Scanning Microscopy Vol. 11, 1997 (Pages 257-276) 0891-7035/97$5.00+0.25
Scanning Microscopy International, Chicago (AMF O'Hare), IL 60666, USA

TRANSFORM NOISE STATISTICS AND FOURIER COMPONENT ESTIMATION

W. O. Saxton

Department of Materials Science and Metallurgy, University of Cambridge, Cambridge, U.K.

Abstract

As a first step towards more absolute quantitative procedures for evaluating images, new and better ways are presented of estimating the spacings and complex amplitudes present in the image of a crystalline specimen by examining the peaks in its calculated transform. So that the expected performance of different estimators can be compared, the statistical properties of the noise in calculated Fourier transforms are established in some detail, and related to those of the noise in the image itself: the variance is found at each pixel, the covariance between pixels, and the actual *distribution* of the transform noise. The role of image windowing in minimising systematic errors due to interference between different Fourier components is made clear, and the properties of three different windows evaluated; the half-cosine window is recommended as a useful compromise between the (trivial) unit window and the von Hann window recommended previously. An alternative approach involving image resampling is shown to have excellent properties for low frequency components, and the degradation of high frequencies arising on re-interpolation is characterised quantitatively.

Key Words: Transform noise, transform statistics, windowing, spectral estimation, lattice spacing measurement, structure factor measurement, interpolation.

*Address for correspondence:
W.O. Saxton
Department of Materials Science and Metallurgy,
University of Cambridge
Pembroke Street, Cambridge
CB2 3QZ, U.K.
Telephone number: +44-1223-334566
E-mail: wos1@cam.ac.uk

Introduction

It is testimony to the potency of visual images that high resolution transmission electron microscope (HR TEM) images have been analysed by purely visual comparison with theoretical images for more than two decades with very little attempt to verify the match quantitatively. It has however become clear gradually over the last decade that there is often a substantial mismatch in absolute contrast levels even when a reasonable visual match is achieved between observed and predicted images (e.g., Hÿtch and Stobbs, 1994); reliable ways of quantifying the degree of match between the two are essential if the reason for the mismatch is to be found.

Independently of this general concern, the possibility of making finer distinctions between structures on the basis of their observed images depends on more accurate ways of estimating parameters such spacings, Fourier component amplitudes, and atomic site intensities from images.

This paper addresses some of the most basic of these questions, dealing particularly with the statistics of noise in calculated image transforms, including the effect of 'windowing'.

The transform noise *distribution* is derived, as this does not appear to be widely familiar; so also are the distributions of its modulus and intensity.

In the light of these statistics, better estimators of both the spacing and the amplitude of the image components are presented, and their performance is evaluated relative to various alternatives. Although this discussion is concerned with *periodic* specimens, partly for simplicity and partly for clarity, some of the findings – and all the results about transform noise statistics – can be extended to the much more important general case.

Finally, the alternative approach of estimating component amplitudes by re-sampling the image to contain whole numbers of unit cells in each direction is examined, and found to be attractive in many respects.

The Problem of Fourier Component Estimation

The image of a crystalline specimen may be reduced to a few numbers only, specifying the spatial frequency and complex amplitude of its components. The frequencies should of course form a lattice, the reciprocal lattice, and are estimated essentially from the *positions* of the peaks in the calculated image transform, while the

amplitudes are estimated from the *value* of the transform at the peak positions. Neither is however trivial to determine *accurately* from an experimental image: if no particular precautions are taken, the calculated transform normally exhibits marked horizontal and vertical streaking around peaks, confusing both positions and values; the phenomenon is obvious in the top row of Figure 1.

The difficulty arises from the fact that the field of view from which the transform is calculated does not normally hold a whole number of unit cells. A component with (h,k) cycles across the image field in the two directions gives rise to a transform peak (h,k) pixels from the origin, and these numbers are not normally integers. The effect is seen clearly in the transform of a field containing a single component, which can be calculated analytically.

An image array

$$f_{pq} = c\exp\{2\pi i(\frac{h_o p}{M} + \frac{k_o q}{N})\} \qquad (1)$$

i.e. comprising a single component with (complex) amplitude c and frequency (h_o,k_o) cycles per field, has a discrete Fourier transform (DFT) that can be evaluated easily and is closely approximated by

$$F_{hk} = MNc\frac{\sin\pi(h-h_o)}{\pi(h-h_o)}\frac{\sin\pi(k-k_o)}{\pi(k-k_o)} \qquad (2)$$

provided only that M,N are much greater than 1. This has a profile that is the product of two 1-D factors of the form

$$M\frac{\sin(\pi h)}{\pi h} \qquad (3)$$

centred at h_o,k_o – the transform in fact of the rectangle function bounding the image field. [See Appendix 1 for a definition of continuous and discrete Fourier transforms, as used here, and the relationship between these.]

The transform values thus sample 'sinc' functions in both directions; the effect of the sampling is shown clearly in the top row of Figure 2. When h_o and k_o are integers, the samples are all zero except for being one at the position (h_o,k_o) itself - the expected isolated peak or 'delta function'. Generally however, the peak region exhibits a cross with arms along the h and k directions, oscillating and decaying no faster than inversely with distance from the peak position. In an image with many Fourier components, one peak may be significantly distorted by the overlapping tails of neighbouring peaks (De Ruijter, 1994).

Table 1: Window transform samples W/M in 1-D

	W_0	$W_{0.25}$	$W_{0.5}$	$W_{0.75}$	W_1	W_2	W_3
Unit	1	0.900	0.637	0.300	0	0	0
½cosine	1	0.943	0.785	0.566	0.333	-0.067	0.029
v.Hann	1	0.960	0.849	0.686	0.5	0	0

Image Window Functions and Their Transforms

The transform peak profile is greatly improved, and the overlap accordingly diminished, if the image is multiplied before transformation by a suitable *window* function that decays slowly to zero at the outside of the field: the peak profile is smoothed, and the oscillation diminished, by convolution with the window transform. While many different forms of window functions have been used, we will consider only three in detail. For each, we set out the window function itself w_p, its DFT W_h (the transform peak profile[1], which determines the signal level in DFT pixels), and the DFT U_h of t.

The squared window (which we shall see below determines the noise level in DFT pixels, and the correlation between these). As two-dimensional (2-D) windows – the case of practical interest – all are *separable* as products of 1-D functions so that $w_{pq} = w_p w_q$, $W_{hk} = W_h W_k$. and $U_{hk} = U_h U_k$.

The first is the trivial case of the *unit* window effectively considered in the previous section

$$\begin{aligned} w_p &= 1 \\ W_h &= M\frac{\sin(\pi h)}{\pi h} \\ U_h &= M\frac{\sin(\pi h)}{\pi h} \end{aligned} \qquad (4)$$

The second is the *half-cosine* window:

$$\begin{aligned} w_p &= \frac{\pi}{2}\cos(\pi\frac{p}{M}) \\ W_h &= M\frac{\cos(\pi h)}{1-4h^2} \\ U_h &= M\frac{\pi\sin(\pi h)}{8h(1-h^2)} \end{aligned} \qquad (5)$$

[1] The transform is given *approximately* only in each case; for example, the exact transform for the unit window case is $\exp(\frac{\pi i h}{M})\frac{\sin\pi h}{\sin(\pi h/M)}$. The approximation error may need considering if parameters are estimated from a small region only of an image, so that M is not very large.

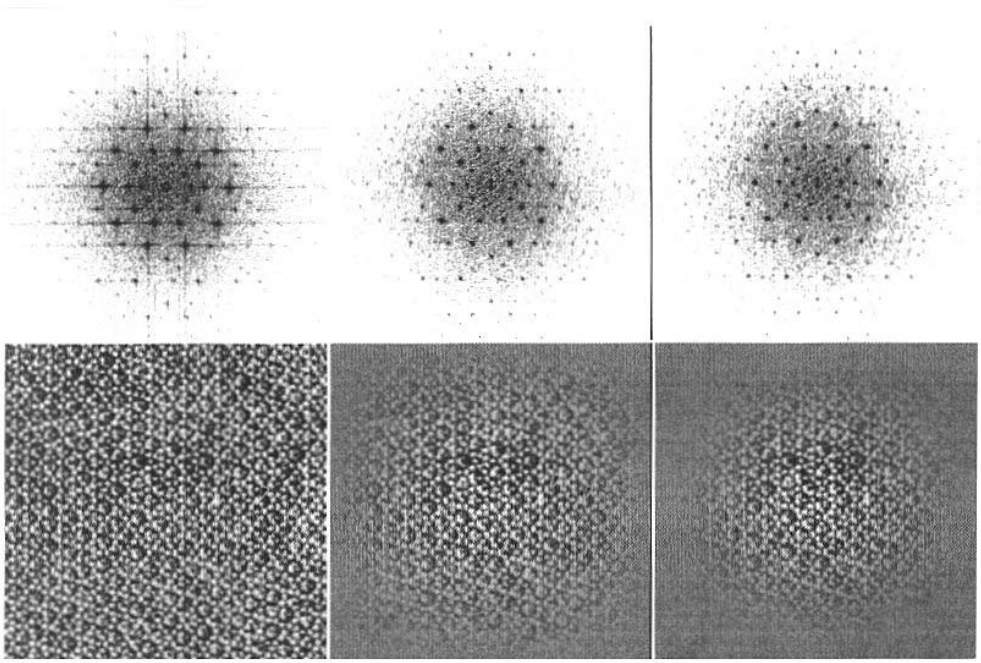

Figure 1. Image of an Al-Mn quasicrystal (left), and calculated transform intensity logarithm (right), with unit, half-cosine and von Hann windows applied, from top to bottom.

Finally, De Ruijter (1994) has recommended the *von Hann* or *raised cosine* window[2].

$$w_p = 1 + \cos(2\pi \tfrac{p}{M})$$
$$W_h = M \frac{\sin(\pi h)}{\pi h(1-h^2)} \qquad (6)$$
$$U_h = M \frac{6\sin(\pi h)}{\pi h(1-h^2)(4-h^2)}$$

The effect of all three windows on an image and the calculated transform is illustrated in Figure 1, while Figure 2 shows the three peak profiles, and the effect of sampling them. Figure 3 shows the three windows w_p and their transforms W_h as line graphs.

[2] This is named after Julius von Hann, to whose original work I have not unfortunately been able to find a reference. The widespread designation 'hanning' window is unfortunate in encouraging confusion with the subtly different 'hamming' window $1.08 + 0.92\cos(2\pi \tfrac{p}{M})$, named after R W Hamming (1977), which achieves a lower first sidelobe in the transform at the expense of a slower decay at large distances.

Table 2: Squared window transform samples U/M in 1-D

Window	U_0	U_1	U_2	U_3
Unit	1	0	0	0
½cosine	1.234	0.617	0	0
v.Hann	1.5	1	0.25	0

The peak profile W has a central value M in each case; the distance from the centre to the first zero is 1, 1.5 and 2 pixels respectively; and the profile decays inversely as the first, second, and third power of the distance respectively. The last is clearly the most effective in suppressing overlaps between transform peaks; however we shall see below that the others are preferable in other respects, and the half-cosine window may indeed be the most useful generally.

The 1-D profiles are tabulated for useful values of h in Table 1; only the half-cosine window has non-vanishing samples outside the central maximum. The 2-D profiles are obtained by multiplying together 1-D profiles in each direction, e.g., $W_{hk} = W_h W_k$. The squared window transforms are tabulated at similar integer values of h in Table 2; the samples vanish outside the "central" maximum in all three cases.

