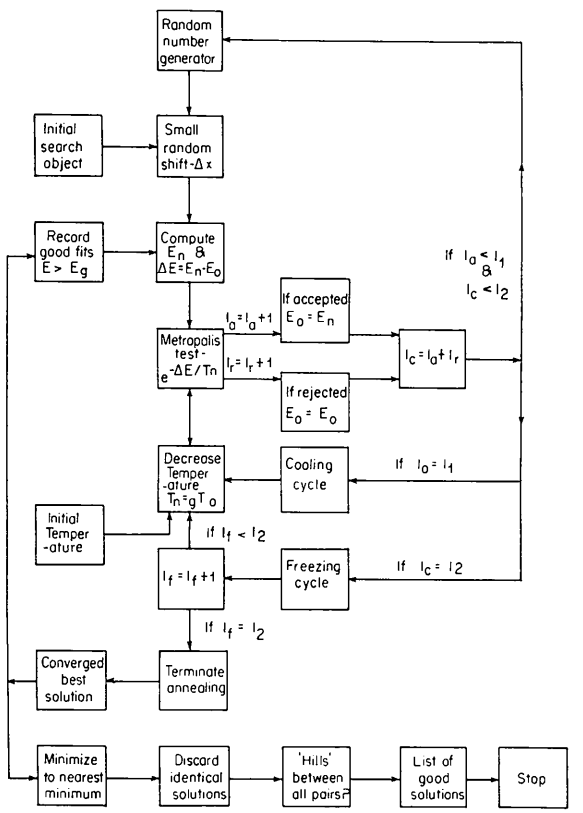

Fig. 2. A flow diagram illustrating the particular implementation of the simulated annealing protocol for the real-space search problem of crystallography. All variable and parameter names are as described in the text. The inputs to the program include initial values for the various parameters (*e.g.* g, T_0, I_1, I_2), the coordinates of the search model and a suitably large search map.

Table 1. *The results of an exhaustive six-dimensional search of a sialic acid model against an influenza hemagglutinin–sialic acid complex e.d. map compared with those from a number of simulated annealing trails*

The number of good solutions to the exhaustive search lying in the top 1, 1·5 and 2 σ ranges are presented in the first row. The results of the five different random trials are presented in the same ranges. (These are also expressed as percentages of the exhaustive search results.) The same data from the five trials are also presented cumulatively.

	Number of peaks		
	Top 1·0σ	Top 1·5σ	Top 2·0σ
Exhaustive	4 (100%)	26 (100%)	67 (100%)
Random trial 1	2 (50%)	7 (27%)	17 (25%)
Random trial 2	2 (50%)	7 (27%)	14 (21%)
Random trial 3	0 (0%)	6 (23%)	12 (18%)
Random trial 4	2 (50%)	9 (35%)	19 (28%)
Random trial 5	0 (0%)	3 (12%)	12 (18%)
Cumulative (1-2)	4 (100%)	13 (50%)	29 (43%)
Cumulative (1-3)	4 (100%)	17 (65%)	36 (54%)
Cumulative (1-4)	4 (100%)	21 (81%)	45 (67%)
Cumulative (1-5)	4 (100%)	22 (85%)	50 (75%)

(*b*) *Results*

Table 1 presents the results of five consecutive trials of the annealing protocol starting from five different random positions. The parameters used were $\langle E \rangle = 0{\cdot}0$, $\sigma_E = 7{\cdot}5$, $g = 0{\cdot}98$, $T_0 = 2\sigma_E$, $I_2 = 10\,000$, $I_1 = 2I_2$. The cumulative success of obtaining the pool of good solutions, as ascertained by the exhaustive search above, is indicated in fractional form. Each trial took approximately 3 c.p.u. min on the MicroVAX II.

Within two trials, the top 1σ peaks were attained. After five trials, approximately 75% of all the top 67 2σ peaks were obtained. Thus, in this case, the top σ solutions were generated with some 80-fold reduction in computer time. For lesser peaks, top 2σ, savings were more on the order of 25-fold.

