

TOMOGRAPHIC MEASUREMENTS OF LONGITUDINAL PHASE SPACE DENSITY

S. Hancock, P. Knaus, M. Lindroos
CERN, Geneva, Switzerland

Abstract

Tomography is now a very broad topic with a wealth of 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 is a hybrid one which incorporates particle tracking, and considerable effort has been invested to optimize the computer code so that it may also be compiled to exploit parallel architectures efficiently. A selection of the results obtained at different CERN accelerators is presented. The starting-point in each case is simply a “mountain range” of digitized bunch profiles.

1 INTRODUCTION

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 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 projected at a slightly different angle. It only remains to combine such profiles tomographically to obtain a two-dimensional picture of phase space density[1].

2 RECONSTRUCTION

Back projection is a key 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 one bin is shared over all the cells that could have contributed to that bin. The back projection of all bins of all profiles yields a first approximation to the original distribution.

Algebraic Reconstruction Techniques[2] (ART algorithms) exploit the fact that the coefficients for

sharing bins in back projection can also be used to project the contents of cells into those bins. Hence a set of projections can be obtained from the first approximation. Back projection of the bin-by-bin difference between the original set of profiles and this new one yields an improved approximation. Further iterations converge more rapidly if any cell whose contents has become negative is reset to zero.

The problem with conventional ART is that its strategies for estimating the redistribution coefficients are based on straight-line back projection geometry. This implies either rigid, circular motion of the two-dimensional distribution or that its projections be measured simultaneously. An alternative approach is to consider how the contents of one cell gets projected into the bins of a particular profile. By launching a small number of test particles which, initially, are uniformly distributed within the cell, the calculation of coefficients becomes a simple matter of counting how many particles end up in each bin at the particular instant when the profile was measured. Thus, a hybrid algorithm which combines particle tracking with ART allows large-amplitude synchrotron motion to be taken into account since the trajectories of the test particles need not be assumed circular. Although iteration proceeds as before, there is a price to be paid: a large map of (projection) coefficients must first be built and its inverse (for back projection) 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 a calculation with fresh data taken under the same conditions.

3 SOME RESULTS

The first part of Fig. 1 shows a simulated distribution comprising an unmatched bunch. These particles were tracked over an entire synchrotron period to obtain 72 profiles as the distribution filamented. Taking these profiles as the input dataset, the new algorithm yields the accurate reconstruction of the second part of Fig. 1.

The method is very robust. Since the iterations converge to the consensus of the information in the profiles, trigger jitter and noise in the data have only a minor effect. In addition, since the test particles are only tracked for a relatively brief period to build the maps, moderate errors in the machine parameters used in the tracking model are not serious.

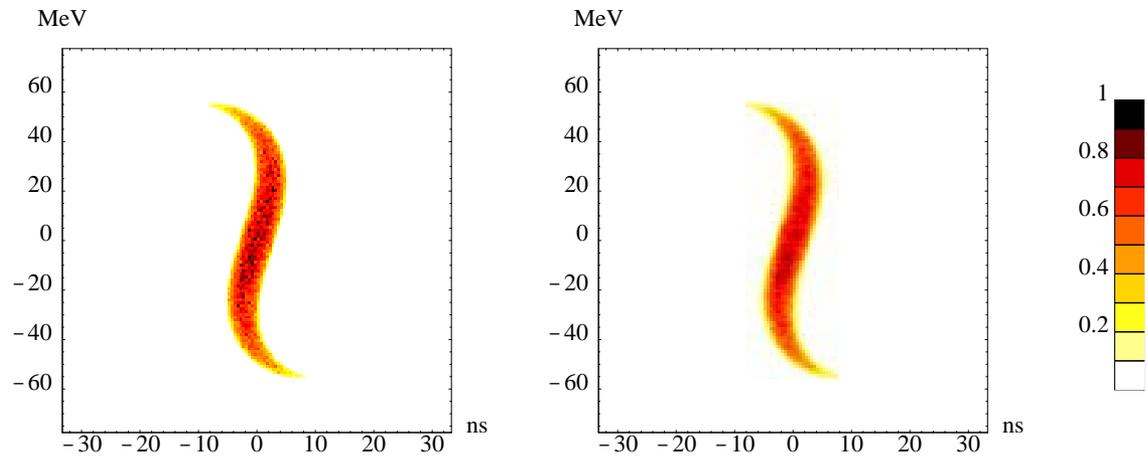


Figure 1: Simulated proton distribution and its reconstruction (on the same density scale).