While they are not examined further below, it is probably useful here to note two other common but non-

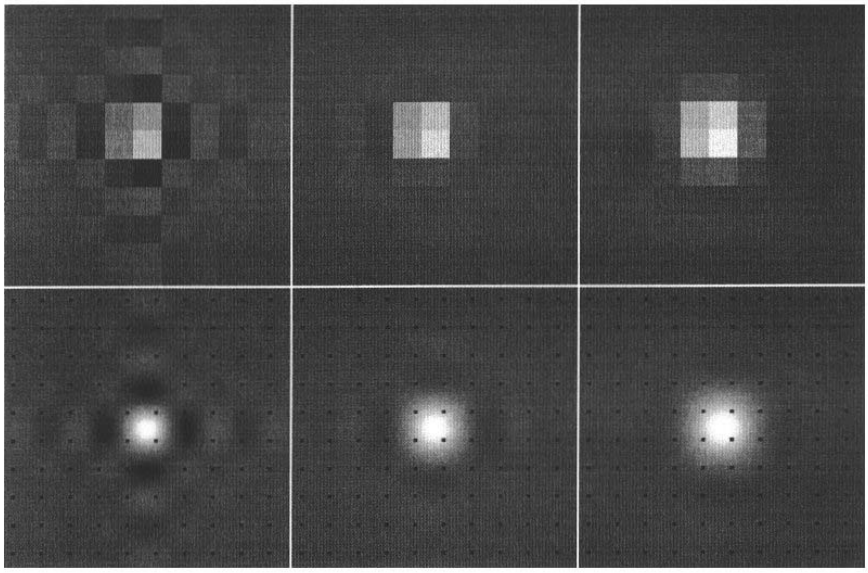

Figure 2. Theoretical peak profiles for unit, half-cosine and von Hann windows, from top to bottom; (left) in continuous form, and (right) as sampled at the positions marked left.

separable window functions. A simple circular mask (unity within, and zero outside, a ccntral circle of radius p_o pixels), gives a transform peak with a profile of the form $2MJ_1(\pi h)/(\pi h)$, decaying inversely as the three-halves power of the distance, not much better than the unit window. However, the same mask extended by a gaussian edge is virtually the same as the mask convolved with a gaussian, so that its transform peak profile, being *multiplied* by a gaussian, decays very rapidly indeed.

A note is necessary on the treatment of the background level of the image (its spatial mean), which may need subtracting before windowing is applied. Without windowing, this level affects only the central pixel of the transform; windowing without background subtraction -- the simplest expedient -- causes that pixel to be replaced by the appropriate peak profile however; as it is commonly orders of magnitude higher than the other transform peaks, significant overlap can arise in spite of a rapid profile decay. Subtracting the background level before windowing eliminates the central peak completely; if desired (e.g., to avoid negative pixels), the original image data range can be maintained by adding the background again after windowing, which affects the central pixel only.

Transform Signal and Noise Statistics

We now present systematically a number of results about the signal and noise levels in the transform of a noisy image with the three different windows applied; these are used subsequently to establish the expected standard deviation (SD) or its square (the variance) of various parameter estimators, and are in any case essential ground-work for later investigations to be reported elsewhere. Appendix 2 explains the generalisation of statistical parameters such as variance to complex variables in general, Appendix 3 derives the expectation (signal level), variance (squared noise level), and covariance (interdependence) of windowed image transform pixels; the results are summarised here. The signal level in transform pixels near a peak depends critically, as noted above, on the exact distance $(u,v) = (h-h_0,k-k_0)$, from pixel to peak. For a Fourier component with complex) amplitude c, the pixel expectation is

$$\overline{F}_{hk} = cW(u,v) \qquad (7)$$

and is in general *increased* by windowing. At one extreme, when $u = v = 0$, $W(u,v)$ is MN for all three windows; at the other extreme, when $u = v = 0.5$, $W(u,v)$ is $0.405MN$, $0.617MN$ and $0.721MN$ for the three cases.

If the image pixels all have a standard deviation (SD) σ, and the noise in different pixels is uncorrelated, then the noise in all transform pixels has the same variance

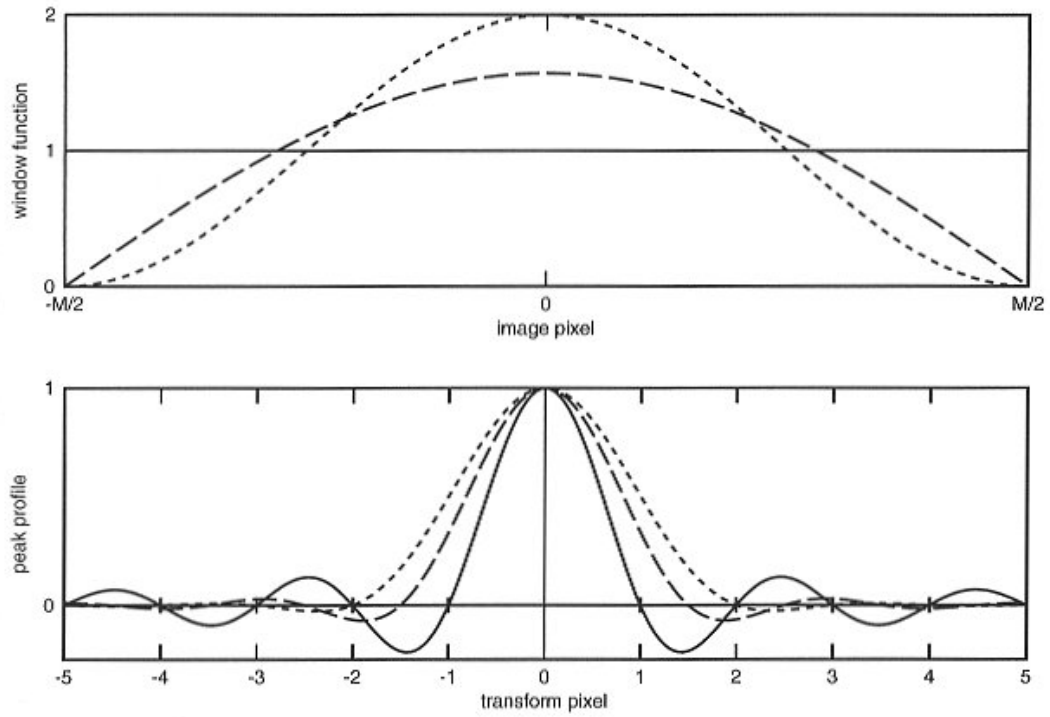

Figure 3. 1-D profiles through the unit, half-cosine and von Hann windows (above) and their transforms (below), shown solid, broken, and dashed respectively.

$$\mathrm{var}\{u\} = \sigma^2 U_{00} \tag{8}$$

where U_{hk} is the DFT of the squared window function w_{pq}^2, as given in Equations (4-6) and Table 2. This is also increased by windowing; U_{00} is MN, $1.52MN$ and $2.25MN$ respectively for the unit, half-cosine and von Hann windows.

When (u,v) are small, the signal-to-noise (s/n) ratio in individual transform pixels, measured as the ratio of the signal to the noise SD, thus deteriorates by a factor $1/\sqrt{1.522} = 0.811$ on application of the half-cosine window, and by $1/\sqrt{2.25} = 0.667$ with the von Hann window; however at the other extreme, when $u = v = 0.5$, there is a slight improvement on windowing: the s/n ratio changes by $0.617/(0.405\sqrt{1.522}) = 1.235$ on application of the half-cosine window and by $0.721/(.405\sqrt{2.25}) = 1.170$ with the von Hann window.

If the image pixels all have a SD σ, and the noise in different pixels is uncorrelated, the covariance between two transform pixels separated by (h,k) is

$$\sigma^2 U_{hk} \tag{9}$$

with U_{hk} obtainable from Table 2 again, for nearest that although neighbouring pixels are uncorrelated for neighbour and for diagonal neighbour pixels. We note the unit window, nearest neighbours are correlated for the half-cosine window, and neighbours up to two pixels away are correlated for the von Hann window.

The way in which windowing introduces correlation between neighbouring transform pixels may be understood simply. The windows are simple superpositions of slowly varying linear phase factors, multiplication by any one of which results in a small displacement of the transform; multiplication by the superposition of the phase factors thus results in a superposition of mutually displaced transforms, so that each pixel involves a superposition of its original close neighbours.

The *real* and *imaginary parts* of a transform pixel each have a variance half that in the complex value (Equation 8), and are uncorrelated with each other (so that their variances simply add to give that in the complex value). For two transform pixels separated by (h,k), the covariance between the two real parts, and also between the two imaginary parts, is half that given in Equation 9,

while the real part of one is also uncorrelated with the imaginary part of the other.

The *modulus* of a transform pixel has a variance once again half that in the complex value (Equation 8), and the *phase* has the same variance too apart from division by the signal intensity; the modulus and the phase are not correlated. For two transform pixels separated by (h,k) the covariance between the moduli is a factor β of half of that in Equation 9, with β depending on the relative phases of the two pixel signals, being the cosine of the phase difference.

Transform Noise Distributions

The previous section has given simple expressions for the common statistical parameters measuring the noise in DFT pixels, independent of the particular *distribution* of the noise in image or transform. This section notes that the transform noise has a gaussian distribution, with independent real and imaginary parts, regardless of the image noise distribution. It also sets out the distributions of the modulus and the intensity of the noise transform; and the distributions of the modulus and phase of a transform pixel comprising both signal and noise.

Appendix 4 shows that if an image f_{pq} has uncorrelated pixels, with zero expectation and variance σ^2 everywhere, the probability distribution of the real and imaginary parts $G_{hk}+iH_{hk}$ of its DFT F_{hk} are independent, with a joint gaussian distribution

$$p(G,H) = \frac{1}{\sigma^2 \pi U_{00}} \exp\{-\frac{G^2+H^2}{\sigma^2 U_{00}}\} \qquad (10)$$

which is illustrated in Figure 4. As noted in the previous section, each part has a variance

$$\mathrm{var}\{G\} = \mathrm{var}\{H\} = \tfrac{1}{2}\sigma^2 U_{00} \qquad (11)$$

To find the distribution of $|F|$, we integrate $p(G,H)$ over annular elements at a given |F|; this gives

$$p(|F|) = \frac{2}{\sigma^2 U_{00}} |F| \exp\{-\frac{|F|^2}{\sigma^2 U_{00}}\} \qquad (12)$$

which is illustrated in Figure 5; and has an expectation and variance

$$E\{|F|\} = \tfrac{1}{2}\sigma\sqrt{\pi U_{00}}\,; \quad \mathrm{var}\{|F|\} = (1-\tfrac{\pi}{4})\sigma^2 U_{00}\,. \qquad (13)$$

The distribution of the intensity $I = |F|^2$ is obtained from this via

$$p(|F|)d|F| = p(I)dI \qquad (14)$$

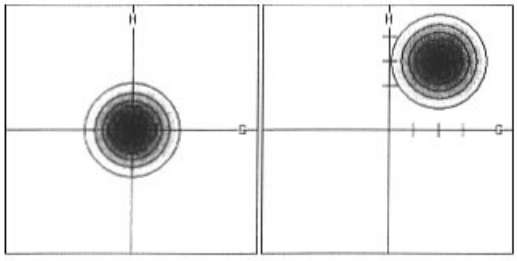

Figure 4. Probability distribution of real and imaginary parts of calculated transform pixel $F = G+iH$. Left: with zero expectation (i.e. transform of pure noise); right: with non-zero expectation.

