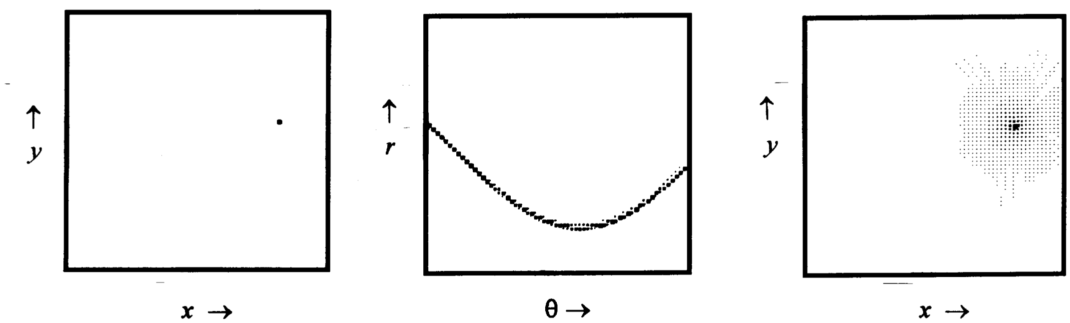

Figure 1.1: Point-type object, its scan data, and the image recovered through immediate back projection.

1.5 Back Projection method

The idea of back projection is conceptually very straightforward, it iseasy to implement, and the computational cost is moderate. In its most direct form, the reconstruction comes out 'smeared'. However, the addition of a simple filter all but resolves this. It is little surprise that Filtered Back Projection (FBP) has become the most widely used reconstruction process in the medical community (where it very well meets the requirements placed on it).

For an $n \times n$ image reconstruction, the FBP method will cost $O(n^3)$ operations. In many contexts, this cost is acceptable (or rather, it became accepted at a time when the major alternatives were even more costly). The computing time using FBP is quite fast in comparison to the other tasks involved (such as patient handling etc.). However, in an application such as micro-tomography, the situation is very different. The resolution can here be as high as 1000 simultaneously recorded slices, each to be imaged on a 1000×1000 grid (giving a 3-D 1μm resolution throughout a cubic millimeter sample). Each full inversion for such a cube of using back projections would cost on the order of 10^{12} operations. This is likely to become a slow process in comparison in comparison with the rapid one of automated sample handling (where data for the 1000 slices are collected simultaneously by using a 2-D rather that a 1-D array of X-ray sensors). We describe in Section 1.6 an inversion algorithm which cuts this cost by some orders of magnitude - in this case to about 10^{10} operations.

1.5.1 Immediate back projection

The left part of Figure 1.1 shows a point-type object, and central part its scan data. Figure 1.2 illustrates the simplest form of back projection: For every angle θ we draw parallel bands across the image area, with darkness corresponding to the absorption that was recorded for that same angle θ. The right part of Figure 1.1 shows the result of this process in this case of the point object. The only error is that the point has turned into 'smeared out' cone-type mound. The sharp edge of the original point object has been lost, as areas near the point are also covered by some of the bands. However, the position of the recovered mound is precisely the same as that of the original point-object. Also, the amplitude and shape of the mound is position invariant - it takes the same values wherever the original point object was located.

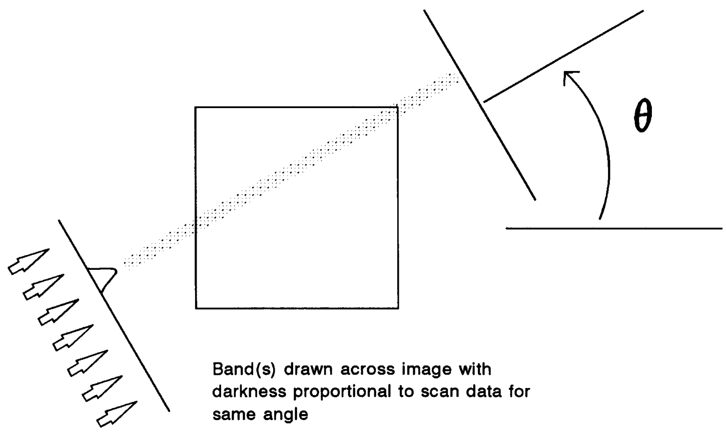

Figure 1.2: Principle behind back projection (when applied immediately to scan data - no filtering) shown here in the case of a point object.

To appreciate the significance of this example with a point object, we need to note that both the scanning and the back projection phases are *linear*.

