



Constructing an Error Estimate from Finite Difference Time Domain Calculations

J. Smith*

December 28, 2008

Abstract

Numerical calculation of Maxwell's equations will give results where the smallest features that can be resolved are normally limited by the resolution of the grid, or mesh, on which the equations are solved. A method to allow meaningful error estimates to be made in 3D and 2.5D finite difference time domain solvers is developed. This report introduces a straightforward prescription which can be applied to large data sets, outlining the free parameters and appropriate choices for wakefield calculations of tapering structures and illustrates how this can be used to assign an estimate of the residual uncertainty in the calculation.

*Cockcroft Institute, Daresbury, Cheshire, UK

1 Introduction

This memorandum is concerned with tapering structures, such as collimators, or transitions to the inside of an undulator, which are typically long in terms of the bunch length. These calculations are important in a number of linear accelerator projects. Finite Difference Time Domain (FDTD) calculations are deterministic, but the discretization means the solution will be systematically slightly different from the problem under study. As the resolution is increased, the system modelled numerically becomes closer to that for which we are ultimately trying to determine the characteristics, however this may not happen in a completely linear way, unless the chosen mesh size is a submultiple of the modelled structure.

2 Resolution studies in 2.5D

Quasi-2D, or 2.5D calculations can be performed when we assume cylindrical symmetry to the geometry. Calculations of the wakefields of tapering structures frequently employ FDTD methods, based on the solution of Maxwell's equations on the Yee cell[1]. Two commonly used tools used for this purpose are ABCI[2] and ECHO[3], and while MAFIA[4] is also capable of performing such calculations a moving mesh feature is present in both ABCI and ECHO, which makes them preferable tools with which to perform the calculation. Furthermore, both contain simple switches to use the Napoly[5] integration contour, which helps avoid problems caused by the 'catch up' problem. The catch up problem is described in detail by Wilson[6].

We wish to avoid numerical dispersion, where errors are caused by the speed of signals being different in different directions. This effect is smallest when the cells have the same grid spacing both axially and radially. Studies for collimator geometries were previously carried out, [7] [8] (Beard and Smith/Beard and Jones) using MAFIA, in which skewed aspect ratios were necessary to reduce the number of cells such that calculation was feasible with limited computing resources available at the time. As no investigation to establish the importance of numerical dispersion was possible, the results of [7] [8] will not be considered further herein.

The scale of the wakefield problem in a tapering structure is normally determined by the bunch length, and with analytic formulae generally expressing wakefields in terms of a parameter which combines this with a taper angle and gap size, the RMS bunch length is the most sensible datum. It contains the smallest scale features which need to be resolved in the numerical calculation. Failure to resolve this assumed Gaussian charge distribution in a FDTD method will produce errors, with non-physical frequencies in the bunch can be excited in the system due to the limited number of cells present in such a description. In the 2D situation the calculations are relatively swift, facilitating a thorough mapping of the parameter space, which should be useful to anyone attempting similar calculations in future.

Many users of FDTD software use their experience to judge when a calculation is being performed with sufficient resolution that the results are correct. The intention here is to




Figure 1: Monopole Longitudinal Wake against dz (left), and against dz^2 (right)

describe a technique which can be applied to aid this sense, and give an inexperienced user more confidence.

2.1 ABCI calculations

The ABCI calculations that have been performed are intended to test thoroughly whether we have a more robust analytical model to calculate the wakefields. Our interest was in both monopole and dipole longitudinal and transverse wakes. For these both the maxima and minima are of interest, both magnitudes and locations relative to the centre of the bunch. Also the integrated maximum for these quantities is of interest.

There are normalisation issues with the longitudinal dipole fields in ABCI 12.5. Users should ensure they are using the correct radius for both drive and witness charge in the .out file, as it is possible the beam pipe radius will be used instead.

Most users of ABCI will at some point have performed regression studies, for example this study of the transverse impedance of axially symmetric tapered structures by Podobedov[9]. In this case there was a linear dependence of the physical quantity under investigation, the transverse impedance, with mesh size.

Figure 1 shows calculated values for the longitudinal monopole wake of a collimator, as a function of cell spacing, dz , and cell area, dz^2 . To parameterize the behaviour of the calculations such that a reliable extrapolation is possible to arbitrarily small mesh size (“infinitesimal mesh size”), both linear and quadratic functional forms are presented. A significantly better description is obtained when cell area, dz^2 , as opposed to 1D cell spacing, dz , is used as the ordinate.