which gives a negative exponential distribution for the noise intensity

$$p(I) = \frac{1}{\sigma^2 U_{00}} \exp\{-\frac{I}{\sigma^2 U_{00}}\} \qquad (15)$$

also illustrated in Figure 5; this has an expectation and variance

$$E\{I\} = \sigma^2 U_{00}\,; \quad \mathrm{var}\{I\} = \sigma^4 U_{00}^2 \qquad (16)$$

If the image does *not* have a zero expectation, the only change in $p(G,H)$ is displacement to the expected position $(\overline{G},\overline{H})$. Provided the noise is small compared with the signal, Appendix 4 shows that the modulus of the pixel |F| also has the same gaussian distribution

$$p(|F|) = \frac{1}{\sigma\sqrt{\pi U_{00}}} \exp\{-\frac{(|F|-|\overline{F}|)^2}{\sigma^2 U_{00}}\} \qquad (17)$$

with a variance

$$\mathrm{var}\{|F|\} = \tfrac{1}{2}\sigma^2 U_{00}$$

The argument (phase) θ of F is also similarly distributed apart from a scaling factor:

$$p(\theta) = \frac{|\overline{F}|}{\sigma\sqrt{\pi U_{00}}} \exp\{-\frac{|\overline{F}|^2 (\theta-\overline{\theta})^2}{\sigma^2 U_{00}}\} \qquad (18)$$

with a variance

$$\mathrm{var}\{\theta\} = \tfrac{1}{2}\frac{\sigma^2 U_{00}}{|\overline{F}|^2} \qquad (19)$$

Spatial Frequency (Spacing) Estimation

Essentially, the spatial frequency of a given Fourier component is of course estimated by the position of the peak; the transform sampling is however often rather coarse, and it is important to estimate this position with sub-pixel accuracy. De Ruijter (1994) has pointed out that

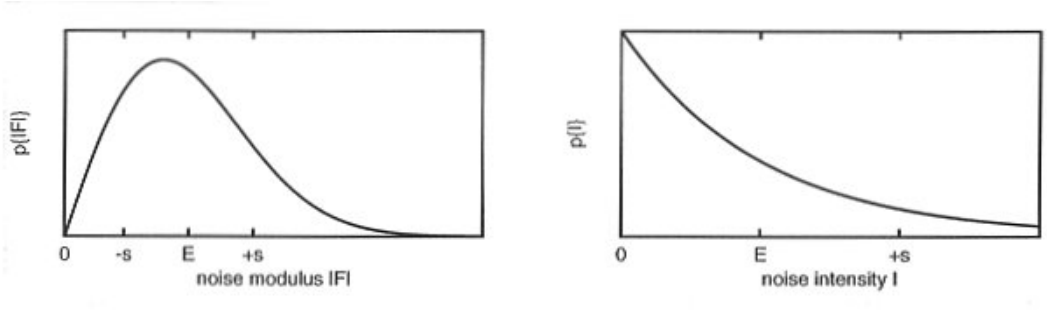

Figure 5. Probability distributions of modulus and intensity (left and right) of noise in calculated transform pixel.

this can be estimated conveniently from the relative values of pixels adjacent to the highest; we give here the estimators appropriate to each of the three windows above, and improve their accuracy by using more of the transform data.

Consider the 2×2 block of pixels around the ideal transform peak position, found by locating the largest modulus pixel[3], and including the larger of the two neighbours horizontally and vertically; for brevity subsequently we call the pixels p_{00}, p_{10}, p_{01} and p_{11}. We seek estimators for the fractional distances (u,v) across this block from p_{00} to the ideal peak position; in fact we only consider u explicitly, as v is equivalent with rows and columns interchanged.

The pixels have expectations

$$\begin{aligned} p_{00} &= F_{hk} = cW(u)W(v); \\ p_{10} &= cW(1-u)W(v); \\ p_{01} &= cW(u)W(1-v); \\ p_{11} &= cW(1-u)W(1-v). \end{aligned} \tag{20}$$

The simplest estimators for u relies on p_{00} and p_{10}. For the unit window, these have the expected form

$$\begin{aligned} p_{00} &= cMN \frac{\sin \pi u}{\pi u} \frac{\sin \pi v}{\pi v} \\ p_{10} &= cMN \frac{\sin \pi(1-u)}{\pi(1-u)} \frac{\sin \pi v}{\pi v} \end{aligned} \tag{21}$$

since $\sin \pi(1-u) = \sin \pi u$, u is easily extracted from the ratio p_{00}/p_{10} to give the estimator

$$u' = \frac{|p_{10}|}{|p_{10}|+|p_{00}|} \tag{22}$$

in which the modulus of the pixels is used in preference to the complex values themselves to ensure the estimate is real[4].

In the same way, we can find estimators for the half-cosine and von Hann windows respectively (the last being De Ruijter's recommendation):

$$u' = \frac{3|p_{10}|-|p_{00}|}{2(|p_{10}|+|p_{00}|)} \tag{23}$$

$$u' = \frac{2|p_{10}|-|p_{00}|}{|p_{10}|+|p_{00}|} \tag{24}$$

These estimators rely on one row only of the 2×2 block; however, the ratio $|p_{01}/p_{11}|$ between pixels in the other row is expected to be the same as $|p_{00}/p_{10}|$, and we expect estimators based on the *average* of the two rows to be more accurate. Accordingly, the estimators we now examine in detail are obtained from *weighted* averages of the two rows, with weighting depending on the true peak position.

If the original and the second row are given weights x and $(1-x)$, the estimators for the three windows become

$$u' = \frac{x|p_{10}|+(1-x)|p_{11}|}{x|p_{10}|+(1-x)|p_{11}|+x|p_{00}|+(1-x)|p_{01}|} \tag{25}$$

$$u' = \frac{3(x|p_{10}|+(1-x)|p_{11}|-(x|p_{00}|+(1-x)|p_{01}|)}{2(x|p_{10}|+(1-x)|p_{11}|+x|p_{00}|+(1-x)|p_{01}|)} \tag{26}$$

$$u' = \frac{2(x|p_{10}|+(1-x)|p_{11}|)-(x|p_{00}|+(1-x)|p_{01}|)}{x|p_{10}|+(1-x)|p_{11}|+x|p_{00}|+(1-x)|p_{01}|} \tag{27}$$

[3] When the transform signal-to-noise ratio is poor, it is of course possible that this procedure does not correctly identify the pixel nearest h_0,k_0; the expressions given subsequently for the accuracy must thus be considered optimistic.

[4] The variance of $|p_{10}+p_{00}|$ is in fact the same as that of $|p_{10}|+|p_{00}|$, so it does not matter which form is used.

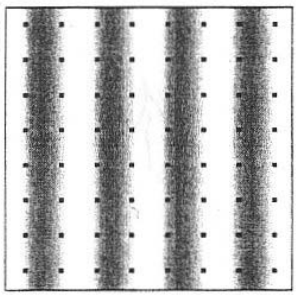

Figure 6. An image fringe system parallel to the sampling lattice, and sampled near its zeros.

The choice $x = 1$ is equivalent to the two-point estimators in (22-24) above; $x = 0.5$ gives all four points equal weight.

The SD in these estimators is not simple to obtain unfortunately (because of the interdependence of neighbouring transform pixels), and is derived in Appendix 5. The four-pixel estimator ($x = 0.5$) is found to have about half the variance of the two-pixel estimator ($x = 1$) when v is 0.5, i.e. when the peak lies mid-way between the two rows, but the relative performance is reversed when v is 0 and the peak lies in the first row. This is essentially because when the peak lies near the first row, the data in the second row are small and contribute more noise than signal to an equally weighted estimator.

The *optimum* weighting of the two rows would ideally be found by minimising the resulting SD with respect to x, for each combination of (u,v) values. However, a little manual exploration shows that weighting

$$x = 1 - v; \qquad x = 1 - 2v^2; \qquad x = 1 - 2v^2 \tag{28}$$

is certainly not far from optimal for the three cases, and this approximation is proposed accordingly; the resulting variance, tabulated in appendix 5 for a range of (u,v) values, is at least as good everywhere as the *better* of the two limiting cases $x = 1$ and $x = 0.5$. The actual value of v must be estimated by a preliminary calculation using a fixed weighting such as $x = 1$.

A reasonable summary of the variance expected in the peak position (Fourier component spacing) is now possible as follows. The SD in the each component of the estimated spacing (in cycles per field) is

$$SD\{u'\} = \frac{\alpha}{\sqrt{MN}} \frac{\sigma}{|c|} \tag{29}$$

with the multiplier a having a value around 0.5 (actually varying from 0.15 to 1.4 depending on the value of (u,v) and the window function used). Applying the half-cosine and von Hann windows increases the SD by factors of around 1.3 and 1.8 respectively; this is the cost of eliminating the systematic error arising when one peak is overlapped by the tails of another.

Fourier Component Amplitude Estimation

Essentially, the complex amplitude of an image Fourier component is estimated from the value of the transform pixel p_{00} nearest the estimated peak position h_0,k_0, . This section examines several particular estimators, establishing the variance of each.

The transform pixels have the expectations (signal) given in Equation 7 above, and variance (noise power) given in Equation 8; the correlation between pixels, when image windowing is employed, is given in Equation 9.

$$F_{hk} = cW_{hk} + N_{hk} \tag{30}$$

The simplest estimator considered is based on the single pixel p_{00}, with the peak profile divided out:

$$c' = \frac{p_{00}}{W_{hk}} \tag{31}$$

This clearly has an expectation c, and so is unbiassed; its variance is simply

$$\mathrm{var}\{c'\} = \frac{\sigma^2 U_{00}}{W_{hk}^2} \tag{32}$$

While this is always about $\sigma^2/(MN)$, the actual multiplier depends on the actual values of (u,v) as well as on the window function used; tables A6.1-3 give some representative values for the three windows, with multipliers between 1 and 6. In all cases, the accuracy is best for $(u,v) = 0,0$, where W_{hk} takes its maximum value of MN, and worst for $(u,v) = ½,½$ – the opposite pattern to that for the spacing estimators; the unit window is the most variable, being the best of the three for (u,v) near (0,0), but easily the worst for (u,v) near (½,½).

Estimators based on more than one pixel may be expected to be more accurate. Accordingly, we consider next the estimator achieving a least-squares fit to several pixels around the peak (though we shall see below this is not in fact optimal):

$$c' = \frac{\Sigma_{ij} W_{ij}^* p_{ij}}{V}, \text{ with } V = \Sigma_{ij} |W_{ij}|^2 \tag{33}$$

specifically, this achieves a minimum summed squared difference between observed and predicted pixel values

$$\Sigma_{ij} |c' W_{ij} - p_{ij}|^2 \tag{34}$$

This estimator is also easily seen to have an unbiassed expectation c. Its variance is not easily calculated however because of the correlation between neighbouring

pixels. Appendix 6 calculates the variance for a four-point estimator in which the sum extends over the four pixels p_{00}, p_{10}, p_{01} and p_{11}.

Tables A6.4-6 give some representative values of the variance for the three windows. The result depends again on the actual values of (u,v) as well as on the window function used; however, the dependence is not strong, and it is a reasonable summary to say that the variance is close to σ^2/MN, $2\,\sigma^2/MN$ and $3\sigma^2/MN$ for the unit, half-cosine and von Hann windows respectively.

