



Fermi National Accelerator Laboratory

FERMILAB-FN-675

Smoothing CASIM Results

A. Van Ginneken

*Fermi National Accelerator Laboratory
P.O. Box 500, Batavia, Illinois 60510*

November 1998

Disclaimer

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

Distribution

Approved for public release; further dissemination unlimited.

Copyright Notification

This manuscript has been authored by Universities Research Association, Inc. under contract No. DE-AC02-76CHO3000 with the U.S. Department of Energy. The United States Government and the publisher, by accepting the article for publication, acknowledges that the United States Government retains a nonexclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this manuscript, or allow others to do so, for United States Government Purposes.

Smoothing CASIM Results

A. Van Ginneken

Fermi National Accelerator Laboratory*
P. O. Box 500, Batavia, IL 60510

November 16, 1998

Abstract

A smoothing algorithm presented earlier is applied specifically to CASIM Monte Carlo outputs. Emphasis is on two applications: estimation of low level radiation dose subject to large statistical uncertainty and estimation of maximum energy deposition density along and near the beam axis.

1 Introduction

A smoothing algorithm—based on the Bernstein polynomials, which have a long history for use in such problems—is presented in [1]. In addition to the basic algorithm, which converts a (multidimensional) histogram to a smoothed histogram, ref. [1] mentions how boundary and other conditions may be used to tailor the algorithm to a specific application. This note presents a set of such specific procedures, encoded in the package ZMOOZ, which comprises a set of routines to smooth outputs of the Monte Carlo program CASIM [2]. No changes, other than the required CALL statement, are needed in CASIM to accommodate ZMOOZ. The basic algorithms of [1] appear to work quite well in a variety of problems, but the large dynamical range of CASIM output—as much as 20 orders of magnitude, with bin-to-bin variations of an order of magnitude or more—puts considerable strain on any smoothing procedure. As an illustration, ref. [1] includes an application to CASIM output but the procedure used there is not very well optimized and was never put in general use.

Two frequently encountered CASIM applications are addressed here: (1) large-scale shielding problems where dose must be estimated at large distances from where beam is lost or dumped, and (2) energy deposition problems where one needs to know the *maximum* of the energy density in the target and its distribution in that vicinity. In both applications it is assumed that cylindrical symmetry

*Fermi National Accelerator Laboratory is operated by Universities Research Association under contract with the US Department of Energy.

applies so that one deals with a 2-D problem. While this frequently does not apply in practice, it usually prevails at the computational level, where a problem which is close to cylindrical (typically rectangular) gets symmetrized or where a strongly asymmetric situation is approximated by a set of cylindrical ones, which together cover the azimuthal range of interest. The smoothing procedures include a ‘radial decline option’ which enforces—in minimalist fashion—the constraint that, at a given depth, dose or energy density declines monotonically with radius. The algorithm presented here is quite specifically geared to CASIM output. In particular, it deals with a 50×50 (z, r)-bins array, *equispaced* in both z and r . Reworking ZMOOZ to accommodate different bin arrays is straightforward. Extension to variable bin size might be more daunting since the Bernstein polynomials specifically assume equispacing [3].

A main concern about smoothing is that the end product will be taken more seriously than warranted [1]. This concern arises because smoothness may be mistaken for precision as is more or less true for (raw) Monte Carlo results [4]. Hence smoothing—it is feared—may become a substitute for CPU time. Nonetheless, some smoothing is almost always done when faced with statistically uncertain results. Whether it is just ‘eyeballing’ or incorporates more laborious methods, like fitting to analytic forms, it is difficult for such an exercise to remain free from prejudices which may arise when, e.g., a shielding problem poses large financial or other challenges. Such bias is likely reduced by using an automated smoothing routine. The aim of smoothing here is to obtain a better estimate of dose or energy deposition at a particular location, essentially by making use of results at nearby locations. It is *not* to produce an esthetically pleasing curve or surface. While some fluctuations may disappear others may just be reduced in size. Since the algorithm has no way to distinguish a genuine feature from a fluctuation in a noisy environment, there is a tendency to hedge between the two—not unlike an unbiased observer might do under these circumstances. However, sometimes the observer has the advantage of knowing with certainty that a particular excursion, e.g., one that occurs deep inside a homogeneous shield, is indeed a fluctuation. In such cases it may still remain for the user to perform some final smoothing. Usually, this final smoothing will be more straightforward and less prone to error and bias than one performed directly on raw CASIM outputs. Regardless of which type or level of smoothing is performed, *it is essential that the user consults the raw outputs and their associated errors in judging the reliability of the results.*

A brief description of CASIM outputs is in the next section, followed by the radial decline option after which smoothing of dose and energy deposition are discussed in turn. Some illustrations of smoothing typical outputs are then given and matters end with a few concluding remarks.