Also one would appear to be justified in neglecting those points with $dz^2 > 2 \times 10^{-8}$, since these appear to be following a different trend. This implies that for this sort of calculation there is no practical benefit of running at a resolution with less than 16 cells/sigma. Further credibility is given to this statement if we look at what is happening with other parameters, for example the integrated longitudinal dipole wake, as shown in Figure 2. We would hope that this would tend to zero, but again it is clear that the four lowest resolution points in this study do not follow the trend of the others. Although




Figure 2: Dipole Longitudinal Integrated Wake (left), Transverse Dipole Integrated Wake (right)

it would be possible to parameterize the full range of this distribution with a functional form, it is more robust to restrict the fitting interval to the higher resolution mesh points and continue with the minimal number of free parameters.

This also indicates something about the uncertainty in the calculation. The difference between where the trend line goes and zero, where it would go if the calculation was perfect gives a guide. It is not obvious that putting a single straight line through the transverse wake is the right thing to do. This is the most useful parameter calculated by the numerical software, as ultimately it tells us about the jitter amplification and emittance growth as a beam traverses the device. Here we are more confident if we throw away the low resolution points and use dz as the ordinate. This report concerns itself with establishing a procedure that can be used without such manual intervention to determine which points should be excluded. The author suggests an appropriate method is to perform the fit, weighting the points based on the cell size, such that the lower resolution results have reduced contribution to the parametrised behaviour used to extrapolate to infinitesimal mesh size.

The error could be calculated in one of a number of ways. Podobedov[9] ran a loop, halving the cell spacing until there were two results within the required accuracy (5%) of each other. For the slowly varying, tapered collimator geometries of interest, this would produce simulations that were impractically large. Effectively, Podobedov is getting an estimate from the difference between the highest resolution result calculated by the numerical software, and the point. This error is dependent on the technique used to do the fit.

Consider one particular bunch length, the one at the bottom of figure 3. To calculate the transverse dipole wake, we perform simulations with increasing resolution. We want to know when our results have converged. Table 1 gives the value of the 'infinite resolution point'. the result of the convergence study, where similar calculations have been performed of the transverse dipole integrated wake with 10, 12, 14, etc. cells per sigma, up to the number displayed in the first column in the table. It shows us the relative benefit of performing extra calculations to increase the knowledge that can be fed into