Thirdly, if the sum in (33) is extended over *all* transform pixels, it is possible to calculate the variance of the resulting estimator by a different method, also given in appendix 6, and it proves to be independent of the values of (u,v):

$$\mathrm{var}\{c'\} = \sigma^2 \frac{\Sigma_{p,q} w_{pq}^4}{(\Sigma_{p,q} w_{pq}^2)^2} \qquad (35)$$

The value of this no longer depends on the actual peak position (u,v); for the unit, half-cosine and von Hann windows respectively, its value is σ^2/MN, $2.25\sigma^2/MN$ and $3.78\sigma^2/MN$, i.e. slightly greater than that of the four-point estimator.

For practical purposes, such an estimator would of course need to be approximated by one in which the sum was truncated so as to exclude other transform peaks (which might introduce bias); the point of practical interest is that it does not matter much exactly how large the area is over which the sum in (33) is extended.

Fourthly, we set out the *optimal* estimate possible from a given set of transform pixels. When the transform data are correlated, as here, a least-squares fit is not in fact the best possible, in the sense that a different choice of the coefficients in (33) achieves a lower variance in c'. A statement of the best estimator demands a more formal vector/matrix notation: given a set of pixels p_i near a transform peak, with expectations cW_i, and a variance-covariance matrix $c_{ij} = \mathrm{cov}(p_i,p_j)$, the optimal estimator of c is

$$c' = \Sigma_i r_i p_i \qquad (36)$$

with a coefficient vector given by

$$r_i = \frac{c_{ik}^{-1} W_k^*}{W_i c_{ik}^{-1} W_k^*} = \frac{\mathbf{C}^{-1}\mathbf{W}*}{\mathbf{W}.\mathbf{C}^{-1}\mathbf{W}*} \qquad (37)$$

and a variance

$$\mathrm{var}\{c'\} = \frac{1}{\mathbf{W}.\mathbf{C}^{-1}\mathbf{W}^*} \qquad (38)$$

Appendix 7 proves these results, illustrating them for a four-pixel estimator. For the unit window, the results are exactly the same as for the least-square estimator (Equation 33); however tables A7.1 and A7.2 list representative values of the variance for the half-cosine and von Hann windows, which can be seen to be up to a third lower than the variance of the corresponding least-squares estimator.

Finally, we consider a paradox that may underline the significance of the results set out above. It is a familiar fact that the samples values obtained from an image fringe system parallel to the sampling lattice and with a period close to two pixels sometimes reflect the full fringe amplitude (when the samples fall near the fringe extrema) but are also sometimes close to zero (when they fall mid-way between these). Over large distances, the relative phase of fringes and sample positions change, ensuring the fringes are detected; but within a given image field they may hardly be registered at all, as in Figure 6. How is this to be reconciled with the statements in this section that the fringe amplitude can be estimated without bias regardless of the spacing and the size of the field of view?

The explanation of the paradox lies in the possibility of overlap between different transform peaks. A peak near the limit of either transform axis will be accompanied by a conjugate peak at the opposite limit; the periodic nature of DFTs means that each is overlaid by a repeat of the other just outside the field of view. Depending on their relative phase (i.e. on the relative position of fringes and samples) they may interfere constructively or destructively. No such problem arises if the fringe system is not parallel to the sampling lattice as the transform peak is not then near the repeat of its conjugate.

Refinement of Lattice Spacings

It will frequently, though not invariably, be possible to measure the positions of many independent transform peaks, and to refine the reciprocal lattice base vectors deduced from any pair by a least-squares fit to all positions. The details of the minimisation are given, for example, by Saxton (1992); initial estimates for the base vectors are needed sufficiently accurate to *index* positions correctly. Clearly, the highest order peaks will define the base vectors most closely; a very rough estimate, based on the usual pattern of error reduction by a factor of $\sqrt{n}$ when n independent values are averaged, is that a total of n high order peak positions measured, with an error σ_h in each component, results in a standard error around

$$\sigma_h\sqrt{(2/n)} \qquad (39)$$

in the fitted base vector components.

It is of course also possible to determine (real-space) lattice vectors directly in real space, by least-squares fitting the peak positions in the auto-correlation function of the image, or its cross-correlation function with a smaller subregion (Saxton and Baumeister, 1982). While

the real-space approach appears likely to be more accurate when the unit cell is large so that the transform peaks are all near the origin and therefore coarsely sampled, the question has still not been examined carefully, and remains open.

Apart from its possible use in identifying small included phases from their lattice spacings, the main reason for seeking high accuracy in measured spacings is only their effect on the component amplitude estimates; this is considered in the next section.

Amplitude Estimation in Real Space

It is easy to estimate Fourier component amplitudes (though not spacings) directly from the image, and their are some advantages in doing so, which we will note in this section. The principal drawback is that, although simple, the amplitude estimators require more computation than their Fourier-space counterparts.

As elsewhere, we ignore initially all but a single component of the image, with h_o,k_o cycles across the field:

$$f_{pq} = c\exp\{2\pi i(\frac{h_o p}{M}+\frac{k_o q}{N})\} + n_{pq} = \bar{f}_{pq} + n_{pq} \quad (40)$$

in which the first term is the image signal (expectation) $\bar{f}_{pq}$, and the second term the image noise n_{pq}, with zero mean and variance σ^2_{pq}. We examine the amplitude estimator

$$c' = \frac{1}{MN}\Sigma_{p,q} w_{pq} f_{pq} \exp\{-2\pi i(\frac{h_o p}{M}+\frac{k_o q}{N})\} \quad (41)$$

which involves a summation over the entire image; w_{pq} is a window functions such as those discussed earlier (and introduced with the same objective of reducing systematic bias arising from other components with similar frequencies, as becomes clear below). Firstly, the estimator is unbiassed, since according to (A2.5), the expectation of the estimator (41) is

$$\begin{aligned} E\{c'\} &= \frac{1}{MN}\Sigma_{p,q} w_{pq} \bar{f}_{pq} \exp\{-2\pi i(\frac{h_o p}{M}+\frac{k_o q}{N})\} \\ &= \frac{1}{MN}\Sigma_{p,q} w_{pq} . c \quad = c \end{aligned} \quad (42)$$

Secondly, according to (A2.7), its variance, for uncorrelated pixels with variance σ^2 everywhere, is

$$\mathrm{var}\{c'\} = \frac{1}{M^2N^2}\Sigma_{p,q} w^2_{pq} . \sigma^2 = \frac{\sigma^2 U_{00}}{M^2N^2} \quad (43)$$

being σ^2/MN, $1.52\,\sigma^2/MN$ and $2.25\,\sigma^2/MN$ for the three windows, and as good as the best of the Fourier-space estimators in each case, no matter where the transform peak lies relative to the DFT pixels; this is one of the virtues of the real-space estimator.

The cross-talk between the signal in different Fourier components – the reason for introducing image window functions above – manifests itself to exactly the same degree in real space. If the image contains a further component with amplitude c_{hk} and with (h,k) cycles across the field, then the expectation of Equation 41 contains an additional term

$$\begin{aligned} &\frac{c}{MN}\Sigma_{p,q} w_{pq} \exp\{2\pi i(\frac{(h-h_o)p}{M}+\frac{(k-k_o)q}{N})\} \\ &\qquad = \frac{c}{MN} W(h-h_o, k-k_o) \end{aligned} \quad (44)$$

equivalent to the contribution to the DFT at h_o,k_o from another transform peak at (h,k). Window functions thus have exactly the same role in reducing cross-talk between components, as they do in Fourier-space amplitude estimation.

The other virtue apparent in the real-space amplitude estimator (Equation 41) is a lower sensitivity to systematic error arising from a mis-estimated spacing. If the estimator (Equation 41) is calculated with an incorrect value (h,k) for the spacing, its expectation is the expression in Equation 44 rather than c, the result is Equation 44 rather than Equation 42. Since the function W passes through a maximum around h_o,k_o however, the errors caused in the amplitude estimate are very small (second error small for first order errors in the spacing). The same is not true of the Fourier-space estimates: although the effect is less simply summarised, the fact that in Fourier space W is usually sampled at points other than its maximum, where it is changing comparatively rapidly, means that errors in the estimated spacing cause larger (first order) errors in the amplitude estimate. In either case, windowing reduces the sensitivity of the amplitude estimate to errors in the spacing estimate, by making W vary more slowly.

Amplitude Estimation by Image Resampling

A final option to be considered is the resampling of the image on a (non-cartesian) lattice with base vectors parallel to those of the crystal lattice, and a whole number of unit cells contained within the field in both directions (e.g., Aebi *et al*., 1973).

The distinctive advantage of this approach is the complete elimination of cross-talk between peaks and direct amplitude estimation without windowing or transform peak profile fitting. Once the image has been resampled in this way, all the components present have an integral number of cycles in each direction across the field; this results in transform peaks restricted to single points only, with no possibility of overlap by other

components (cf. the section on the Problem of Fourier Component Estimation). The component amplitudes are obtainable directly from the DFT pixels: the amplitude of the component with (h,k) cycles across the (interpolated) field is simply

$$c = \frac{F_{hk}}{MN} \tag{45}$$

which has a variance σ^2/MN under the assumptions made elsewhere about the image noise.

Computationally, this is clearly the most efficient way of determining the Fourier component amplitudes. There are other less important virtues in the approach also: unit cells may be extracted individually and averaged before transformation, or transformed individually and averaged in Fourier space; in either case, the (h,k) pixel of the resulting transform immediately gives the amplitude of the component with (h,k) cycles across the unit cell. This local approach (e.g., Saxton and Baumeister, 1982) makes it possible to exploit irregularly shaped regions of crystal that do not fill a rectangular field (or a parallelogram before interpolation) efficiently. If the squared unit cells are averaged too, the noise level can be measured directly in real and Fourier space, via Equation A2.2. Moreover, very modest resources are sufficient, as only small arrays need to be transformed.

These benefits are offset by one serious drawback however: most forms of interpolation *smooth* the image, and so reduce the amplitude of high frequency components. Loosely, this means that image components with periods shorter than four pixels may be seriously underestimated.

The effect of interpolation defies precise description, as it depends on the relative positions of original and final samples, which vary in an irregular way across the field. However, its effect can be roughly modelled, for any particular interpolation method, by assuming an average over a uniform random distribution of relative placements between original and final samples; we illustrate this for two simple forms of interpolation in 1-D.

The value obtained by bilinear interpolation between two given pixels f_p and f_{p+1} at a fractional distance x from the first is

$$f'_{p+x} = (1-x)f_p + xf_{p+1} \tag{46}$$

Accordingly, a sample value of a continuous function $f(p)$ obtained by bilinear interpolation from two samples x' below it and $1-x'$ beyond it is

$$f'(p) = (1-x')f(p-x') + x'f(p+1-x') \tag{47}$$

We now model the effect of repeating such an interpolation for many equivalent values of p (i.e., positions connected by the image periodicity) by averaging over a uniform probability distribution for x' from 0 to 1:

$$\begin{aligned}&\int_0^1 (1-x')f(p-x')dx' + \int_0^1 x'f(p+1-x')dx' \\ &= \int_0^1 (1-x')f(p-x')dx' + \int_{-1}^0 (1+x')f(p-x')dx'\end{aligned} \tag{48}$$

This is a convolution with a triangle function extending from -1 to 1, with maximum value 1; its effect is accordingly to multiply the Fourier transform by the transform, i.e.

$$M^2 \frac{\sin^2(\pi \frac{h}{M})}{\pi^2 h^2} \tag{49}$$

The effect on the highest spatial frequency present, $h = M/2$, corresponding to a two-pixel period, is thus attenuation by a factor of 0.41 (and the square of this in 2-D, i.e. 0.16); at half this maximum frequency (i.e. a four-pixel period, the attenuation is by a factor of 0.81 (0.66 in 2-D).