6. Experiment 2

(*a*) *Problem*

The crystal structure of HLA-A2 has been recently determined to 3·5 Å resolution (Bjorkman *et al.*, 1987). The structure was obtained after averaging e.d. maps from two space groups, monoclinic $P2_1$ and orthorhombic $P2_12_12_1$. The initial orientation was established in the monoclinic form by two independent methods. The relatively poor map was traced by eye, while concurrently an exhaustive computer search was performed using the 101 α-carbon locations of the homologous immunoglobulin constant region (*i.e.* C_{H3} from human Fc was expected on grounds of sequence similarity to have the same fold as the β_2-microglobulin domain of HLA-A2) (Bjorkman *et al.*, 1987). The same trace was confirmed. Following this, the relative orientation of the partially traced $P2_1$ model was sought in the $P2_12_12_1$ e.d. map. Packing considerations that suggested a possible relative orientation failed to give clear results. Thereafter an exhaustive search based on a fortuitous choice of a coarse grid using a model of 89 α-carbon atoms resulted in the unambiguous discovery of this relative orientation (Bjorkman, personal communication). This search was conducted on a very coarse grid spacing of 3 Å and required approximately 150 MicroVAX II c.p.u. h. However, a grid spacing of at least 1 Å would certainly have been necessary if the model had not been in a favourable location relative to the coarse grid spacing used. Such a search could not have been contemplated as estimates of the required MicroVAX II processor time would have been on the order of 18 years. The parameters of this hypothetical 18-year search are $a = 60$, $b = 40$, $c = 60$, $k = 3$, $n = 89$ and $C = 2{\cdot}9 \times 10^{18}$ b.c.o.

Can simulated annealing address this problem? If so, then how effectively?

(*b*) *Results*

Five trials were attempted from five different random positions, with the following parameters: $\langle E \rangle = 0{\cdot}0$, $\sigma_E = 32{\cdot}5$, $g = 0{\cdot}98$, $T_0 = 3\sigma_E$, $I_2 = 25\,000$ and $I_1 = 2I_2$ (Table 2). In all cases the optimum, at 10σ, corresponding to the true (crystal-packing) solution was attained in approximately 4 h on the MicroVAX II. Other good solutions were found. However, since the exhaustive search cannot be performed in practice, no assessment of the relative success in attaining these can be made. These results nominally suggest a 88 000-fold increase in speed.

Table 2. *The results of a six-dimensional simulated annealing search of a partially built HLA-A2 model of* 89 α-C *atoms against an e.d. map of the complete protein in space group* $P2_12_12_1$

An exhaustive search to find the top peak at 10σ would have taken some 18 years of MicroVAX II c.p.u. time. Five simulated annealing searches from five different random starting points all converged on the correct/top solution. Each trial took some 4 h of MicroVAX II c.p.u. time and found on the order of 100 good peaks above 6σ.

	Number of peaks above 6σ	Includes correct/top peak?
Random trial 1	65	Yes
Random trial 2	87	Yes
Random trial 3	126	Yes
Random trial 4	96	Yes
Random trial 5	110	Yes

7. Concluding remarks

The application of the simulated annealing procedure to the crystallographic search problem results in producing a list of good solutions in much reduced time scales. In a real application where the exhaustive search can itself be performed it finds a major proportion of all good solutions that may be suitable candidates for the true fit. Moreover, this list includes the best fit. It is not surprising that substantial savings, on the order of 88 000 fold, accompany the application of this protocol. It is a fundamental property of simulated annealing that, as the size of the problem grows, the time required for annealing scales in a manner much less than linear. Perhaps, more important than this sheer improvement in speed, is the possibility of conducting searches that were previously impossible. Since the protocol is completely general, we expect this method to be as effective in reciprocal space. Accordingly, work is under way in assessing the suitability of this method in addressing the molecular replacement problem of protein crystallography. [The program, which is written in Fortran 77, is available in its present form upon request. Please address such requests to SS *via* Bitnet: subbiah@huxtal.]

This work would have been impossible without the help of Dr P. J. Bjorkman, Dr W. Weis, Dr A. Mondragon, Dr M. A. Saper and Professor D. C. Wiley. This work was supported by NSF grant CHE8509574 awarded to S. C. Harrison, D. C. Wiley & M. Karplus.

References

BJORKMAN, P. J., SAPER, M. A., SAMRAOUI, B., BENNETT, W. S., STROMINGER, J. L. & WILEY, D. C. (1987). *Nature (London)*, **329**, 506-512.