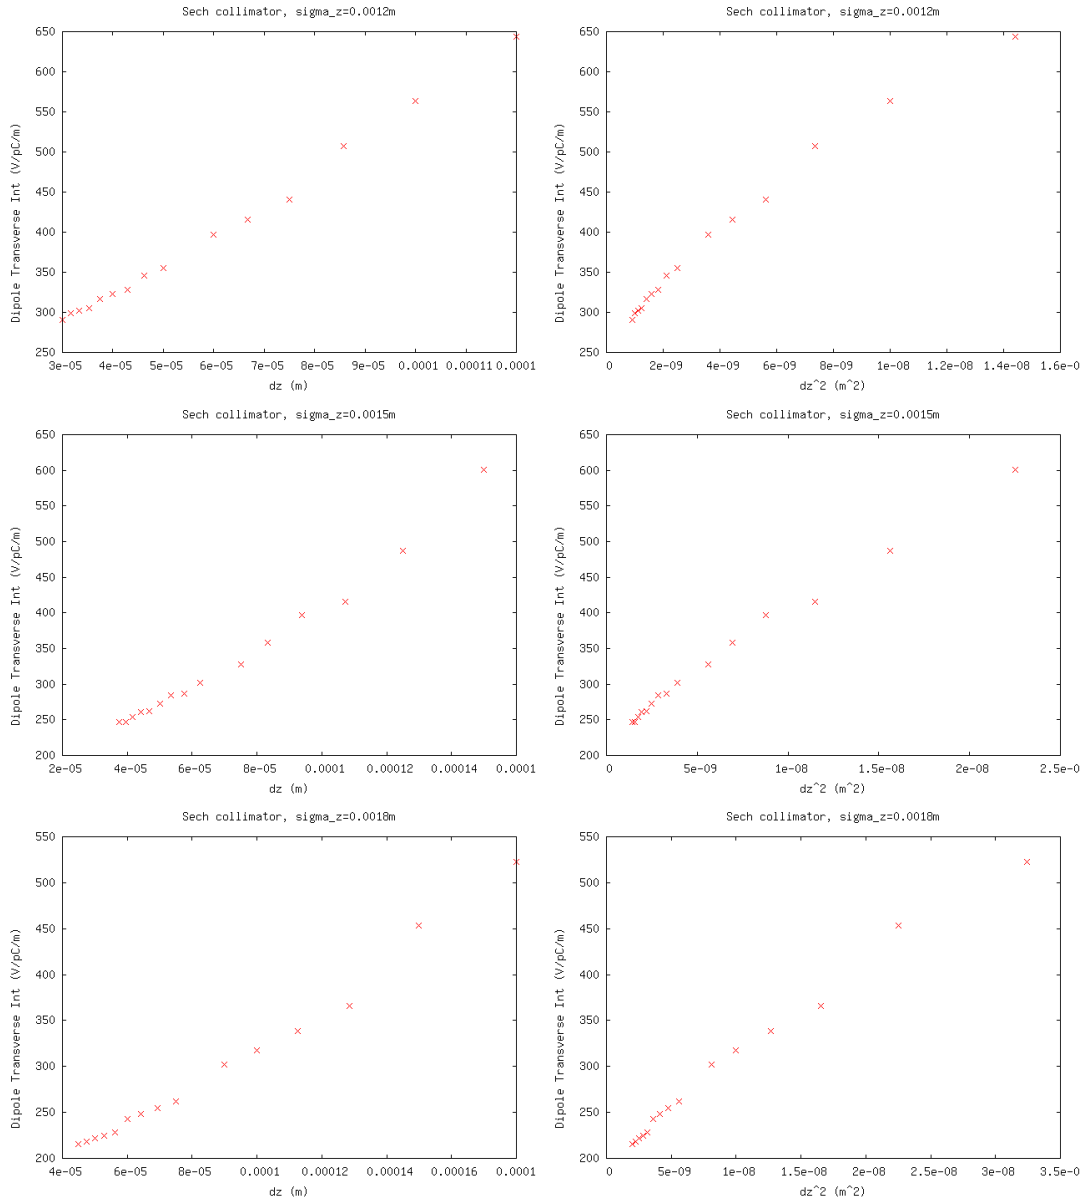


Figure 3: Transverse Dipole Integrated Wake against dz (left) and dz^2 (right), differing bunch length

Table 1: Convergence study using ABCI with sech collimator

Max. resolution	Linear fit (V/pC/m)	dz weighted (V/pC/m)	dz^2 weighted (V/pC/m)
32	91.1	105.1	116.8
34	93.1	105.1	113.1
36	95.9	107.5	115.1
38	98.9	110.6	119.0
40	101.8	113.9	123.4

Table 2: Convergence study using ECHO-2D with sech collimator

Resolution (cells/sigma)	Transverse wake (V/pC/m)
5	149.7
10	147.0
15	146.0
20	145.7

such a resolution study.

Having identified that low resolution points harm the convergence study, these could be excluded. Were the results from calculations with fewer than 20 cells/sigma removed as they would be if the task was being performed manually, with a linear fit the kick would be 140V/pC/m. This should be taken as the benchmark against which an automated procedure would be evaluated. Likewise trimming away calculations with fewer than 18 cell/sigma would give a converged value for the transverse dipole integrated wake of 127.37V/pC/m with the simple linear fit. We can see in table 1 that the linear fit with dz^2 weighted results are much closer. Furthermore, if we remove the 18 cell/sigma and lower resolution results, but use the dz^2 weighting to those points we still have, it gives a wake of 140.4V/pC/m. This weighting has saved us from the harm done by the lowest resolution point. In summary, it would appear prudent to apply this weighting whenever performing a 2.5D resolution study, and where possible not to include results with such low resolution that they taint the study.

2.2 Calculations with ECHO-2D

ECHO-2D features a conformal algorithm, improved from what has passed before. Even with a fairly coarse mesh it produces results with relatively little spread, so the resolution study required is much less involved. For the same collimator featured in the ABCI section above the following calculations were made, illustrated in table 2.2 for the integrated transverse dipole wake.

The lowest resolution result is again higher than the others, but by a much smaller amount. The calculated value of 145 compares favourably with the ABCI results in the

section before.