Curiously, nearest-neighbour interpolation, in which the value at the required position is simply replaced by the nearest available sample, while producing a markedly less uniform image appearance, causes less attenuation of the transform components than bilinear interpolation. This can be modelled by convolution with a simple rectangle function extending 0.5 pixels in each direction, i.e. by less than the kernel modelling bilinear interpolation, which has a transform

$$M \frac{\sin(\pi \frac{h}{M})}{\pi h} \tag{50}$$

The attenuation at the highest spatial frequency present is by a factor 0.64 (0.41 in 2-D), and at half this frequency by a factor of 0.90 (0.81 in 2-D). Simple numerical simulation confirms these expressions, suggesting that it may be possible substantially to compensate for the attenuation by division by (49) or (50) as appropriate; simulation also shows that bicubic interpolation performs much better than either.

The real-space resampling approach remains useful for the lower spatial frequencies, where all the advantages listed above are available at no cost. It is perhaps curious that the problems of the earlier approach of transform peak profile fitting are completely independent of spatial frequency: while they are no worse for the very high frequencies, they are equally no better for the low frequencies.

Summary of Findings

After so many particular statements, it may be helpful to summarise the more important conclusions.

Transform peak profiles, from which the precise frequency and complex amplitude of image Fourier components are most commonly estimated, can be seriously distorted by the overlapping tails of neighbouring peaks.

As pointed out previously by De Ruijter (1994), these tails may be greatly reduced by applying window functions in real space before transformation; however, windowing also causes degradation of the transform signal-to-noise ratio, and introduces correlation between the noise in neighbouring transform pixels (which greatly complicates the theoretical comparison of different estimators).

The more rapidly the tail decays with distance from the peak, the worse the signal-to-noise degradation: roughly speaking, windows with inverse square and cube decays lead to estimates with roughly two and three times the variance of those with no window and a consequent simple inverse decay.

The half-cosine window (inverse square decay) may be a better practical compromise than the von Hann window recommended by De Ruijter.

The familiar tools of means, variances and covariances are easily generalised to accommodate complex values such as occur in calculated image transforms.

The distribution of calculated transform pixels in the presence of random image noise does not appear to be widely familiar; the distribution is fact *normal* under very general conditions, even when the image is windowed. The real and imaginary parts are distributed normally, independently of each other; the same applies to the modulus and phase where the signal-to-noise ratio is good. The intensity in the transform of the image noise has a negative exponential distribution.

The position of a transform peak can be more closely estimated (by factors between 1.2 and 2) from four pixels than from two as recommended by De Ruijter.

The complex amplitude of a transform peak can be more closely estimated (by factors between 1.1 and 2) from four pixels than from one as recommended by De Ruijter. Several estimators for this are available, one with statistically optimal properties.

The complex amplitude of a Fourier component can also be estimated in real space. One approach, involving component-by-component summation over the entire image, and still requiring windowing to eliminate overlap by other components, provides an estimate as good as any Fourier space estimator, with lower sensitivity to errors in the frequency estimate – attractive in all respects except computational efficiency. An alternative, involving commensurate image resampling, solves the overlap problem completely without windowing and provides better estimates accordingly; however, high frequencies are underestimated (by 10-20% in a component with a 4-pixel period).

Acknowledgements

I am grateful to the Leverhulme Trust for research support, to my Department for laboratory facilities, and to the reviewers of this article for comments and careful proof-reading.

Appendix 1: Continuous and Discrete Transforms

To be explicit, we define the Fourier transform here by

$$F(\mathbf{k}) = \iint f(\mathbf{x}) \exp(-2\pi i \mathbf{k}.\mathbf{x}) d^2\mathbf{x} \qquad \text{(A1.1)}$$

so that the spatial frequency **k** measures cycles per unit distance, and the discrete Fourier transform (DFT) of an array of (*M*,*N*) pixels by

$$F_{hk} = \sum_{p=-\frac{1}{2}M}^{\frac{1}{2}M-1} \sum_{q=-\frac{1}{2}N}^{\frac{1}{2}N-1} f_{pq} \exp\{-2\pi i(\frac{hp}{M} + \frac{kq}{N})\} \qquad \text{(A1.2)}$$

so that (*p*,*q*) measure cycles per *field*, i.e. per *M* or *N* pixels. [We take (*M*,*N*) to be even for simplicity, and the origin to be at the *centre* of the array rather than at one corner, in contrast to the unfortunate choice still made by most DFT subroutines]. While the inverse transform only requires F_{hk} at integer values of (*h*,*k*), the expression in Equation A1.2 defines it for other values too, and we assume such definition in this paper.

The relationship between continuous and discrete transforms has two aspects: sampling and aliassing. If a continuous image $f(\mathbf{x})$ is sampled on a lattice with base vectors (**a**,**b**) so that

$$f_{pq} = f(p\mathbf{a} + q\mathbf{b}) \qquad \text{(A1.3)}$$

then the DFT provides samples of its transform $F(\mathbf{k})$ sampled on a *reciprocal* lattice spanned by base vectors

$$\mathbf{a}^* = \frac{1}{M} \frac{\mathbf{b} \times \mathbf{n}}{(\mathbf{a} \times \mathbf{b}).\mathbf{n}} \qquad \mathbf{b}^* = \frac{1}{N} \frac{\mathbf{n} \times \mathbf{a}}{(\mathbf{a} \times \mathbf{b}).\mathbf{n}} \qquad \text{(A1.4)}$$

where **n** is a unit vector normal to the plane of the image, so that

$$F_{hk} = F(h\mathbf{a}^* + k\mathbf{b}^*) \qquad \text{(A1.5)}$$

In the common case of a square image field ($M = N$) sampled at an interval *a* in both directions, the transform samples are provided at intervals of $1/(Ma)$ in both directions.

Aliassing is repetition and superposition at all sites of a lattice. The samples f_{pq} above are in fact taken from the image $f(\mathbf{x})$ aliassed on a lattice spanned by ($M\mathbf{a}$,$N\mathbf{b}$), and the transform samples F_{hk} from $F(\mathbf{k})$ aliassed on the reciprocal lattice ($M\mathbf{a}^*$,$N\mathbf{b}^*$). The effect is most obvious near edges, where features extending beyond one edge reappear "wrapped round" into the field at the opposite

edge. It is also indirectly apparent in the transform of images where – as is normally the case – opposite edges do not match exactly: the abrupt transition generated at the edge in the aliassed image gives rise to strong high frequency components normal to the edge, i.e. "streaking" along both transform directions.

For more information, see Saxton (1978), where these results are proved and explained.

Appendix 2: Statistics of Random Complex Variables

As noise in complex numbers is less familiar than noise in real numbers, it may be useful to note the main properties of their statistics - all of which are simply obtainable by treating the real and imaginary parts independently.

Definitions. The distribution of a random complex variable $z = x+iy$ is described by a probability distribution dependent on its real and imaginary parts $p(x,y)$; its *mean* (or *expectation*) is defined by

$$\mathrm{E}\{z\} = \bar{z} = \iint zp(x,y)dxdy \qquad \text{(A2.1)}$$

and its *variance* (the mean squared modulus of the deviation from the mean) by

$$\mathrm{var}\{z\} = \mathrm{E}\{|z-\bar{z}|^2\} = \mathrm{E}\{|z|^2\} - |\bar{z}|^2 \qquad \text{(A2.2)}$$

The standard deviation (SD), or RMS deviation, is simply the square root of the variance. Two variables are said to be *independent* if their probability distribution has the form

$$p(z_1,z_2) = p_1(z_1)p_2(z_2) \qquad \text{(A2.3)}$$

Interdependence (correlation) between two variables is measured by the *covariance*, defined by

$$\mathrm{cov}\{z_1,z_2\} = \mathrm{E}\{z_1 z_2^*\} - \bar{z}_1\bar{z}_2^* \qquad \text{(A2.4)}$$

which is zero for independent, or simply *uncorrelated,* variables, and equals the variance when the two variables are identical.

Theorems. For any random variables z, z_i, z_{1i} and z_{2i}, and fixed (complex) numbers a, a_i, a_{1i}, a_{2i}, and b, the following basic results apply:

$$\mathrm{E}\{\textstyle\sum_i a_i z_i\} = \sum_i a_i\, \mathrm{E}\{z_1\} \qquad \text{(A2.5)}$$

$$\begin{aligned}&\mathrm{cov}\,\{\textstyle\sum_i a_{1i}z_{1i}, \sum_i a_{2i}z_{2i}\}\\ &= \textstyle\sum_i\sum_j a_{1i}a_{2j}^*\,\mathrm{cov}\{z_{1i},z_{2j}\}\end{aligned} \qquad \text{(A2.6)}$$

Several simpler results are included in these, amongst them

$$\begin{aligned}&\mathrm{var}\,\{\Sigma_i a_i z_i\} = \Sigma_i |a_i|^2\ \mathrm{var}\,\{z_i\}\\ &\quad + \mathrm{Re}\,[\Sigma_i\Sigma_j a_i a_j^*\ \mathrm{cov}\{z_i,z_j\}]\end{aligned} \qquad \text{(A2.7)}$$

$$\mathrm{cov}\{a_1+z_1, a_2+z_2\} = \mathrm{cov}\{z_1,z_2\} \qquad \text{(A2.8)}$$

$$\mathrm{var}\{a+z\} = \mathrm{var}\{z\} \qquad \text{(A2.9)}$$

Appendix 3: Signal and Noise Statistics in Windowed Image Transforms

This appendix presents a substantial number of detailed results about the statistics of DFT pixels, viewed as complex numbers, as real and imaginary parts, and as modulus and phase values. If the image consists of a single Fourier component with complex amplitude c, or more realistically if we ignore the contribution of other Fourier components to one peak neighbourhood, then the image may be separated without loss of generality into two additive terms

$$f_{pq} = c\exp\{2\pi i(\frac{h_o p}{M} + \frac{k_o q}{N})\} + n_{pq} = \bar{f}_{pq} + n_{pq} \qquad \text{(A3.1)}$$

in which the first term is the image signal (expectation) $\bar{f}_{pq}$, and the second term the image noise n_{pq}, with zero mean and variance σ_{pq}^2. Its DFT, after multiplication by a window function w_{pq}, is

$$F_{hk} = \Sigma_{p,q} w_{pq} f_{pq} \exp\{-2\pi i(\frac{hp}{M} + \frac{kq}{N})\} \qquad \text{(A3.2)}$$

this appendix establishes the expectation and variance of this transform, and the covariance between any two transform pixels.

Firstly, the expectation of Equation A3.2, according to Equation A2.5, is

$$\mathrm{E}\{F_{hk}\} = \Sigma_{p,q} w_{pq} \bar{f}_{pq} \exp\{-2\pi i(\frac{hp}{M} + \frac{kq}{N})\} \qquad \text{(A3.3)}$$

or simply

$$\mathrm{E}\{F_{hk}\} = cW(h-h_o, k-k_o) \qquad \text{(A3.4)}$$

where $W(h,k) = W_{hk}$ is the DFT of the window function w_{pq}, if (h,k) are taken to be *continuous* variables.

