

**A SIMPLE ALGORITHM FOR LONGITUDINAL PHASE SPACE
TOMOGRAPHY**

S. Hancock

Abstract

Tomography is now a very broad topic with a wealth of different algorithms for the reconstruction of both qualitative and quantitative images. One of the simplest algorithms has been modified to take into account the non-linearity of large-amplitude synchrotron motion. This permits the accurate reconstruction of longitudinal phase space density from one-dimensional bunch profile data. The method may be further extended to treat, for example, multi-harmonic systems and self-fields.

Geneva, Switzerland
May 1997

1. Introduction

Strictly, tomography was born in 1917 when the Austrian mathematician J. Radon published a proof[1] that any two- (three-) dimensional object can be reconstructed from the infinite set of all its one- (two-) dimensional projections. This result has been repeatedly rediscovered. Indeed, the 1979 Nobel Prize for Medicine was awarded to the English computer engineer Sir G.N. Hounsfield for his pioneering contribution to computer-assisted tomography, which is employed in the body scanners of modern hospitals.

The underlying principle of tomography is to combine the information in a sufficiently large number of profiles to be able to reconstruct unambiguously the fuller picture with the extra dimension reinstated. Thus, for example, many one-dimensional profiles of x-ray transparency taken from different angles can give doctors an image of a two-dimensional slice through a patient. The ingenuity of the mathematical algorithms[2] employed derives from a need for reconstructions which are simultaneously stable and unique despite a necessarily finite set of data profiles.

The application of tomography to longitudinal phase space becomes obvious once it is realised that a bunch performing synchrotron motion is analogous to a patient rotating in a stationary body scanner. On each turn around the machine, a longitudinal pick-up provides a snapshot of the bunch taken from a slightly different angle. It only remains to combine such profiles tomographically to obtain a two-dimensional picture of phase space density[3].

2. The ART Algorithm[4]

The key to reconstruction is back projection. This is the process by which the contents of the bins of a one-dimensional histogram are redistributed over the two-dimensional array of cells which comprise the reconstructed image. Given no *a priori* knowledge of the original two-dimensional distribution, the contents of a bin is shared by all the cells that could have contributed to that bin. This is shown in Figure 1 for a small number of cells and an oblique projection angle. The thick line indicates the centre of one particular bin. The length of the segment of this line that lies inside a cell is used to scale how much of the bin contents is back projected into that cell. These coefficients are depicted by a grey-scale in the second part of the figure. The result is a jagged “track” of shaded cells whose total contents is that of the original bin. Note that the coefficients obtained in this way are only approximate, so, for example, nothing gets into the first cell shown even though the bin clearly impinges upon it. However, the calculation of all the areas of overlap with the full bin width (thin lines) would be much more cumbersome.

The back projection of all bins of all projections yields a first approximation to the two-dimensional distribution that produced those projections. The so-called Algebraic Reconstruction Technique (ART) is an iterative algorithm which exploits the fact that the coefficients for back projection can also be used to project the contents of cells into bins. Hence a new set of profiles can be obtained from the approximate distribution. Back projection of the difference between the original set of profiles and this latest one yields an improved reconstructed image. The iterations converge more rapidly if any cell whose contents becomes negative upon back projection of the difference profiles is reset to zero.

Here, it has been tacitly assumed that the different projection angles of the profiles correspond to a rigid, circular rotation of the two-dimensional distribution. For only then can the centreline of a back projected bin be drawn straight and the necessary coefficients be readily recalculated at each iteration.




Figure 1: The geometry of back projection.

3. A Modified Algorithm

Instead of back projection, consider how the contents of a cell gets divided among the bins of a particular profile. In general, one cell will contribute to only a couple of bins. By launching a number of test particles which, initially, are uniformly distributed within the cell, the estimation of coefficients becomes a simple matter of counting how many particles end up in each column of cells at the particular instant when the profile was measured. Provided that the trajectories of the test particles can be followed as a function of time, a set of maps can thus be built giving the redistribution of cells into bins for all projections. Furthermore, one can extract from those maps the information of the reverse transformation, namely the redistribution of bins into cells - i.e., back projection. The point of all this is that the trajectories of the test particles need not be assumed circular. Also, since the origin in time at which the test particles are launched can be chosen freely within the measurement interval, the distribution reconstructed can be at any instant during that period.

Although iteration proceeds as before, there is a price to be paid: a large map of coefficients must first be built and its inverse derived for every profile in the set of measured data. On the other hand, since most of the computational effort is invested in building the maps, it becomes trivial to repeat the calculation with fresh data taken under the same conditions.