ECHO-2D does not run as a batch process, consequently large automated resolution studies are not possible. The results I have seen, including those with differing bunch lengths suggest that detailed resolution studies are relatively unnecessary.

3 Collimators, and Resolution Studies in 3D

Over the course of the EUROTeV project, the author of GdfidL developed a version of the software where the fields are only calculated in the vicinity of the bunch. This ‘window’ in which the fields are calculated can move at the speed of light along with the bunch. This assumes that the fields in front of the bunch are zero, and that any fields outside the window cannot exceed the speed of light, and reenter the window. We refer to this arrangement as moving-mesh or window-wake. In order to limit the numerical errors our initial choice would have been to base every calculations on a homogenous grid, however without moving mesh representation, there are a prohibitively large number of cell elements involved in such a calculation, and it would use more computational resource than we have available. Our calculations are for ultrarelativistic bunches of electrons, so the assumption that particles remain fixed in a co-moving frame is valid. While calculations have been performed both with and without the window-wake technique, we will focus largely on the window-wake enabled calculations, since this allows us a greater parameter space, and the fixed grid calculations were performed largely with older versions of GdfidL, in which bugs have been found.

Consider the bunch moving along the z-axis with a speed c , through a collimator designed to spoil the halo particles with high y excursions, and is there are no variations in x apart from the side walls. In calculations without moving mesh whole structure, memory forced us to reduce the resolution in the dimension with the most slowly changing field, in the same way that was described for the MAFIA calculations[7]. In our case this is the x dimension, and we have no choice other than to accept numerical dispersion and instability introduced by reducing this. The introduction of the windowwake method allows us to circumvent this as the computational requirements are drastically reduced, homogenous mesh becomes possible but for a completely homogenous mesh and the bunch lengths of interest to us ($\sim 300\text{--}500\mu\text{m}$) we are limited to a maximum resolution not much in excess approximately $50\mu\text{m}$, compared with global dimensions of $\approx 40\text{mm} \times 40\text{mm} \times 1\text{m}$. Given that we are interested in the wakefield response for various collimators, to various bunch lengths for bunches with a range of offsets, we already find ourselves in a multi-dimensional parameter space. Inclusion of simulations with varying resolution in order to ascertain convergence makes this a further parameter. We run completely through this parameter set, making use of the symmetries available to us. The minimum resolution considered was 5 cells/sigma, as it provides just enough sample points to adequately describe a Gaussian bunch. Initially, anything more than about 10 cells/sigma proved unstable, though later versions have overcome this to a degree. This report focusses more on the technique than the results.

3.1 Tools required

Large scale calculations place large demands on computational hardware. Simulations involving tens of cores at a time inevitably fail periodically. Hardware can fail, the network can become busy causing timeouts that prematurely end the job, and delays due to activities of other users can all cause jobs to fail. Sometimes this is due to failures of the PVM or MPI parallel environment, or of the queue manager. As a consequence a suite of tools was developed to monitor the running jobs, determine which had failed, propose changes to submission files and so on. Such tools are necessary for this scale of calculation. Manually postprocessing each of these jobs each is not a viable option. For the collimator project, Perl was used to create this suite, on account of its ease of use for text processing and pattern matching.

Automated scripts have been written in Perl:

- For submission of the jobs (so to sweep over the large parameter sets, with automatic estimation of job size, hence number of nodes).
- For filtering jobs to identify those which have successfully completed (e.g. reassuring message “The end of the simulation has been reached”).
- For performing secondary postprocessing, and extraction into a csv file readable by spreadsheet of the pertinent information to the job.
- For taking this spreadsheet and performing a fit to establish approximate error on runs of the same collimator, with the same bunch length, with same mesh aspect ratio, with the same offset, with the same jaw depth and determining a single, supposedly mesh independent value to be carried through the rest of the postprocessing.
- For taking these values for the loss factors, and turning this into an estimation of the kick factor, for comparison with analytic and experimental measurements.

Some further Perl and bash scripting was put in place in order to identify which jobs actually ran, how close to finishing other jobs were, enabling swift rerunning of those jobs which had not completed normally.