Secondly, if the image noise is uncorrelated from point to point so that $\mathrm{cov}(n_{pq}, n_{p'q'}) = \sigma_{pq}^2 \delta_{p-p',q-q'}$, the transform variance is, according to Equation A2.7

$$\begin{aligned}\mathrm{var}\{F_{hk}\} &= \Sigma_{p,q} w_{pq}^2\ \mathrm{var}(f_{pq})\\ &= \Sigma_{p,q} w_{pq}^2 \sigma_{pq}^2\end{aligned} \qquad \text{(A3.5)}$$

If in addition the image variance has the same value σ^2 everywhere, the transform variance is simply

$$\mathrm{var}\{F_{hk}\} = \sigma^2 \Sigma_{p,q} w_{pq}^2 = \sigma^2 U_{00} \qquad \text{(A3.6)}$$

where U_{hk} is the DFT of the squared window function w_{pq}^2.

Thirdly, Equation A2.6 shows that the covariance between two transform pixels F_{hk} and $F_{h'k'}$ is

$$\Sigma_{p,q}\Sigma_{p',q'} w_{pq}w_{p'q'} \operatorname{cov}\{n_{pq},n_{p'q'}\} \times \exp\{-2\pi i(\frac{(hp-h'p')}{M}+\frac{kq-k'q'}{N})\} \quad \text{(A3.7)}$$

If the image noise is uncorrelated from point to point as before, this is

$$\Sigma_{p,q} w_{pq}^2\sigma_{pq}^2 \exp\{-2\pi i(\frac{(h-h')p}{M}+\frac{(k-k')q}{N})\} \quad \text{(A3.8)}$$

If in addition the noise variance has the same value σ^2 everywhere, the covariance is simply

$$\operatorname{cov}\{F_{hk},F_{h'k'}\} = \sigma^2 U_{h-h',k-k'} \quad \text{(A3.9)}$$

Next, we establish the corresponding statistics for the real and imaginary parts of transform pixels, writing $F_{hk} = G_{hk}+iH_{hk}$. These prove to be slightly more complicated.

$$\begin{aligned} G_{hk} &= \Sigma_{p,q} w_{pq} f_{pq} \cos\{2\pi(\frac{hp}{M}+\frac{kq}{N})\} \\ H_{hk} &= -\Sigma_{p,q} w_{pq} f_{pq} \sin\{2\pi(\frac{hp}{M}+\frac{kq}{N})\} \end{aligned} \quad \text{(A3.10)}$$

from which

$$\begin{aligned} \operatorname{var}\{G_{hk}\} &= \Sigma_{p,q} w_{pq}^2\sigma_{pq}^2 \cos^2\{2\pi(\frac{hp}{M}+\frac{2kq}{N})\} \\ &= \Sigma_{p,q} w_{pq}^2\sigma_{pq}^2 \tfrac{1}{2}[1+\cos\{2\pi(\frac{2hp}{M}+\frac{2kq}{N})\}] \end{aligned} \quad \text{(A3.11)}$$

If the image variance is σ^2 everywhere, this simplifies to

$$\operatorname{var}\{G_{hk}\} = \frac{\sigma^2}{2}[U_{00}+\tfrac{1}{2}(U_{2h,2k}+U_{-2h,-2k})] \quad \text{(A3.12)}$$

Except when $|h|$ is near 0 or $M/2$, and $|k|$ is near 0 or $N/2$ (i.e., except only for pixels at the transform centre, its corners, or the middle of an edge), U_{hk} is zero for the windows considered (cf. Table 2), permitting the simplification

$$\operatorname{var}\{G_{hk}\} = \tfrac{1}{2}\sigma^2 U_{00} \quad \text{(A3.13)}$$

In the same way, $\operatorname{var}\{H_{hk}\}$ is found to be the same as $\operatorname{var}\{G_{hk}\}$ under the same assumptions about (h,k); and we note that each is exactly half of the previously obtained variance in the complex value (Equation A3.5), (Equation A3.6).

The covariance $\operatorname{cov}\{G_{hk},H_{hk}\}$ is found to be zero under the same assumptions about (h,k); and indeed without restriction on (h,k) if the image window is centro-symmetric.

Similar arguments also show the following exact results about the covariance between the real and imaginary parts of *different* pixels.

$$\begin{aligned} \operatorname{cov}\{G_{hk},G_{h'k'}\} &= \tfrac{1}{2}\sigma^2[U_{h-h',k-k'} + U_{h+h',k+k'}] \\ \operatorname{cov}\{H_{hk},H_{h'k'}\} &= \tfrac{1}{2}\sigma^2[U_{h-h',k-k'} - U_{h+h',k+k'}] \\ \operatorname{cov}\{G_{hk},H_{h'k'}\} &= \operatorname{cov}\{H_{hk},G_{h'k'}\} = 0 \end{aligned} \quad \text{(A3.14)}$$

These expressions imply correlation between opposite pixels $(h',k'\approx -h,-k)$ as well as neighbouring pixels $(h',k'\approx h,k)$ – unsurprising given the conjugate symmetry of the transform. However, if we are concerned only with correlation between neighbouring pixels, $U_{h+h',k+k'}$ is zero under the same assumptions about (h,k) as cited above, and the term can be discarded leaving

$$\begin{aligned} \operatorname{cov}\{G_{hk},G_{h'k'}\} &= \tfrac{1}{2}\sigma^2 U_{h-h',k-k'} \\ \operatorname{cov}\{H_{hk},H_{h'k'}\} &= \tfrac{1}{2}\sigma^2 U_{h-h',k-k'} \\ \operatorname{cov}\{G_{hk},H_{h'k'}\} &= \operatorname{cov}\{H_{hk},G_{h'k'}\} = 0 \end{aligned} \quad \text{(A3.15)}$$

and we note that the first two of these are exactly half of the previously obtained covariance between the complex values (A3.9).

Finally, we obtain the corresponding statistics for the modulus of the transform pixels; further approximation is necessary for this purpose, with an assumption that its SD is much smaller than the modulus of its expectation, i.e. that the noise is much smaller than the signal.

Under this assumption, we can obtain an approximation to a change in the modulus $|F_{hk}|$ in the form

$$\delta|F_{hk}| = \frac{G_{hk}}{|\overline{F}_{hk}|}\delta G_{hk} + \frac{H_{hk}}{|\overline{F}_{hk}|}\delta H_{hk} \quad \text{(A3.16)}$$

by differentiation. In view of (A2.7), (A3.13) and (A3.15), the variance of this is given by

$$\begin{aligned} \operatorname{var}\{|F_{hk}|\} &= \frac{G_{hk}^2}{|\overline{F}_{hk}|^2}\operatorname{var}\{G_{hk}\} + \frac{H_{hk}^2}{|\overline{F}_{hk}|^2}\operatorname{var}\{H_{hk}\} \\ &= \tfrac{1}{2}\sigma^2 U_{00} \end{aligned} \quad \text{(A3.17)}$$

being the same as the variance of the real and imaginary parts. The variance of the phase $\theta_{hk} = \arg(F_{hk})$ is similarly obtained from

$$\delta\theta_{hk} = -\frac{H_{hk}}{|\overline{F}_{hk}|^2}\delta G_{hk} + \frac{G_{hk}}{|\overline{F}_{hk}|^2}\delta H_{hk} \quad \text{(A3.18)}$$

as

$$\mathrm{var}\{\theta_{hk}\} = \tfrac{1}{2}\frac{\sigma^2}{|\overline{F}_{hk}|^2}U_{00} \qquad \text{(A3.19)}$$

while the covariance between the two, according to (A2.6), is zero.

Within the scope of the same low noise approximation, the covariance between the moduli of two different pixels is

$$\mathrm{cov}\{\frac{G_{hk}}{|\overline{F}_{hk}|}\delta G_{hk} + \frac{H_{hk}}{|\overline{F}_{hk}|}\delta H_{hk}, \frac{G_{h'k'}}{|\overline{F}_{h'k'}|}\delta G_{h'k'} + \frac{H_{h'k'}}{|\overline{F}_{h'k'}|}\delta H_{h'k'}\} \qquad \text{(A3.20)}$$

which in view of (A2.6) and (A3.15) may be simplified to give

$$\mathrm{cov}\ \{|F_{hk}|,|F_{h'k'}|\} = \tfrac{1}{2}\beta\sigma\, U_{h-h',k-k'} \qquad \text{(A3.21)}$$

in which the parameter β is

$$\beta = \frac{\mathrm{Re}\{\overline{F}_{hk}\overline{F}^*_{h'k'}\}}{|\overline{F}_{hk}\overline{F}^*_{h'k'}|} = \cos(\theta_{hk} - \theta_{h'k'}) \qquad \text{(A3.22)}$$

showing that the covariance vanishes when the pixels have phase differing by π/2.

Finally, the covariance between the phases of two different pixels is

$$\mathrm{cov}\{\theta_{hk}, \theta_{h'k'}\} = \tfrac{1}{2}\beta\frac{\sigma^2}{|F|^2}U_{h-h',k-k'} \qquad \text{(A3.23)}$$

Appendix 4: Noise Distributions in Windowed Image Transforms

This appendix establishes first the distribution of the real and imaginary parts of the DFT F_{hk} of a windowed image f_{pq}, by showing that central sections through the distribution have the same gaussian profile regardless of the section direction. A section in a direction at an angle ϕ to the real axis can be considered as the real part of the transform multiplied by an arbitrary phase-shifting factor $\exp\{-i\phi\}$:

$$G'_{hk} = \mathrm{Re}\{\Sigma_{p,q} w_{pq} f_{pq} \exp[-2\pi i(\phi + \frac{hp}{M} + \frac{kq}{N})]\}$$
$$= \Sigma_{p,q} w_{pq} f_{pq} \{\cos\phi\cos[2\pi(\frac{hp}{M} + \frac{kq}{N})] - \sin\phi\sin[2\pi(\frac{hp}{M} + \frac{kq}{N})]\} \qquad \text{(A4.1)}$$

The first step in the argument is to show that this has a gaussian (normal) distribution, regardless of the value of ϕ. This follows immediately from the well-known central limit theorem, which asserts that the sum (or weighted sum) of a large number of random variables has a near-gaussian distribution, regardless of the distribution of the individual variables summed.