2 CASIM Outputs

The basic output of the Monte Carlo code CASIM consists of a 50×50 array of either star densities, i.e., densities of nuclear interactions above a threshold of 50 MeV, $\rho_S(z, r)$, or energy deposition densities, $\rho_E(z, r)$. Attached to each member of this array is an estimated relative error $\sigma(z, r)$ which is calculated by dividing the Monte Carlo run into 10 equal sub-runs which are then treated as 10 independent measurements [2]. When the error is small, below about 25%, $\sigma(z, r)$ can be interpreted as the usual error associated with a normal distribution but above this level it gradually loses this meaning. In particular, when all contributions are gathered during a single sub-run—and an error cannot be calculated properly—the computation returns $\sigma(z, r) = 1$. This is like rare-event-statistics where a single event is usually quoted as 1 ± 1 . However, in a weighted Monte Carlo, where individual contributions may differ by several orders of magnitude, the error analysis is more complicated. While a relative error of unity appropriately covers the entire downside (the result must be non-negative), to the up-side it is essentially an artifact of the error calculating procedure. In the smoothing, $\sigma(z, r)$ is used in the stop criteria (see below) of the iteration and in this context it becomes clear empirically that when $\sigma(z, r)$ approaches unity it must be enlarged to avoid long iterations which bring the smoothed values closer to these very uncertain results and thereby produce rather un-smooth output. The remedy adopted is to replace the relative standard deviation by $\sigma_{mod} = \sigma_{rel} / [1 - (1 - 1/\kappa)\sigma_{rel}]$ where $\kappa \geq 1$ is a constant to be specified. Here the choice $\kappa = 10$ is made which means that $\sigma_{mod} = 10$ when $\sigma_{rel} = 1$. This choice is quite arbitrary but the procedure is not very sensitive to the precise value of κ or to changes in the functional form of $\sigma_{mod}(\sigma_{rel})$ provided $\sigma_{mod} = \sigma_{rel}$ for $\sigma_{rel} \ll 1$ and σ_{mod} approaches roughly the same multiple of σ_{rel} when $\sigma_{rel} \rightarrow 1$. It appears best also to avoid dealing with very small σ_{mod} and ZMOOZ provides a low-error cut-off with default set, somewhat arbitrarily, at 1%. Typical *systematic* accuracy of CASIM results is well above 1%, so smoothing below that level is hard to justify even when statistical errors are much smaller. Also, the estimated errors are themselves subject to fluctuation so if, by chance, an abnormally low error is computed at some location, one avoids undue iteration trying to match the value to which this error attaches. Both κ and the low-error cut-off are at the disposal of the user.

When smoothing output of a large shielding configuration, problems may arise in connection with the column of bins on the beam axis. In such a problem, typically, beam size is much smaller than the radial extent of the bins resulting in large intra-bin variation in the on-axis bins. Averaging over bins typically leaves an on-axis bin much larger than its off-axis neighbors and hides the continuity which is present on a smaller scale. This complicates the smoothing and the simplest way around it is to replace—temporarily—the on-axis bins by radial extrapolations from its neighbors, except if, at some distance into the shield, the extrapolated

value exceeds the direct CASIM result in which case the latter is retained. (Eventually, after the 50×50 array is smoothed with the extrapolated results in place, the 50×1 array of inner bins is smoothed separately and replaces the corresponding bins in the 50×50 array.)