BRUNGER, A. T., KURIYAN, J. & KARPLUS, M. (1987). *Science*, **235**, 458-460.
COLMAN, P. M. & WEBSTER, R. G. (1985). *Biological Organisation: Macromolecular Interactions at High Resolution.* P&S Biomedical Sciences Symposia.
HUBER, R. (1965). *Acta Cryst.* **19**, 353-356.
KIRKPATRICK, S., GELATT, C. D. JR & VECCHI, M. P. (1983). *Science*, **220**, 671-220.
METROPOLIS, N., ROSENBLUTH, A., ROSENBLUTH, M., TELLER, A. & TELLER, E. (1953). *J. Chem. Phys.* **21**, 1087-1089.
REYNOLDS, R. A., REMINGTON, S. J., WEAVER, L. H., FISHER, R. G., ANDERSON, W. F., AMMON, H. L. & MATTHEWS, B. W. (1985). *Acta Cryst.* **B41**, 139-147.
VAN HEMMEN, J. L. & MORGENSTERN, I. (1983). Editors. *Heidelberg Colloquium on Spin Glasses. Lecture Notes in Physics*, Vol. 192. Berlin: Springer-Verlag.
VAN KAMPEN, N. G. (1981). In *Stochastic Processes in Physics and Chemistry.* Amsterdam: Elsevier.
WEIS, W., BROWN, J. H., CUSACK, S., PAULSON, J. C., SKEHEL, J. J. & WILEY, D. C. (1988). *Nature (London)*, **333**, 426-431.
WILSON, I. A., SKEHEL, J. J. & WILEY, D. C. (1981). *Nature (London)*, **289**, 366-373.

Acta Cryst. (1989). A**45**, 342-346

Observation of Dependent to Independent Bloch Wave Transition in Kikuchi Patterns

BY A. G. WRIGHT AND D. M. BIRD

School of Physics, University of Bath, Bath BA2 7AY, England

(*Received* 4 *October* 1988; *accepted* 16 *January* 1989)

Abstract

The transition from dependent to independent Bloch waves in high-energy electron diffraction theory is demonstrated by observing the disappearance of subsidiary fringes in Kikuchi patterns as the crystal thickness is increased. Comparison is made between experimental and computed Kikuchi band profiles in Si. It is shown that the subsidiary fringes provide a method for thickness determination in zone axis convergent-beam electron diffraction patterns from relatively thin crystals.

1. Introduction

Electron channelling effects have been observed in a wide range of electron microanalytical techniques including X-ray emission (*e.g.* Cherns, Howie & Jacobs, 1973; Taftø & Spence, 1982), energy loss spectroscopy (*e.g.* Taftø & Krivanek, 1982; Taftø, 1987), backscattering (Hagemann & Reimer, 1979), secondary electron emission (Reimer, Badde, Seidel & Buhring, 1971), cathodoluminescence (Pennycook & Howie, 1980), Compton scattering (Williams & Bourdillon, 1982) and thermal diffuse scattering (Rossouw & Bursill, 1985). The feature common to all these is that the interaction between the fast electrons and the crystal is *localized* about the atomic sites $\mathbf{r}_\kappa$. The magnitude of the interaction therefore depends on the electron density at the atoms $[n(\mathbf{r}_\kappa)]$, and it is the variation of this quantity with orientation that gives rise to the channelling effects. In principle, the electron density must be calculated by taking the modulus squared of the wavefunction, which consists of a sum of Bloch states $\psi^{(j)}(\mathbf{r}_\kappa)$ with amplitudes $\varepsilon^{(j)}$:

$$n(\mathbf{r}_\kappa) = \left| \sum_j \varepsilon^{(j)} \psi^{(j)}(\mathbf{r}_\kappa) \right|^2. \quad (1)$$

$n(\mathbf{r}_\kappa)$ therefore contains cross terms between different Bloch states; in other words, (1) represents a *dependent*-Bloch-wave result. It is well known, however, that in thicker crystals the effects of the cross terms diminish owing to summation over all the atoms in the crystal (*i.e.* thickness integration). The electron density is then well represented by the independent-Bloch-wave result [for a full discussion see, for example, Bird & Wright (1989)]

$$n(\mathbf{r}_\kappa) = \sum_j |\varepsilon^{(j)} \psi^{(j)}(\mathbf{r}_\kappa)|^2. \quad (2)$$

In this paper we show that the subsidiary fringes which are observed in Kikuchi bands are a dependent-Bloch-wave phenomenon, and that their disappearance with increasing thickness provides a simple and clear demonstration of the transition to independent waves. The word 'transition' is a little misleading here, but we use it because it has appeared widely in the literature. In fact, the dependent-Bloch-wave result is always correct; the independent-Bloch-wave treatment simply becomes an increasingly good approximation as thickness increases. We also show that the subsidiary fringes may be used to provide an estimate of thickness in zone axis convergent-beam patterns (CBPs) from relatively thin crystals. The occurrence of these fringes has been known for some time (Uyeda, Fukano & Ichinokawa 1954), together with their theoretical explanation (*e.g.* Fujimoto &

0108-7673/89/050342-05$03.00