The second step is to show that the expectation and variance are also independent of ϕ. If the image has zero expectation (as is the case for noise images), the expectation of G'_{hk} is zero. If the image pixels are uncorrelated, and their variance is σ^2 everywhere, the variance of G'_{hk} is

$$\sigma^2\{\cos^2\phi\Sigma_{p,q}w_{pq}^2\cos^2[2\pi(\frac{hp}{M}+\frac{kq}{N})] + \sin^2\phi\Sigma_{p,q}w_{pq}^2\sin^2[2\pi(\frac{hp}{M}+\frac{kq}{N})] - 2\cos\phi\sin\phi\Sigma_{p,q}w_{pq}^2\cos[\ldots]\sin[\ldots]\} \qquad \text{(A4.2)}$$

This can be re-expressed in the form

$$\sigma^2\{\cos^2\phi\Sigma_{p,q}w_{pq}^2\tfrac{1}{2}[1+\cos\{2\pi(\frac{2hp}{M}+\frac{2kq}{N})\}] + \sin^2\phi\Sigma_{p,q}w_{pq}^2\tfrac{1}{2}[1-\cos\{2\pi(\frac{2hp}{M}+\frac{2kq}{N})\}] - 2\cos\phi\sin\phi\Sigma_{p,q}w_{pq}^2\tfrac{1}{2}\sin\{2\pi(\frac{2hp}{M}+\frac{2kq}{N})\}\} \qquad \text{(A4.3)}$$

and so as

$$\tfrac{1}{2}\sigma^2\{\cos^2\phi[U_{00} + \tfrac{1}{2}(U_{2h,2k} + U_{-2h,-2k})] + \sin^2\phi[U_{00} - \tfrac{1}{2}(U_{2h,2k} + U_{-2h,-2k})] - i\cos\phi\sin\phi[U_{2h,2k} - U_{-2h,-2k}]\} \qquad \text{(A4.4)}$$

Now except when $|h|$ is near 0 or $M/2$, and $|k|$ is near 0 or $N/2$ (i.e., except only for pixels at the transform centre, its corners, or the middle of an edge), U_{hk} is zero for the windows considered (cf. Table 2), permitting the simplification

$$\mathrm{var}\{G'_{hk}\} = \tfrac{1}{2}\sigma^2 U_{00} \qquad \text{(A4.5)}$$

regardless of ϕ.

The distribution $p(G,H)$ thus has the same zero-mean gaussian section in all directions, and must be the 2-D gaussian distribution[5]:

$$p(G,H) = A\exp(-\frac{G^2+H^2}{\sigma^2 U_{00}}) \qquad \text{(A4.6)}$$

[5] In the absence of windowing, the exceptional cases F_{00}, $F_{-M/2,0}$, $F_{0,-N/2}$ and $F_{-M/2,-N/2}$ are the few transform pixels with distributions restricted to being *real*; with windowing, the distribution at neighbouring pixels is also affected.

The separability of this shows that G,H are uncorrelated (and indeed independent).

Next we establish the distribution of the modulus and phase of a transform pixel when the expectation is not zero. We write $F=G+iH$ again, but taking G,H to have non-zero expectations. Provided their variation is a small fraction of their mean, we can obtain an approximation to the change in the modulus |I|in the form

$$\delta|F| = \frac{G}{|\overline{F}|}\delta G + \frac{H}{|\overline{F}|}\delta H \qquad \text{(A4.7)}$$

by differentiation. As the sum of two independent gaussian variables, this also has a gaussian distribution; its variance is given in (A3.17) above. We note that the distribution is simply that of any section through the complex pixel distribution[6].

The distribution of the phase $\theta = \arg(F)$ is obtained in the same way from

$$\delta\theta = -\frac{H}{|\overline{F}|^2}\delta G + \frac{G}{|\overline{F}|^2}\delta H \qquad \text{(A4.8)}$$

being once again gaussian; its variance is given in (A3.19) above.

Appendix 5: Variance of Spacing Estimators

Exact calculation of the expected SD in the spacing estimators is complicated by the correlation introduced by windowing the image.

For *small* changes at least in $|p_{00}|$, $|p_{10}|$, $|p_{01}|$ and $|p_{11}|$, an approximation to the change in the estimator u can be obtained in the form

$$\delta u' = c_{00}\delta|p_{00}| + c_{10}\delta|p_{10}| + c_{01}\delta|p_{01}| + c_{11}\delta|p_{11}| \qquad \text{(A5.1)}$$

by differentiation. For the unit window and the estimator (13), the coefficients are

$$\begin{aligned}
c_{00} &= \frac{x[x|p_{10}|+(1-x)|p_{11}|]}{[x|p_{10}|+(1-x)|p_{11}|+x|p_{00}|+(1-x)|p_{01}|]^2}\\
c_{10} &= \frac{x[x|p_{00}|+(1-x)|p_{01}|]}{[x|p_{10}|+(1-x)|p_{11}|+x|p_{00}|+(1-x)|p_{01}|]^2}\\
c_{01} &= \frac{(1-x)[x|p_{10}|+(1-x)|p_{11}|]}{[x|p_{10}|+(1-x)|p_{11}|+x|p_{00}|+(1-x)|p_{01}|]^2}\\
c_{11} &= \frac{(1-x)[x|p_{00}|+(1-x)|p_{01}|]}{[x|p_{10}|+(1-x)|p_{11}|+x|p_{00}|+(1-x)|p_{01}|]^2}
\end{aligned} \qquad \text{(A5.2)}$$

and they happen to be exactly twice and thrice these expressions for the half-cosine and von Hann windows. Under the assumption of *small* changes, we treat these coefficients subsequently as constants, equal to their expectations, approximated in turn by the values of the coefficients when the pixels p_{00} etc. take their expected values as given in (9).

The variance of (A5.1) can be evaluated using (A2.7) to reduce it to a sum of variances in and covariances between the transform pixel increments $\delta|p_{00}|$ etc., which according to (A2.8) and (A2.9) are equal to those in and between the pixel moduli themselves. For image noise uncorrelated from point to point, and with the same variance σ^2 everywhere, these are in turn given by (A3.17) and (A3.21); the parameter β in (A3.21) is here 1, as the pixel expectations all have the same phase. In this way, we obtain

$$\begin{aligned}
\mathrm{var}\{u'\} = {} & [|c_{00}|^2+|c_{10}|^2+|c_{01}|^2+|c_{11}|^2]\tfrac{1}{2}\sigma^2 U_{00}\\
& +2\,\mathrm{Re}\{[c_{00}c_{10}+c_{01}c_{11}]\}\tfrac{1}{2}\sigma^2 U_{10}\\
& +2\,\mathrm{Re}\{[c_{00}c_{01}+c_{10}c_{11}]\}\tfrac{1}{2}\sigma^2 U_{01}\\
& +2\,\mathrm{Re}\{[c_{00}c_{11}+c_{10}c_{01}]\}\tfrac{1}{2}\sigma^2 U_{10}
\end{aligned} \qquad \text{(A5.3)}$$

The variance for the simpler two-point estimators (10-12) is of course simply the value of this when $x = 1$; algebraically the result is

$$\begin{aligned}
\mathrm{var}\{u'\} = {} & [|c_{00}|^2+|c_{10}|^2]\tfrac{1}{2}\sigma^2 U_{00}\\
& +2\,\mathrm{Re}\{c_{00}c_{10}\}\tfrac{1}{2}\sigma^2 U_{10}
\end{aligned} \qquad \text{(A5.4)}$$

with coefficients given, for the unit window, by the simpler expressions

$$\begin{aligned}
c_{00} &= \frac{|p_{10}|}{[|p_{10}|+|p_{00}|]^2}\\
c_{10} &= \frac{|p_{00}|}{[|p_{10}|+|p_{00}|]^2}
\end{aligned} \qquad \text{(A5.5)}$$

The value of the variance depends (via the pixels p_{00} etc.) on the actual value of (h,k) (i.e., where the true peak lies) as well as on the relative weight given to the two rows. Tables A5.1 gives values for $u,v = 0$, 0.25 and 0.5 for the unit window, and those following give them for the half-cosine and von Hann windows, in each case with the approximate optimum weighting proposed in (28).

Table A5.1. Spacing estimator variance: unit window $|c|^2 MN\mathrm{var}\{u\}/\sigma^2$ for weighting factor $x = 1\text{-}v$

	u=0	u=0.25	u=0.5
v=0	0.5	0.22	0.15
v=0.25	0.56	0.24	0.17
v=0.5	0.62	0.27	0.19

[6] This is arguably obvious geometrically, as is the corresponding result for the distribution of the phase.

Table A5.2. Spacing estimator variance: half-cosine window
$|c|^2 MN\mathrm{var}\{u\}/\sigma^2$ for weighting factor $x = 1\text{-}2v^2$

	u=0	u=0.25	u=0.5
v=0	0.75	0.40	0.31
v=0.25	0.83	0.44	0.34
v=0.5	0.91	0.48	0.38

Table A5.3. Spacing estimator variance: von Hann window
$|c|^2 MN\mathrm{var}\{u\}/\sigma^2$ for weighting factor $x = 1\text{-}2v^2$

	u=0	u=0.25	u=0.5
v=0	1.17	0.71	0.59
v=0.25	1.26	0.77	0.63
v=0.5	1.35	0.82	0.68

Table A6.1. One-pixel amplitude estimator variance: unit window
$MN\mathrm{var}\{c'\}/\sigma^2$

	u=0	u=0.25	u=0.5
v=0	1	1.23	2.47
v=0.25	1.23	1.52	3.04
v=0.5	2.47	3.04	6.09

Table A6.2. One-pixel amplitude estimator variance: half-cosine window
$MN\mathrm{var}\{c'\}/\sigma^2$

	u=0	u=0.25	u=0.5
v=0	1.52	1.71	2.47
v=0.25	1.71	1.93	2.78
v=0.5	2.47	2.78	4.00

Table A6.3. One-pixel amplitude estimator variance: von Hann window
$MN\mathrm{var}\{c'\}/\sigma^2$

	u=0	u=0.25	u=0.5
v=0	2.25	2.44	3.12
v=0.25	2.44	2.64	3.39
v=0.5	3.12	3.39	4.33

Table A6.4. Four-pixel LS amplitude estimator variance: unit window
$MN\mathrm{var}\{c'\}/\sigma^2$

	u=0	u=0.25	u=0.5
v=0	1	1.11	1.23
v=0.25	1.11	1.23	1.37
v=0.5	1.23	1.37	1.52

Table A6.5. Four-pixel LS amplitude estimator variance: half-cosine window
$MN\mathrm{var}\{c'\}/\sigma^2$

	u=0	u=0.25	u=0.5
v=0	2.08	2.12	2.17
v=0.25	2.12	2.16	2.21
v=0.5	2.17	2.21	2.25

Table A6.6. Four-pixel LS amplitude estimator variance: von Hann window
$MN\mathrm{var}\{c'\}/\sigma^2$

	u=0	u=0.25	u=0.5
v=0	3.39	3.23	3.19
v=0.25	3.23	3.08	3.05
v=0.5	3.19	3.05	3.01

Appendix 6: Variance of Amplitude Estimators

This appendix collects various results about amplitude estimators. Firstly, table A6.1 gives values of the variance (32) of the one-pixel amplitude estimator (31) for u,v = 0, 0.25 and 0.5 for the unit window, and the two following give values for the half cosine and von Hann windows respectively.

Secondly, the estimator achieving a least-squares fit to an arbitrary set of transform pixels p_i with expectations cW_i is established. We seek the c' that minimises the summed squared difference between observed and predicted pixels, i.e.

$$\Sigma_i (c'W_i - p_i)(c'^*W_i^* - p_i^*) \qquad \text{(A6.1)}$$

setting the derivative of this w.r.t. c'^* equal to zero gives

$$\Sigma_i (c'W_i - p_i)W_i^* = 0 \qquad \text{(A6.2)}$$

from which we obtain the estimator in (33), namely

$$c' = \frac{\Sigma_i W_i^* p_i}{V}, \text{ with } V = \Sigma_i |W_i|^2 \qquad \text{(A6.3)}$$

We now calculate the variance of the particular estimator involving the four pixels p_{00}, p_{10}, p_{01} and p_{11} surrounding the transform peak. If the image pixels are uncorrelated, and their variance is σ^2 everywhere, this is given, according to (A2.7), (A3.6) and (A3.9) by

$$\begin{aligned}\mathrm{var}\{u\} = (&[|W_{00}|^2+|W_{10}|^2+|W_{01}|^2+|W_{11}|^2\,]\sigma^2U_{00}\\ &+2\,\mathrm{Re}\{[W_{00}W_{10}^*+W_{01}W_{11}^*]\}\sigma^2U_{10}\\ &+2\,\mathrm{Re}\{[W_{00}W_{01}^*+W_{10}W_{11}^*]\}\sigma^2U_{01}\\ &+2\,\mathrm{Re}\{[W_{00}W_{11}^*+W_{10}W_{01}^*]\}\sigma^2U_{10})/V^2\end{aligned} \tag{A6.4}$$

Table A6.4 gives values of this for u,v = 0, 0.25 and 0.5 for the unit window, and the two following give values for the half cosine and von Hann windows respectively.