CASIM outputs for energy deposition also have a 50×50 array which predicts energy deposition as a function of location over the entire target. However since (1) on-axis intrabin variation is expected to be even larger here and (2) energy deposition along the axis—where the maximum energy density occurs—is of special interest in many applications, the innermost CASIM bins are further divided into 10 radial regions. For an accurate result on maximum energy deposition, radial dimensions should be specified so that beam size well exceeds radial size of the bins in the 50×10 array [5]. Both the 50×50 and the 50×10 array are processed by ZMOOZ.

3 Radial Decline Option

The main difference between smoothing as performed here *versus* that by an unbiased observer lies in the use of *a priori* information. Even without a specific model the observer may know certain qualitative features about the spatial distribution of dose or energy density which might be of use in smoothing. In many cases such features depend on details of the geometry and are not readily incorporated into a general smoothing algorithm. However, one such feature applies almost universally: dose or energy density at any given depth declines monotonically with radius. Other regularities which could be exploited, e.g., a single peak in the longitudinal dependence, would require considerably more finesse to introduce into an algorithm and exceptions would be far more often encountered. Eventually, inclusion of a lot of such information begins to resemble fitting—not smoothing. However, given the rather minimal nature of the monotonic radial constraint it can still be regarded as part of the smoothing and is included as an *option*. The radial constraint is enforced by redistributing the results—prior to any smoothing—in a manner that conserves total (dose or energy deposited) as well as $\langle r \rangle$, the first moment of the distribution, at a given depth. This redistribution is kept as local as possible by involving in it only the bin where the violation occurs and its two immediate neighbors. Enforcing the radial constraint is independent of the rest of the algorithm and could even be implemented without any smoothing.

Assume that at a given depth $\rho_n < \rho_{n+1}$, contrary to a monotonic decline. Densities of the n , $n + 1$, and $n + 2$ bins are then redistributed among themselves so that afterwards $\rho'_n = \rho'_{n+1}$ while total stars as well as the ρ -weighted mean of the radius, $\langle r \rangle$, over the three bins is unchanged—which specifies the problem completely. When applying these conservation conditions, the dose, etc., is multiplied by the corresponding bin volume to arrive at the proper weighting for that bin. This leads to the three equations

$$\begin{aligned}
\rho'_n &= \rho'_{n+1} \\
&= (2n-1)\rho'_n + (2n+1)\rho'_{n+1} + (2n+3)\rho'_{n+2} \\
&= (2n-1)\rho_n + (2n+1)\rho_{n+1} + (2n+3)\rho_{n+2} \\
&= [n^3 - (n-1)^3] \rho'_n + [(n+1)^3 - n^3] \rho'_{n+1} + [(n+2)^3 - (n+1)^3] \rho'_{n+2} \\
&= [n^3 - (n-1)^3] \rho_n + [(n+1)^3 - n^3] \rho_{n+1} + [(n+2)^3 - (n+1)^3] \rho_{n+2}. \quad (1)
\end{aligned}$$

The second equation conserves total dose in the three bins while the last enforces the equality

$$\begin{aligned}
\int_{n-1}^n r^2 \rho'_n dr &+ \int_n^{n+1} r^2 \rho'_{n+1} dr + \int_{n+1}^{n+2} r^2 \rho'_{n+2} dr \\
&= \int_{n-1}^n r^2 \rho_n dr + \int_n^{n+1} r^2 \rho_{n+1} dr + \int_{n+1}^{n+2} r^2 \rho_{n+2} dr \quad (2)
\end{aligned}$$

which keeps $\langle r \rangle$ the same. Depending on the statistical accuracy of the original results several passes may be required before the required monotonicity is achieved.

The result flattens out any bumps in the radial distribution. Subsequent smoothing then produces a decline through these plateaus though perhaps less than might be surmised from the global distribution. The monotonic option is especially beneficial at large depths and small radii where one expects the radial distribution to show a broad maximum centered on the beam axis and this result follows quite naturally from exercising the option. This is an important region which may determine, e.g., how long a shield must be, but which often has very poor statistics—especially in the bins close to the axis which have relatively small volume.