3.2 Resolution dependency

When one is attempting to obtain a mesh-independent result, there are a number of potential ways to fit the data, as discussed in the ABCI section above. We would aim to reduce the dependence of the final result on the (unphysical) arbitrary choice of mesh resolution. From general considerations, it is logical that numerical errors are reduced as one creates smaller and smaller cells until one reaches the limit where the cells disappear. Without a good reason to do anything else it would seem an unweighted linear fit would be the selection, however as resolution tends towards the infinite resolution point one frequently observes a convergence of the results to a given value, and a exponential fit

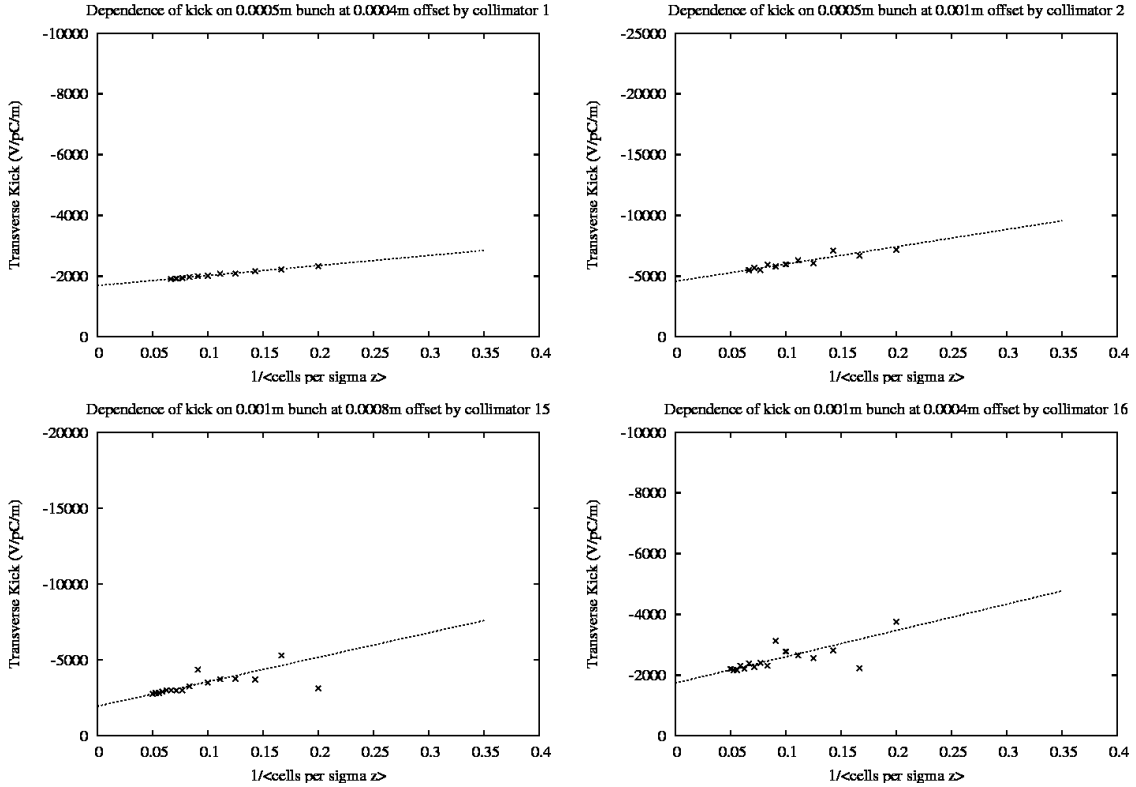


Figure 4: Transverse kicks for different collimators, with dz^3 weighting

looks more suitable for the kicks. Forms considered which do not provide adequate descriptions include cubic spline fitting and polynomial fitting particularly when the beam is close to the collimator jaws, for higher order terms not to be present, tainting the value of the infinite resolution point. The results I will present will consider just linear fits, with weighting of the points.

Also one can question whether all the results should carry equal weight, or whether those with higher resolution, which are expected to be more accurate, should then be given more weight than the other points. This can be achieved by introducing an effective uncertainty proportional to the size of the cells, so this is lower for the higher resolution models.

It is an assumption that the highest resolution point is the most accurate, and the difference between that and our result as we tend to infinite resolution gives a crude estimate of error. GdfidL does not provide an uncertainty on the calculated loss values, therefore attaching an error to our result is somewhat arbitrary, however it is possible to counter this to a limited extent when one is calculating the kick factor. From our list of successfully completed jobs created by the script above, we then run over these runs, find like runs with differing resolutions, and merge these into a single value for a given collimator, with given bunch length, offset, etc. Some examples of resolution studies are shown below:

In order to determine an appropriate weighting, a typical example resolution study is performed. This suggests enforcing larger errors on lower resolution points.