Appendix 7 gives a more general form of the variance applicable to least-squares estimators as in (A6.3) for arbitrary sets of pixels; here we give the result when the set is extended to embrace the entire transform field. For this purpose, we write transform pixels in the form

$$p_{hk} = F_{hk} = cW_{hk} + N_{hk} \tag{A6.5}$$

where N_{hk} is the DFT of the *windowed* image noise $w_{pq}n_{pq}$, so that the estimator takes the form

$$c' = \frac{1}{V}\Sigma_{h,k}W_{hk}^*cW_{hk} + \frac{1}{V}\Sigma_{h,k}W_{hk}^*N_{hk} \tag{A6.6}$$

in which the signal and noise terms are separated. We seek the variance of the latter.

We avoid extending the double sum of cross terms appearing in (A6.4) to the entire transform field, by transferring the sum to real space, using Parseval's theorem,

$$\Sigma_{h,k}W_{hk}^*N_{hk} = MN\Sigma_{p,q}w_{pq}.w_{pq}n_{pq} \tag{A6.7}$$

This re-expresses the noise term in the form

$$\frac{MN}{V}\Sigma_{p,q}w_{pq}^2n_{pq} \tag{A6.8}$$

in which (on the same assumptions that image pixels are uncorrelated) the terms are no longer correlated. In addition, the same theorem allows us re-express V as a real-space sum:

$$V = \Sigma_{h,k}|W_{hk}^*|^2 = MN\Sigma_{p,q}w_{pq}^2 \tag{A6.9}$$

Thus, if the image pixels have the same variance σ^2 everywhere, the variance in the amplitude estimator, according to (A2.7), can finally be seen to be

$$\mathrm{var}\{c'\} = \sigma^2\frac{\Sigma_{p,q}w_{pq}^4}{(\Sigma_{p,q}w_{pq}^2)^2} \tag{A6.10}$$

Appendix 7: Optimal Amplitude Estimators

This appendix derives the optimal linear estimate of a transform component amplitude from a set of transform pixels, allowing for their inter-correlation, and illustrating the general results by the particular case of an estimate based on the four pixels in a 2×2 block around the peak.

We define a 4-component coefficient vector r_i with a data vector

$$p_i = \begin{pmatrix}p_{00}\\ p_{10}\\ p_{01}\\ p_{11}\end{pmatrix} \tag{A7.1}$$

a similar vector W_i whose elements are the corresponding peak profile values (20), and a (real symmetric) *variance-covariance* matrix whose elements are, according to (9),

$$c_{ij} = \mathrm{cov}\{p_i, p_j\} = \begin{pmatrix}U_{00} & U_{10} & U_{01} & U_{11}\\ U_{10} & U_{00} & U_{11} & U_{01}\\ U_{01} & U_{11} & U_{00} & U_{10}\\ U_{11} & U_{01} & U_{10} & U_{00}\end{pmatrix} \tag{A7.2}$$

We consider a linear estimator

$$c' = r_1p_{00} + r_2p_{10} + r_3p_{01} + r_4p_{11} = r_ip_i \tag{A7.3}$$

if we adopt the usual *summation* convention (summation over repeated suffices), which has a variance, according to (A2.6)

$$\mathrm{var}\{c'\} = r_ir_j^*c_{ij} \tag{A7.4}$$

The constraint that the expectation of (A7.3) is c becomes

$$\begin{aligned}E\{c'\} = r_iE\{p_i\} = cr_iW_i = c\\ \Rightarrow r_iW_i = 1\end{aligned} \tag{A7.5}$$

The minimisation of (A7.4) subject to (A7.5) is equivalent to the minimisation of

$$r_ir_j^*c_{ij} + \lambda(r_iW_i - 1) + \mu(r_i^*W_i^* - 1) \tag{A7.6}$$

(with *real* multipliers λ,μ) simultaneous with (A7.5)[7]; setting the derivative w.r.t. r_k equal to zero gives

$$c_{kj}r_j^* + \lambda W_k = 0 \tag{A7.7}$$

Applying the inverse matrix c_{ik}^{-1} (for which $c_{ik}^{-1}c_{kj} = \delta_{ij}$) elicits the solution

$$r_i = -\lambda c_{ik}^{-1}W_k^* \tag{A7.8}$$

[7] Although a fully complex minimisation has been performed for generality, the problem considered here actually involves real values only.

Table A7.1. Four-pixel optimal amplitude estimator variance: half-cosine window
$MN\mathrm{var}\{c'\}/\sigma^2$

	u=0	u=0.25	u=0.5
v=0	1.42	1.63	1.78
V=0.25	1.63	1.88	2.05
V=0.5	1.78	2.05	2.25

and the multiplier λ is chosen as $\lambda = -1/(r_i W_i)$ to satisfy (A7.5): thus the required coefficients for the optimal estimator are

$$r_i = \frac{c_{ik}^{-1} W_k^*}{W_i c_{ik}^{-1} W_k^*} = \frac{\mathbf{C}^{-1}\mathbf{W}*}{\mathbf{W}.\mathbf{C}^{-1}\mathbf{W}*} \qquad \text{(A7.9)}$$

in which we have switched finally to a conventional vector/matrix notation. The variance is given by

$$\mathrm{var}\{c'\} = \mathbf{r}.\mathbf{C}\mathbf{r}* = \frac{\mathbf{C}^{-1}\mathbf{W}^*.\mathbf{C}\mathbf{C}^{-1}\mathbf{W}}{[\mathbf{W}.\mathbf{C}^{-1}\mathbf{W}^*]^2} = \frac{1}{\mathbf{W}.\mathbf{C}^{-1}\mathbf{W}^*} \qquad \text{(A7.10)}$$

When the data p_i are uncorrelated (e.g. for the unit window), the matrix c_{ij} is diagonal, and if the pixels all have the same variance the inverse is also diagonal with equal elements; in these circumstances, (A7.9) reduces to

$$r_i = \frac{W_i^*}{\Sigma_i W_i W_i^*} \qquad \text{(A7.11)}$$

as in the least-squares estimate (33). When other windows are employed however, the estimator (A7.9) has different coefficients from those of (33). For example, with the half-cosine window, the least-squares estimator at $(u,v) = (0,0)$ is
$c' = 0.81p_{00} + 0.27p_{10} + 0.27p_{01} + 0.09p_{11}$,
with variance $2.08\sigma^2/MN$, while the optimal estimator is
$c' = 1.15p_{00} - 0.23p_{10} - 0.23p_{01} + 0.05p_{11}$,
with variance $1.42\sigma^2/MN$.

Values of the variance (A7.10) in the optimal estimators for u,v = 0, 0.25 and 0.5 are given in tables A7.1 and A7.2 for the half-cosine and von Hann windows respectively; these may be compared with tables A6.5 and A6.6 for the least-squares estimator.

Table A7.2. Four-pixel optimal amplitude estimator variance: von Hann window
$MN\mathrm{var}\{c'\}/\sigma^2$

	u=0	u=0.25	u=0.5
v=0	2.04	2.31	2.48
V=0.25	2.31	2.62	2.81
V=0.5	2.48	2.81	3.01

References

Aebi U, Smith PR, Dubochet J, Henry C, Kellenberger E (1973) A study of the structure of the T-layer of *Bacillus brevis*. J Supramol Struct **1**: 498-522.

De Ruijter WJ (1994) Measurement of lattice-fringe vectors from digital HREM images: Theory and simulations. J Comp Assisted Microsc. **6**: 195-212.

Hamming RW (1977) Digital Filters, Prentice-Hall, New Jersey, §5.10.

Hÿtch MJ, Stobbs WM (1994) Quantitative comparison of high resolution TEM images with image simulations. Ultramicroscopy **53:** 191-203.

Saxton WO (1978) Computer Techniques for Image Processing in Electron Microscopy; Academic Press, New York, Ch. 2.

Saxton WO (1992) Imaging deformed proteins: characterising the state fully; Appendix D. Scanning Microscopy Suppl. **6:** 53-70.

Saxton WO, Baumeister W (1982). The correlation averaging of a regularly arranged bacterial cell envelope protein. J Microsc **127:** 127-138.

Discussion with Reviewers

M.J. Hÿtch: It is implicitly assumed that the image intensity in a pixel f_{pq} takes the value of the electron density function I_{pq} sampled at position (p,q). In the case of a CCD, for example, the intensity collected in each pixel is in fact the electron density function integrated over the area covered by the pixel. It is not strictly speaking I_{pq}. Evidently, for slowly varying functions this will not matter much, but for fringes having wavelengths of the order of 4 pixels or less this must become important. Has this problem been addressed and if so, what effect does this have on our measurement of F_{hk}?

Author: In the case of a CCD at least, the sole effect is attenuation of measured amplitudes by convolution of the true image intensity with the CCD pixel shape; the reduction is by 10% for 4-pixel periods parallel to either axes and 36% for 2-pixel periods. More serious problems arise with other forms of digitisation and pre-processing however which can invalidate the assumption that noise is

uncorrelated between neighbouring pixels; some of these are taken up in Saxton (1998).

M.J. Hÿtch: Fringe spacings and phases can be measured in real-space (and not just the amplitudes) using the methods of holographic reconstruction; this will be something to look into.
Author: I agree.

H. Kohl: Further references to textbooks or review articles dealing with the statistical properties of images would probably help the non-specialist reader to follow the discussion. One article coming to mind is Slump and Ferwerda (1986).
Author: The area is not well written up; see however Dainty and Shaw (1974) or Rosenfeld and Kak (1976); and some of the ideas appear in textbooks on basic statistics
.
L.D. Marks: There are other methods of determining a power spectrum rather than simply taking a discrete FFT, for instance the maximum entropy or all-poles method (see for instance Press *et al.* (1992). Are there any advantages in using these methods, particularly for noisy data?
Author: That approach has indeed been little used in electron microscopy so far (but see Anderson *et al*., 1989), and I cannot comment usefully except to say I believe it almost certainly deserves proper examination.

Additional References

Anderson DM, Martin DC, Thomas EL (1989) Maximum-entropy data restoration using both real- and Fourier-space analysis. Acta Cryst **A45**: 686-698.

Dainty JC, Shaw R (1974). Image Science. London, Academic Press, Ch.8.

Press WH, Flannery BP, Teukolsky SA, Vetterling WT, (1986) Numerical Recipes. Cambridge University Press, Cambridge, §12.8.

Rosenfeld A, Kak AC (1976) Digital Picture Processing. Academic Press, New York, Ch.2.

Saxton WO (1998). Quantitative comparison of images and transforms J. Microsc **190**: 52-60.

Slump CH, Ferwerda HA (1986) Statistical aspects of image handling in low-dose electron microscopy of biological material. Electronics & Electron Physics **66**: 201-308.