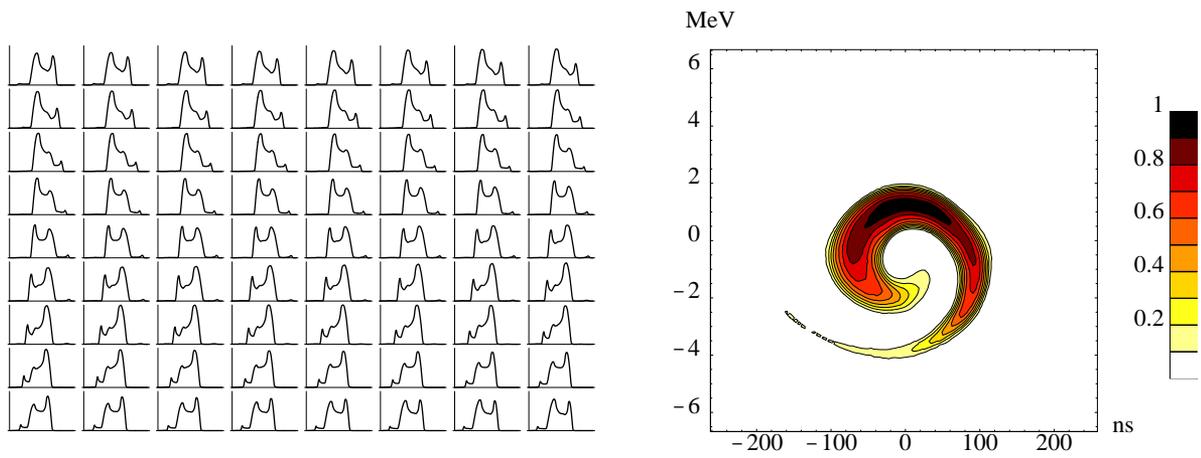


Figure 2: Measured profiles and reconstructed contour plot of a phase-modulated bunch at 1 GeV.

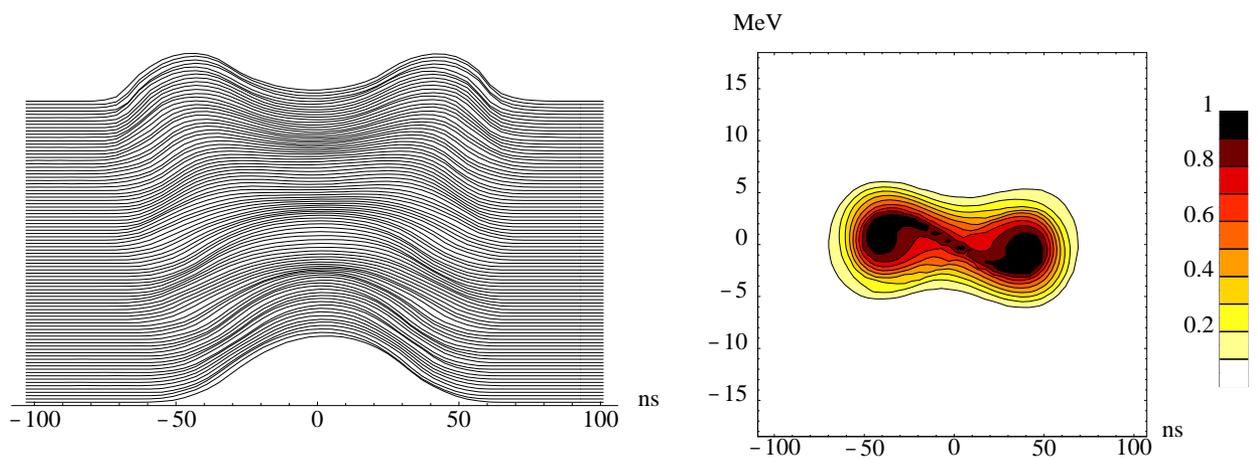


Figure 3: Quasi-adiabatic bunch splitting at 3.5 GeV/c.