4 Low Dose

A typical shielding problem involves specifying the outer radius of concrete or soil surrounding a (potential) beam loss point so that maximum dose outside the shield remains below a given limit. For such cases smoothing is performed directly on the Monte Carlo results—not on their cumulative distribution, as might be desirable [1] if one were not faced with such a large range of values. Unless the shielding is completely homogeneous, the procedure starts with converting star density (stars/cm³) in each bin into dose equivalent (in rem) by multiplying the star density by the ratio of the mean-free paths of the bin material to that of soil and then applying a conversion factor (soil star density to rem) [6]. Unlike the star density, the dose thus obtained is a *continuous* function of location which is essential for smoothing to be applicable. In plain CASIM discontinuities are always present at the boundary of a void where ρ abruptly goes to zero. To avoid $0 \times \infty$ problems in converting to dose, voids must be replaced with a very low density

Table 1: Patches and Output Sets for Low Dose Smoothing

region	Patch z or r	Output z or r
1	1–20	1–15
2	11–30	6–25
3	21–40	16–35
4	31–50	26–45
5	41–60	36–50

material. Where appropriate, explicitly including air at $\rho = 0.00129 \text{ g/cm}^3$ adds some realism. A ‘vacuum’ may be replaced by an actual estimate of its density or by a substance with density chosen so low that it will not affect cascade development, yet will register stars during the Monte Carlo [7]. Its Z and A are best taken as an average over nearby materials. For a homogeneous shield the star density is already a continuous function of location so that no conversion is needed and results are reported as star density instead of dose.

Upon completion of the Monte Carlo run the array of star density or dose is smoothed in two stages—the second of which is optional. The 50×50 (z,r) CASIM output is first enlarged to a 60×60 array by simple (exponential) extrapolation: a set of 5×50 bins is added to the front and back, 60×5 bins are added at large radii and a 60×5 array is added at small ‘negative’ radii by reflection of the five innermost bins [8]. This 60×60 array is now divided into a 5×5 set of patches, each containing an array of 20×20 bins. Table 1 defines patches for each of five z-regions (independent of r) and for each of five r-regions (independent of z) in the augmented array. Also shown is where the patches fit in the 50×50 input/output array. For example: in the $Z_3\text{--}R_4$ -patch the lowest diagonal element, $P_{1,1}$, corresponds to (z,r) bin $A_{21,31}$ of the augmented array and the highest, $P_{20,20}$, to $A_{40,50}$. After smoothing this patch contributes to $A_{21,31}\text{--}A_{40,50}$ of the augmented array, which corresponds to bins $C_{16,26}\text{--}C_{35,45}$ of CASIM output. The patches overlap with their neighbors so as to minimize ‘seams’ in the final smoothed output which is obtained as a linear combination of the overlapping smoothed results. In the central patches the weights decrease linearly with distance from center in both z- and r-directions—from unity down to zero at the edges. Border patches carry unit weight on the side(s) (z, r, or both) which abut a border. These weights also come into play in the stop criterion. Patching—as opposed to smoothing the full 60×60 array at once—has the advantage of better preserving local structure in the outputs as well as reducing the range of values over which to smooth. Extending the array beyond its borders avoids the end-point interpolation property of the Bernstein polynomials. Thus the corners of the patch, which remain unchanged through the procedure, are moved outside of the region of interest and all results inside the

region participate in the smoothing.

The smoothing starts by taking the 2-D Bernstein polynomial of each patch and testing the results *versus* a stop criterion. If not satisfied, ratios of raw to smoothed outputs are taken and the Bernstein polynomial of this ratio is used to multiply the (previous) smoothed output, etc., until the stop criterion is met. Such an iteration step may be one of three types: (1) repeat 2-D on the ratios $Q(z, r) = \rho(z, r)/\rho_i^{sm}(z, r)$ which are smoothed to obtain Q^{sm} and $\rho_{i+1}^{sm}(z, r) = Q^{sm}(z, r) \cdot \rho_i^{sm}(z, r)$, where $\rho(z, r)$, $\rho_i^{sm}(z, r)$ are raw and smoothed contents of the bin at (z, r) , and $\rho_{i+1}^{sm}(z, r)$ becomes the smoothed estimate of the next iteration, (2) ratios of the 1-D marginal-z distribution $Q(z) = \sum_r \rho(z, r)(2r - 1)/\sum_r \rho_i^{sm}(z, r)(2r - 1)$ are smoothed whereupon $\rho_{i+1}^{sm}(z, r) = Q^{sm}(z) \cdot \rho_i^{sm}(z, r)$, for all r , and (3) similarly, for the marginal-r distribution $Q(r) = \sum_z \rho(z, r)/\sum_z \rho_i^{sm}(z, r)$, after which $\rho_{i+1}^{sm}(z, r) = Q^{sm}(r) \cdot \rho_i^{sm}(z, r)$ for all z .