4. Some Mathematica Calculations

The ideas of the previous section have been realized and tested using Mathematica. For simplicity, a straightforward Runge-Kutta integration of the steady-state equations of motion was performed in order to construct the maps from the trajectories of 64 test particles, which were launched in an array of 128×128 cells. The machine parameters chosen were those of the PS immediately prior to debunching for slow extraction. The same conditions were applied in a simulation program to generate 36 profiles of 128 bins at intervals of 17 turns by tracking 50,000 particles. It should be emphasised that this code incorporated an entirely different tracking method than the one used when building the maps. The initial, completely unmatched distribution is shown in Figure 2. Note that, even with 50,000 particles, there are "hot spots" due to statistical fluctuations. The width of both bins and cells is 0.5 ns.

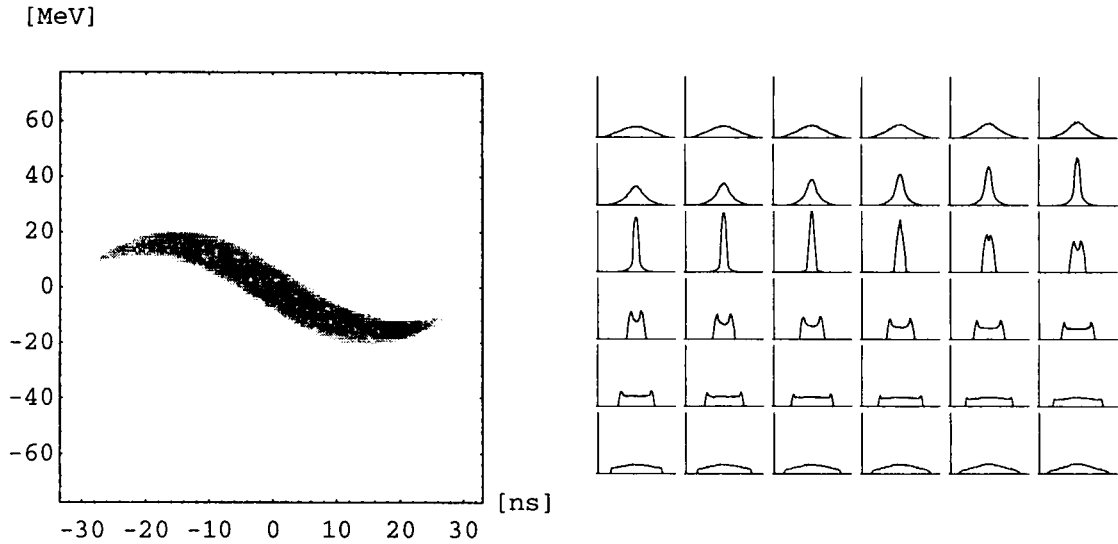


Figure 2: The initial phase space distribution and thirty-six of its projections at intervals of 5° of synchrotron phase. The latter constitute the set of “measured” data.