The 72 profiles of Fig. 2 were measured every 48 turns at the CERN Booster during attempts to move empty phase space into the centre of a normal bunch by phase modulation of the rf. Without tomography, it would be difficult to deduce anything from the bunch shape information.

Arbitrarily complex rf systems can be treated. The “mountain range” plot of Fig. 3 shows a series of measurements every 50 turns during bunch splitting[3] in the CERN PS. Splitting is achieved by increasing the second-harmonic voltage component of a dual-harmonic rf system while reducing the fundamental. The data span 10 ms and the details of the two voltage programmes during this period were incorporated in the tracking model of the algorithm to produce the corresponding reconstructed image. Since the origin in time at which the test particles are launched can be chosen freely within the measurement interval, the reconstruction can be at any instant during that period. In Fig. 3, the reconstruction is at the time of the 72nd profile (out of 92). An animation of the entire process can be made by repeating the calculation for all time slices.

Discrepancy[2] expresses in a single figure of merit the residual differences between the projections of a reconstructed distribution and the original profiles. As such it may be used to monitor convergence. It is an obvious conjecture that, in the limit of many iterations, the lowest discrepancy should be obtained when the parameters used in the tracking model match most closely the reality of the machine. Reconstruction of the simulated data of Fig. 1 assuming different rf voltages supports this idea. The discrepancy after 100 iterations exhibits a sharp minimum at precisely the value of the rf voltage that was used to generate the data in the first place (see Fig. 4). Arguably, this constitutes a voltage measurement.

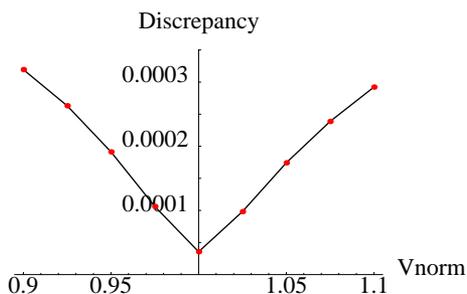


Figure 4: Discrepancy versus normalized rf voltage.

4 COMPUTER CODE

The algorithm was originally developed in Mathematica™, a choice mainly motivated by the rich set of built-in functions for graphics and for the manipulation of arrays. However, the objective was merely to establish a proof of principle. Consequently, with the aim of reducing the execution time by a factor of at least 100,

the code was translated into Fortran 90. This was done with the view ultimately to use parallel architectures.

In the original Mathematica™ code there is extensive use of floating-point operations and many very large and very sparse matrices. In the Fortran 90 version, the number of floating-point operations is reduced by the use of integer representation until the final step in each number manipulation, while a pointer-like re-allocatable vector representation was created to deal with the sparse matrices. The built-in pointer facility was avoided because it would make the use of parallel architectures with shared memory impossible.

The execution time of the present version of the code is some 400 times faster than that of the Mathematica™ one. This enables the production of well-resolved images within minutes, which is entirely adequate for machine development work. It also makes possible the production of animated sequences of dynamic processes within reasonable execution times. Extensive work is now in progress to further optimize the code for faster processors and parallel architectures. The long-term goal is to integrate the measurement as an operational tool within the controls system of the PS Complex.

5 CONCLUSIONS

An algorithm has been developed for longitudinal phase space tomography in which the contents of the reconstructed 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.

Leptons are excluded due to synchrotron radiation, but simulated proton data have shown the method to be both accurate and robust. Valuable tomographic measurements have been made during machine experiments.

The algorithm is a hybrid one and, consequently, arbitrarily complex rf systems can be catered for by modifying a small part of the code.

6 ACKNOWLEDGEMENTS

The authors wish to thank Mike Metcalf for helpful discussions and advice on programming in Fortran 90.

REFERENCES

- [1] S. Hancock, “A simple algorithm for longitudinal phase space tomography”, CERN/PS/RF/Note 97-06, 1997.
- [2] R. Gordon, “A tutorial on ART”, IEEE Trans. Nucl. Sci., NS-21, pp78-93, 1974.
- [3] A. Blas, R. Cappi, R. Garoby, S. Hancock, K. Schindl, J-L. Vallet, “Beams in the CERN PS Complex after the rf upgrades for LHC”, these Proceedings.