In [1] the stop criterion is a linear combination of the square of the difference between raw and smoothed results (as a measure of precision) and second differences between adjacent bins (as a measure of smoothness). Here a somewhat more complicated criterion is adopted which makes use of five χ^2 -statistics of smoothed *vs* raw results for each array: (a) $\chi_1^2 = \sum_z \sum_r w_z w_r (\rho^{sm}(z, r) - \rho(z, r))^2 / \sigma_{mod}^2(z, r)$ where w_z, w_r are the same weights, referred to above, as are used to combine results of neighboring patches, and $\sigma_{mod}(z, r)$ is the modified relative error (see sec. 2) of the raw CASIM output times the average of $\rho(z, r)$ and $\rho^{sm}(z, r)$. Next, raw and smoothed marginal z-distributions of the patch are compared: (b) $\chi_2^2 = \sum_z (D(z) - D^{sm}(z))^2 / \sigma_D^2(z)$ where $D(z) = \sum_r \rho(z, r)(2r - 1)$, $D^{sm}(z) = \sum_r \rho^{sm}(z, r)(2r - 1)$ are the raw and smoothed marginal distributions of the patch and $\sigma_D(z)$ is the relative error of $D(z)$ —obtained as a combination of the errors on $\rho(z, r)$ —times the average of $D(z)$ and $D^{sm}(z)$. There is a similar comparison of raw and smoothed marginal r-distributions: (c) $\chi_3^2 = \sum_r (D(r) - D^{sm}(r))^2 / \sigma_D^2(r)$. The last two statistics to be compared are the distributions of the z- and r-moments, $M(z) = \sum_r \rho(z, r)(2r - 1)^2$ and $M(r) = \sum_z \rho(z, r)z$: (d) $\chi_4^2 = \sum_z (M(z) - M^{sm}(z))^2 / \sigma_M^2(z)$ and (e) $\chi_5^2 = \sum_r (M(r) - M^{sm}(r))^2 / \sigma_M^2(r)$ where $\sigma_M(z)$ and $\sigma_M(r)$ are again obtained by straightforward compounding of the $\sigma(z, r)$. The sums over z and r exclude all values outside the 50×50 CASIM array, i.e., those obtained by extrapolation. The stop criterion demands that $\sum_{i=1}^5 \chi_i^2 < 5.9$. (For a 5-dimensional uncorrelated Gaussian this contains the same fraction of all events as a 1-dimensional Gaussian does within one standard deviation.) A few exceptions to this are made after the iteration is well under way (some ten iterations): (i) when the largest change in any χ_i^2 is less than 5% which aborts slow convergence in which little is gained, and (ii) when χ_1^2 *increases* by more than a factor of two over its previous value, which may occur on rare occasions when dealing with very poor statistics. In this last case the prior iteration is adopted. When the number of iterations reaches 30 the iteration is likewise halted.

The χ_i^2 of the stop criteria are also used to decide what type of iteration—

if any—comes next. Smoothing always starts with full 2-D, i.e., type (1) above. Thereafter the choice is made according to which *cumulative* χ_i^2 ($i = 1, 2, 3$), i.e., which $C_i = \Sigma\chi_i^2$ summed over iteration number, is largest. After the largest C_i is identified, the next smoothing is of type i and the corresponding C_i is set to zero to begin accumulating anew. Smoothing the 50×1 array is much like for the 50×50 case: after augmenting to 60×1 , five 20×1 patches are smoothed and then linearly combined. The stop criterion is based on χ^2 of individual bins and of the average z as a function of radius of the patch.