3.3 Kick calculations

GdfidL calculations consider a source which has a Gaussian charge profile longitudinally, and is a delta function in the transverse plane. Consequently we are not simply exciting a single mode, as we were in 2D case. Having determined a resolution independent result, we need to combine results for the same collimator with different offsets in order to find the dominant dipole kick, but also to determine the higher order modes present in calculations where the bunch has a large offset. This process is similar to that which was necessary when examining experimental data, such as that taken by the T480 project [10][11] at SLAC's End Station A [12]. We tested two different methods there, and I will describe those same methods here.

The kick is defined on the assumption that near the axis the beam will be deflected proportionally with the offset from the design axis, so the transverse kick

$$\kappa_T = \lim_{y \rightarrow 0} \left\langle \frac{y'}{y} \right\rangle$$

Numerically we could attempt to calculate this limit, though the closest we can get to this approximation is simply to take the point closest to the axis, that a 0.2mm, and calculate the deflection there, and turn that into a kick factor. This is somewhat wasteful of the other calculation results. If this single datapoint, formed from only about 5 or 6 GdfidL calculations with varying resolution, is not good, then the reported kick will likewise be poor. A second method is to take the points with offsets of 0.2, 0.4, 0.6mm and construct a linear fit. The difference between the infinite resolution point and highest resolution point described above can be used as a guide for the errors to be used for this fit.

We can then renormalise these so that the χ^2 of the fit is of the order of one. We can take the multiplication factor to give a slightly more credible estimate of the error in the GdfidL simulations. A similar mechanism allows us to perform a linear plus cubic fit on the data recorded at all offsets, treating errors in the same way. We prefer this method, as it utilises all the data, which can then be used to reduce the error and the linear term in such a fit tends to what it would be in the limit described above.

In practice the choice of using the central three points was preferred for the experimental data, as data with large offsets was particularly unreliable.

In the GdfidL calculations, a symmetry plane is enforced in the electrical centre of the device. Therefore a bunch travelling on the symmetry plane cannot experience any kick. While it would be possible to leave the kick with a drive charge on axis to be determined by a fit as a zero order term along with the linear and cubic terms, it makes more physical sense to enforce that this is zero. The constant term in the fits should therefore be removed. For comparison with the ESA data, we should really use the constant term found in that fit. This allows comparison of GdfidL results, in which

the electrical centre is known, with experimental results, where the electrical centre is dependent on precise alignment of components. The constant term in the experimental results tells us where that zero position is relative to the set positions of the mover that was used there.

That describes the post-processing steps required to turn raw GdfidL results into a meaningful comparison, including errors, with experimental data.