If there had been two point objects, the scan data would just have been the *sum* of the scan data for the two objects if recorded independently. Similarly, the back projection produces in that case a result which is the sum of the back projections of the two objects, if they were treated separately. Finally, the darkness of the back projected result of each object (by itself) is proportional to the original darkness. The process of scanning followed by back projection satisfies the two criteria of a function (here a matrix-valued function with a matrix input) to be *linear*:

$$\begin{array}{ll} f(x+y) = f(x) + f(y) & \text{, and} \\ f(\alpha x) = \alpha\, f(x) & . \end{array}$$

From this linearity follows the completely critical conclusion that the reconstruction must also work for full images (and not just point objects).With some repetition of the observations above: Wherever the imaged object x had a gray pixel, the image f will feature a smeared one centered at the same location, and with a darkness proportional to the darkness of the one in the original. From the linearity follows that if x was a sum of two images with a different gray pixel in each, f will become the sum of the two corresponding images. Continuing this observation: Since every image is a combination of pixels, this linearity implies that f must become a (smeared) representation of the original object. Immediate back projection using the scan data shown in Figures I1.3.1.4 leads to the reconstruction seen in Figure 1.3.

1.5.2 Filtered back projection

Comparing the original object in Figure I.1.3.1.2 with Figure 1.3 (both featuring the same grid density of 63×63 points), we clearly see a loss in sharpness. Hence, we look for some way to enhance the output from direct back projection in order to reduce the smearing. The step from

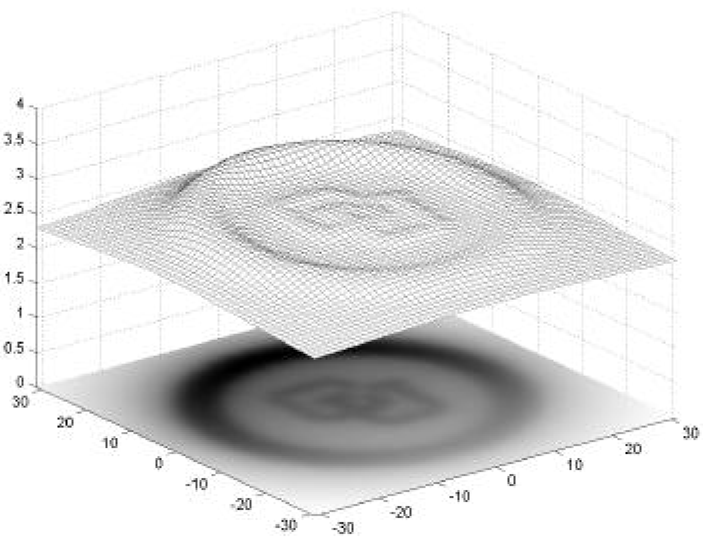

Figure 1.3: Immediate back projection of the scan data for the test object.

original object to scan data is essentially outside our control (dictated by X-rays and physics). Options remaining include

- Based on the smeared image, apply some filter which sharpens all gradients (such filters can for ex. be based on FFTs), and
- Since the cause of the smearing is understood (for the point image), try to alter the scan data in a way that off-sets the back projection smearing.

Both options above are viable; filtered back projection pursues the second one. Two key questions become:

1. Does there exist any special type of (simulated) scan data for which the back projection method will give a nearly pointlike result?
2. Is there any operator (linear, location preserving, and not altering the darkness at the point itself) that we can apply to turn the actual the scan data for the point-type test object into the form that we looked for in the point above?

To address issue 1 first; we note that we can replace the 'single-hump' data by a 'hump' at the same place, but with a bright band on each side of it. The contributions for all the angles will still superimpose to an equally dark spot at precisely the desired location, but at nearby locations, the bright sidebands just might cancel some of the undesired darkness, as the contributions for different angles θ are superposed.

To turn the scan data for a fixed angle, (0,0,...,0,1,0,...,0)^T, into a vector with a bright (negative) entry on each side of the "one" - wherever it is located - can be achieved by multiplying the scan data from the left by a symmetric $n \times n$ tri-diagonal Toeplitz band matrix

$$E = \begin{bmatrix} 1 & -\beta & & & & & & \\ -\beta & 1 & -\beta & & & & & \\ & -\beta & 1 & -\beta & & & & \\ & & \ddots & \ddots & \ddots & & & \\ & & & \ddots & \ddots & \ddots & & \\ & & & & -\beta & 1 & -\beta & \\ & & & & & -\beta & 1 & -\beta \\ & & & & & & -\beta & 1 \end{bmatrix} \tag{1.1}$$

To get Figures 4, we multiplied every single scan data vector with this matrix E, before back projecting (of course, does not require a full matrix×vector multiplication).