The absence of a specific smoothness measure in the stop criterion need not be objectionable. Since smoothness can be expected to decrease with each iteration, one obtains—under the terms of the algorithm—the smoothest distribution consistent with the raw output. The first stage thus yields a smoothed 50×50 array composed of linear combinations of the 25 patches which have been smoothed separately. The smoothed array is then normalized to the total content of the raw one. To improve the smoothing along the borders, the entire procedure is repeated but with smoothed results replacing the extrapolated values outside the 50×50 array while the raw values still apply within. The number of such ‘supercycles’ is currently set at three and this parameter is at the disposal of the user.

A CASIM run concludes by drawing a set of contours of equal star density spaced in powers of ten. In standard CASIM contour locations corresponding to 10^n are estimated by grouping together results between $10^{n-\frac{1}{2}}$ and $10^{n+\frac{1}{2}}$ at a fixed depth. A weighted average of the radii associated with these bins provides an estimate of the contour location at that depth and a symbol is printed at that point. Such a plot then serves as a guide to draw a set of contours ‘by eye’ directed by the symbols of each contour as well as by location and shape of neighboring ones. The present algorithm permits the contour-drawing routine to be sharpened somewhat. In lieu of a weighted average the contour is found by (logarithmic) interpolation of the smoothed results at each z -location. By symmetry, the contours should meet the beam (or z -) axis at right angles. To accomplish this the contours are extended into ‘negative r ’ by reflecting the two points closest to $r=0$ about the z -axis. The program writes the results as a table of (z,r) -pairs for input to a graphics program. The above is referred to as ‘level-1’ smoothing.

A further (level-2) smoothing is available as an option. This consists of smoothing the level-1 contours—considered as functions $r(z)$. A relative error is assigned to each point of $r(z)$ by averaging the error on the CASIM outputs in that vicinity. The stop criterion is based on the single χ^2 measure between level-1 and level-2 contours at each z -location. There is no attempt to align with neighboring contours. Once level-2 smoothing is completed the ρ in the 50×50 array are adjusted to the new contours by interpolation assuming an exponential decline with r between contours at each z -location. Since this is a rather strong assumption no overall normalization to the raw results is attempted.

5 Energy Density

Compared to dose, smoothing of energy deposition distributions present different type problems. In particular one wishes to sharpen estimates of energy deposition along and near the beam axis at a considerably elevated level of precision. However, statistical uncertainty may yet enter into the problem. Beyond a certain depth in the target the increased angular spread of the cascades causes these small central bins to be frequented less and less, leading to poor statistics. Smoothing might thus be helpful at these larger depths. From the assumption of cylindrical symmetry—circular beam incident along target cylinder axis—it follows that maximum energy density occurs along the axis, usually some distance into the target. In addition to the (smoothed) energy density in this central region, it is desirable to estimate its on-axis distribution along with value and location of the maximum energy density in the target. When intra-bin variations or statistical errors are significant these cannot be simply identified as the largest value encountered among the on-axis bins and the corresponding location.

When the calculation involves more than one material, energy density—like star density—is not a continuous function of location. In contrast with star density, the problem involves *two* scaling parameters: nuclear interaction length (λ) as well as radiation length (L) and their influence varies with location so that converting to a continuous function is more complicated. While one could conceive of schemes to accomplish this [9], it is opted to convert to an approximately continuous energy density (in material with index=1) upon multiplying all energy densities by $[\lambda(n)L(n)/(\lambda(1)L(1))]^{1/2}$ where n is the material index of the bin location. While this choice is of scaling quite arbitrary, it should yield sufficient continuity across material boundaries. No large errors are expected from this since (1) in most instances the crucial regime in the inner core of the target or dump is homogeneous, and (2) the exact inverse operation is used at the end to convert smoothed results back to energy density. After this conversion, smoothing of the 50×50 array proceeds like for the dose/star density case.