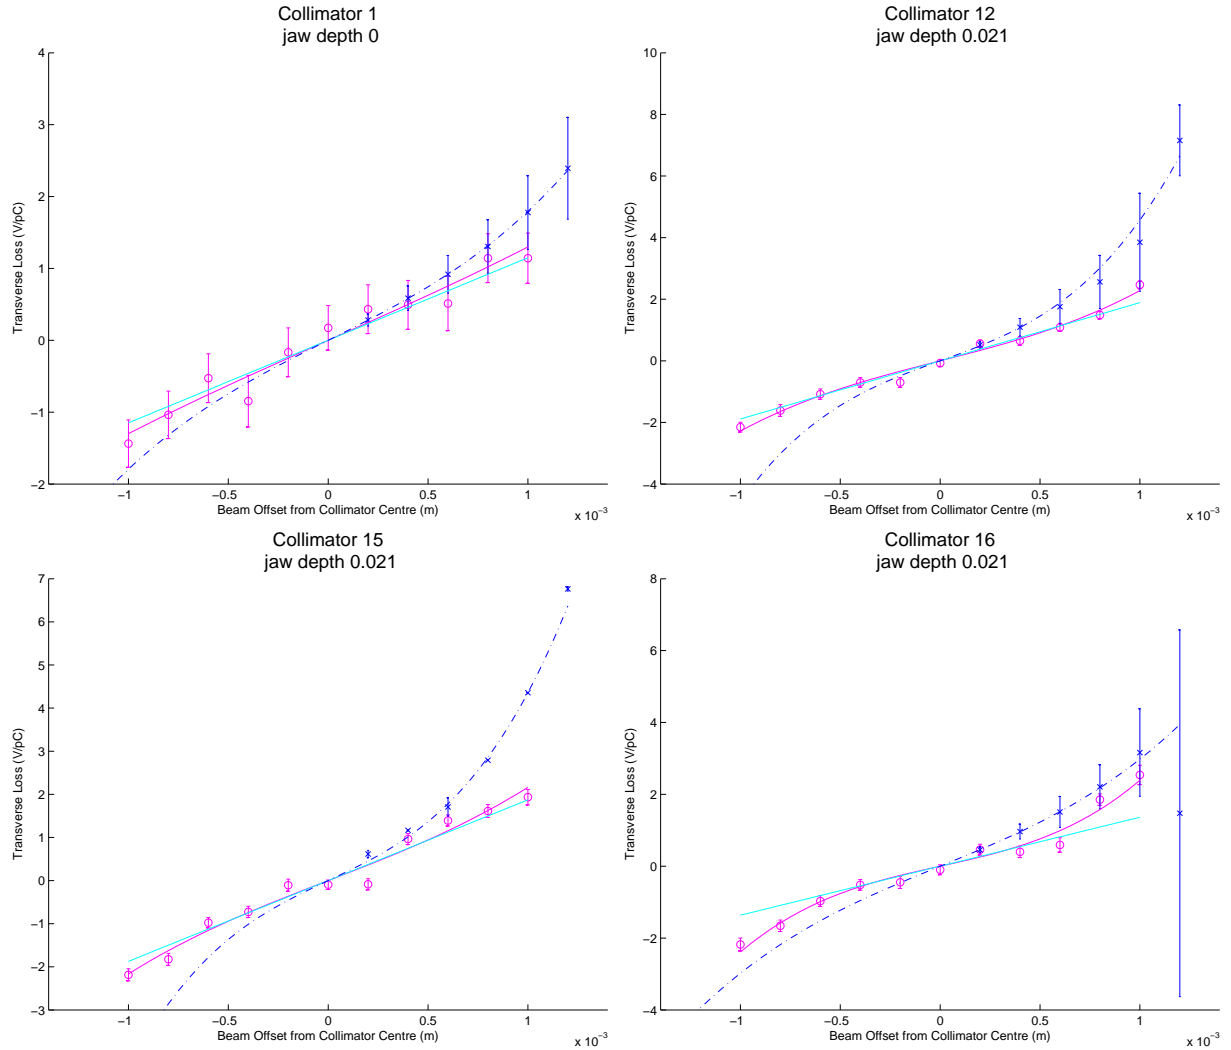


Figure 5: Experimental and calculated loss factors for collimator 1.

Some example calculations are shown in figure 5, but more data is still required for the final T480 calculations. On these plots, the pink points are preliminary experimental data [10][11], the blue curve represents numerical calculations with a $500\mu\text{m}$ bunch, and a fit which includes linear and cubic terms. This is probably shorter than the experimental data, which would make the numerical calculation have greater kick than the equivalent experimental result, however it only includes geometric wakefield kick, and not the resistive wakefield effects which will affect the experimental data. The

cyan line is a fit to the GdfidL data using the central three points, as was the preferred method for handling experimental results[11]. The agreement between experimental and centrally fitted GdfidL results is good.

4 Conclusions

A prescription to allow meaningful error estimates to be made in 3D and 2.5D FDTD solvers has been introduced, for the specific example of slowly varying, tapered collimators. This straightforward prescription can be applied to large data sets and the choice of parameterization of the functional form used to extrapolate calculations performed over a range of finite mesh resolutions to infinitesimal mesh resolution has been justified. This applicability of this procedure has been illustrated in comparison to preliminary data from the T480 testbeam project at SLAC, and will be used by that project to perform the calculations to accompany the experimental data in their final paper (in preparation).

Acknowledgement

This work is supported by the Commission of the European Communities under the 6th Framework Programme "Structuring the European Research Area", contract number RIDS-011899.