This modification of (1) preserves

- location and strength of images of point objects - they will appear less smeared - and
- linearity (meaning that a general image will be as good as is the treatment of point objects).

Choosing β= 0.3, 0.4, 0.5 and 0.6 respectively in (1.1) gives, with the scan data of our test object, the reconstructions that are shown in Figures 1.4. In spite of the filter being narrow (tridiagonal; wider filters could be 'tuned' better) and the optimization of the filter coefficient being crude (simply picking the best-looking of the four cases), we get excellent inversions.

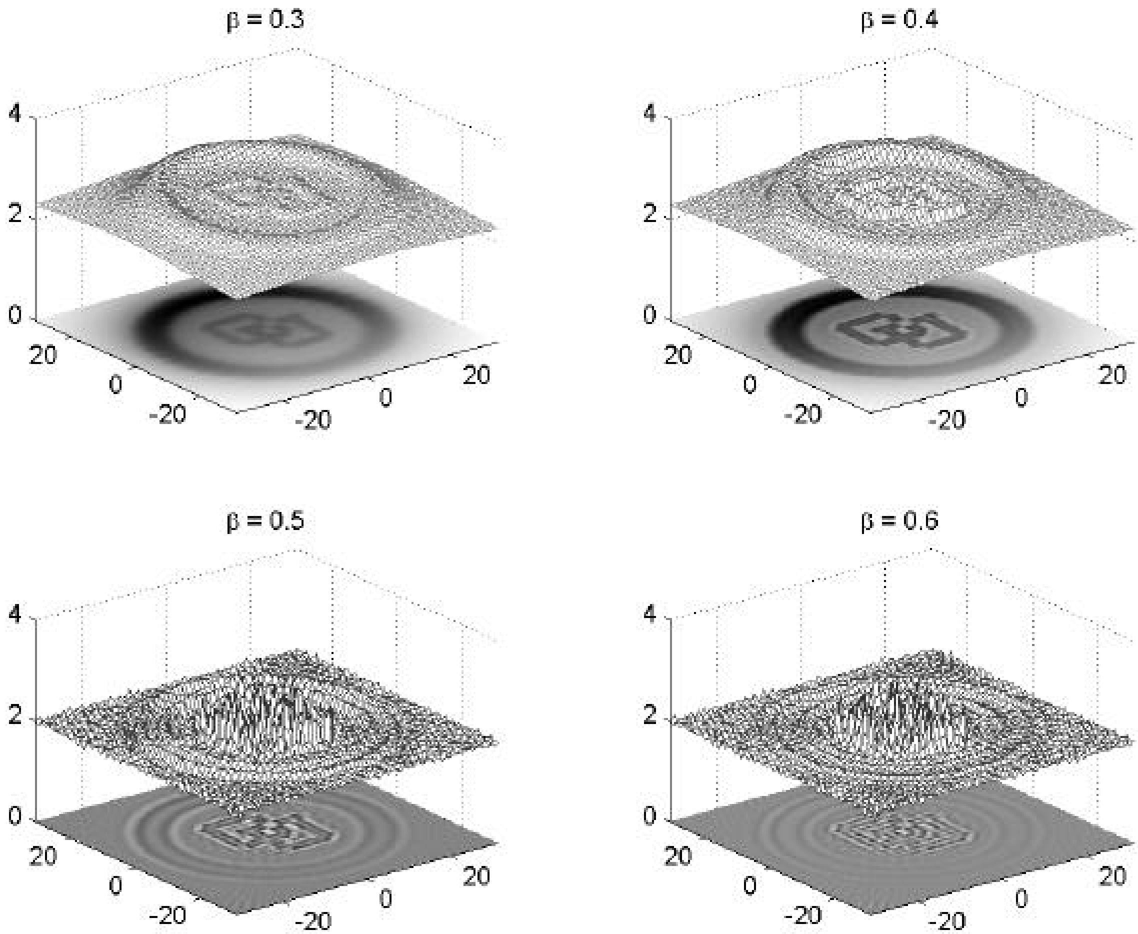

Figure 1.4: Filtered back projection with some choices of simple tridiagonal filters.