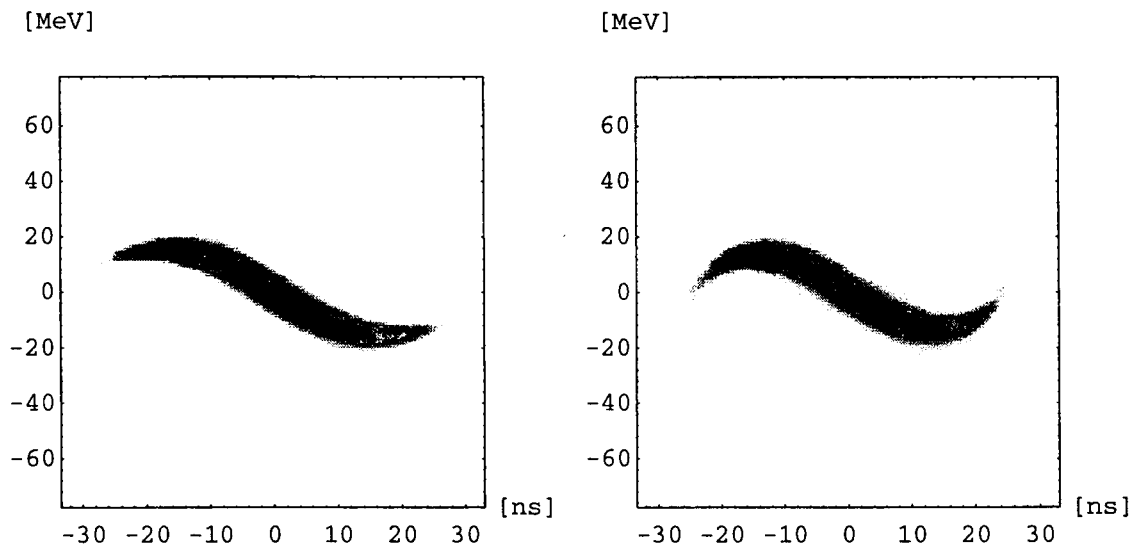


Figure 3: Reconstructed distributions assuming (i) non-linear and (ii) linear synchrotron motion.

Figure 3(i) is the picture after ten iterations of the new algorithm. Here, convergence has been accelerated by multiplying the back projected difference profiles by a factor which was increased at each successive iteration. Naively, one might expect that this weighting factor should remain fixed (at, in this case, $1/36$), but, as the reconstruction improves and the mutual interference between the corrections that are back projected diminishes, it may be steadily increased. Although this is dissatisfyingly empirical, the result is remarkably good. Indeed, there is evidence that the projections of the final distribution have started to follow the statistical fluctuations in the original profiles.

The method is very robust. Trigger jitter was simulated by shifting each profile by a random amount in time. A jitter of 1 ns degraded only slightly the reconstructed image. This might, anyway, be compensated by increasing the number of profiles. Similarly, by making the machine parameters used to generate the profiles different from those used to build the maps, plausible measurement errors were found to have only minor effect.

Figure 3(ii) is an indication of what the original algorithm would produce given the same profile data. However, the reconstruction was still achieved using maps, but these were first rebuilt on the basis of purely linear synchrotron motion. Otherwise, the data treatment was identical in both cases. The two images may therefore be compared directly since the question of the precision of the maps versus that of coefficients calculated from line segments does not arise. The significance of the non-linearity is evident.

The reconstructions of Figure 3 each took more than a day of computation on a 133 MHz Pentium Pro PC. However, the objective was merely to establish a proof of principle; Mathematica could never be considered for processing data in any practical implementation of the method.

5. Discussion

Obviously, the new algorithm has to be translated into a language that will run appreciably faster. Then the number of test particles can be increased to investigate the influence of the precision of the maps. If it proves prohibitively slow to build maps using many test particles per cell, one might consider launching just a few test particles (less than one per cell) from along the centreline of a column of cells and following these backwards in time. The line integral of a spline fit to the resulting points would yield the length of the segment inside each cell - cf. the original ART algorithm.

Some improvement in execution time might also be expected by suitably modifying the first approximation before iterating.

If full-blown tracking were used to follow the motion of the test particles, then arbitrarily complex RF systems could be considered. Even debunching in the absence of an RF voltage could be treated. Likewise, with due attention to normalization, particle densities outside the bucket could be reconstructed. Finally, a knowledge of the machine impedance and of the pick-up calibration would permit self-fields to be included in the analysis. Thus the shape of the measured profiles could be used to take into account the inductive wall effect.

6. Conclusions

An algorithm has been developed for longitudinal phase space tomography in which the contents of the reconstructed phase space array is effectively rotated instead of inclining profile bins in order to make a projection. This allows a different mapping to be applied to each cell in the array so that rigid, circular motion of the phase space distribution need not be assumed.

Simulated data have shown the method to be both accurate and robust. A measurement would require a "mountain range" of digitized bunch profile data. Modern storage oscilloscopes with fast sampling and segmentable memory with short trigger re-arm deadtime make this entirely feasible.

Acknowledgement

I am grateful to Erk Jensen for convincing me to abandon the notion of synchrotron frequency and to adopt a "brute force" approach when following the motion of test particles in my early attempts to build maps.

References

- [1] J. Radon, "On the determination of functions from their integrals along certain manifolds", Ber. Verh. Sächs. Akad. Wiss. Leipzig, Math.-Phys. Kl. 69, pp262-277 (1917).
- [2] http://wwwmath.uni-muenster.de/math/num/inst/natterer/publ/pap/natterer_4/12html/AlgoTom.html
- [3] G. Jackson, "A phase space tomography (PST) monitor for adjusting bunch rotation during coalescing", FN-469, FNAL (1987).
- [4] R. Gordon and G.T. Herman, "Three-dimensional reconstruction from projections: a review of algorithms", Int. Rev. Cytol. 38, pp111-151 (1974).