References

- [1] K.S. Yee. Numerical solution of initial boundary value problems involving maxwells equations in isotropic media. *IEEE Trans. Antennas Propagat.*, 14:302307, 1966.
- [2] Y. H. Chin, K. Takata, and Y. Shobuda. ABCI progresses and plans: Parallel computing and transverse Napoly-Shobuda integral. In *Proceedings of 2007 Particle Accelerator Conference, Albuquerque, New Mexico*, pages 3306–3308, 2007.
- [3] I. Zagorodnov and K.L.F. Bane. Wakefield calculations for 3D collimators. In *Proceedings of 2006 European Particle Accelerator Conference, Edinburgh, UK*, pages 2859–2861, 2006.
- [4] CST AG. [http://www.cst.com/content/products/MAFIA /overview.aspx](http://www.cst.com/content/products/MAFIA/overview.aspx).
- [5] O. Napoly, Y. H. Chin, and B. Zotter. A generalized method for calculating wake potentials. *Nuclear Instruments and Methods in Physics Research A*, 334:255–265, March 1993.
- [6] P. B. Wilson. Introduction to wakefields and wake potentials. Introductory SLAC-PUB-4547, Stanford Linear Accelerator Center, Stanford Linear Accelerator Center, Stanford University, Stanford, CA 94309, January 1989.

- [7] C.D. Beard and J. D. A. Smith. Numerical calculations of collimator insertions. In *Proceedings of 2006 European Particle Accelerator Conference, Edinburgh, UK*, pages 709–711, 2006.
- [8] C. Beard and R. M. Jones. Numerical simulations of collimator insertions using mafia. Technical Report EUROTeV-Report-2006-055, EUROTeV, 2006.
- [9] B. Podobedov and S. Krinsky. Transverse impedance of axially symmetric tapered structures. *Physical Review Special Topics - Accelerators and Beams*, 9(054401):054401–1–8, May 2006.
- [10] S. Molloy, Sergei Seletskiy, Mike Woods, Jonathan David Andrew Smith, Carl David Beard, Juan Luis Fernandez-Hernando, Nigel Watson, Adriana Bungau, and Andre Sopczak. Measurements of the transverse collimator wakefields due to varying collimator characteristics. In *Proceedings of 2007 Particle Accelerator Conference, Albuquerque, New Mexico*, pages 4207–4209, 2007.
- [11] J. L. Fernandez-Hernando, S. Molloy, J.D.A. Smith, and N. K. Watson. Measurements of collimator wakefields at End Station A. In *Proceedings of 2008 European Particle Accelerator Conference, Genoa, Italy*, pages 2868–2870, 2008.
- [12] M. Woods, C. Adolphsen, R. Arnold, G. Bowden, B. Bower, R. Erickson, T. Fieguth, J. Frisch, C. Hast, R. Iverson, Z. Li, T. Markiewicz, D. McCormick, S. Molloy, J. Nelson, M. Pivi, M. Ross, S. Seletskiy, A. Seryi, S. Smith, Z. Szalata, P. Tenenbaum, D. Angal-Kalinin, C. Beard, C. Densham, L. Fernandez-Hernando, G. Ellwood, J. Greenhalgh, F. Jackson, A. Kalinin, J. O'Dell, F. Zimmerman, J.H. Brownell, H.J. Schreiber, M. Viti, I. Zagorodnov, M. Wendt, V. Duginov, S. Kostromin, N. Morozov, S. Boogert, Y. Sugimoto, S. Walston, D. Burton, N. Shales, J. Smith, A. Sopczak, R. Tucker, R. Barlow, A. Bungau, A. Mercer, G. Kurevlev, M. Albrecht, M. Hildreth, W. Allison, V. Blackmore, P. Burrows, G. Christian, C. Clarke, G. Doucas, A. Hartin, M. Johnston, B. Ottewell, C. Perry, C. Swinson, G. White, S. Boogert, P. Huggard, W. Mueller, T. Weiland, D. Adey, M. Stockton, N. Watson, B. Chickering, O. Khainovski, Y. Kolomensky, T. Orimoto, C. Thu Hlaing, M. Slater, M. Thomson, D. Ward, M. Kimmitt, F. Gournaris, A. Liapine, B. Maiheu, S. Malton, D.J. Miller, M. Wing, N. Sinev, and E. Torrence. Test beam studies at slac's end station a and for the international linear collider. In *Proceedings of 2006 European Particle Accelerator Conference, Edinburgh, UK*, 2006. EUROTeV-Report-2006-060 SLAC-PUB-11988.