For the inner 50×10 array, which is used with objective to estimate energy deposition at *points*, it is preferred to work with the cumulative distribution. The inner 50×10 array is first extended to a 60×10 array by extrapolation to 5×10 arrays to the front and the back. To enforce the condition that the energy density as a function of radius has a maximum at $r = 0$ the entire array is reflected about the z-axis into negative r . The array is divided into five (overlapping) 20×20 patches which are separately smoothed and then combined with linearly varying weights. Total energy deposited, $t_{i,j}$, in the (i,j)th bin $z_i \leq z < z_{i+1}$, $r_j \leq r < r_{j+1}$ is

$$t_{i,j} = \int_{z_i}^{z_{i+1}} dz \int_{r_j}^{r_{j+1}} \rho_E(z, r) 2\pi r dr \quad (3)$$

with a cumulative distribution

$$P(Z, R) = \int_0^Z dz \int_{R_T}^R \rho_E(z, r) 2\pi r dr = \sum_i \sum_j t_{i,j} \quad (4)$$

where ρ_E is the energy density. Going from outside, R_T , (low ρ_E) to inside (high ρ_E) improves numerical accuracy of the cumulative distribution. Only 2-D smoothing is performed. Following smoothing of the cumulative array, the $t_{i,j}$ are recovered and tested with a stop criterion entirely similar to that for the 50×50 case. Since errors are typically much smaller, it may take many more iterations to obtain an acceptable result. Therefore, the maximum number of iterations is increased to 100 and—to speed convergence when $\Delta\chi^2$ is small between iterations—smoothing may be performed on the square or cube of the raw-to-smoothed ratio [10]. When the stop criteria are satisfied, the final smoothed version is renormalized to total content of the raw output. The central energy density and the region nearby is calculated at a set of 251 (z) by 11 (r) equispaced *point locations*. The z -range spans the target length while the r 's are at $0, 0.0002, \dots, 0.002$ times the target radius which thus cover the innermost bin of the 10 bins which in turn divide the innermost bin of the 50×50 array. Thus in each patch the density is to be calculated at 101×11 points. From eq. (4) it follows [1] that at $R = 0$

$$\rho_E(Z, 0) = \frac{1}{2\pi} \frac{\partial^3 P(Z, R)}{\partial Z \partial^2 R}. \quad (5)$$

After the initial step $P(Z, R)$ is replaced by its Bernstein polynomial, $B_0(P)$, where the subscript denotes iteration number. For the smoothed central ρ_E , $B_0(P)$ replaces P in eq. (5), which becomes an expression in terms of finite differences of the P , to be evaluated at $R = 0$, at 101 Z -locations. The next iteration smooths the ratios $Q = P/B_0(P)$ to arrive at a new cumulative value $B_1 = B_0 B(Q)$ and a new estimate for $\rho_E(Z, 0)$:

$$\rho_E(Z, 0) = \frac{1}{2\pi} \left[B_0 \cdot \frac{\partial^3 B(Q)}{\partial Z \partial^2 R} + \frac{\partial B_0}{\partial Z} \cdot \frac{\partial^2 B(Q)}{\partial^2 R} + \frac{\partial^2 B_0}{\partial^2 R} \cdot \frac{\partial B(Q)}{\partial Z} + \frac{\partial^3 B_0}{\partial Z \partial^2 R} \cdot B(Q) \right] \quad (6)$$

and so on until the stop criterion on the array is met. Two terms from the general expression for $\partial B_0 B(Q) / \partial z \partial^2 r$ are absent from eq. (6) because they are identically zero at $R = 0$, by symmetry. Away from the axis, the expression

$$\rho_E(Z, R) = \frac{1}{2\pi R} \frac{\partial^2 P(Z, R)}{\partial Z \partial R} \quad (7)$$

is used in similar fashion to evaluate ρ_E at (Z, R) .

Central densities from neighboring patches are again combined linearly with weights which vary from unity at the center of a patch to zero at the edges. At each radius ($0, 0.0002, \dots, 0.002$) the 251 values of the energy density are smoothed— independent of one another—in nine overlapping patches of 51 points each. From

the on-axis energy density obtained in this fashion, the overall maximum energy is located and its value estimated by parabolic interpolation. At the z -value corresponding to the maximum on axis energy density the energy densities at the points are evaluated at the same ten radii to provide the near-axis radial dependence of ρ_E at the maximum.

6 Illustrations

No attempt is made here at a detailed analysis of the capabilities of the smoothing algorithm. Experience with a wide variety of realistic problems is needed to judge its usefulness. By way of illustration, sample results from the two cases elaborated upon earlier are presented here: estimates of dose in and around a large shield and energy density calculated along and around the axis of a target.

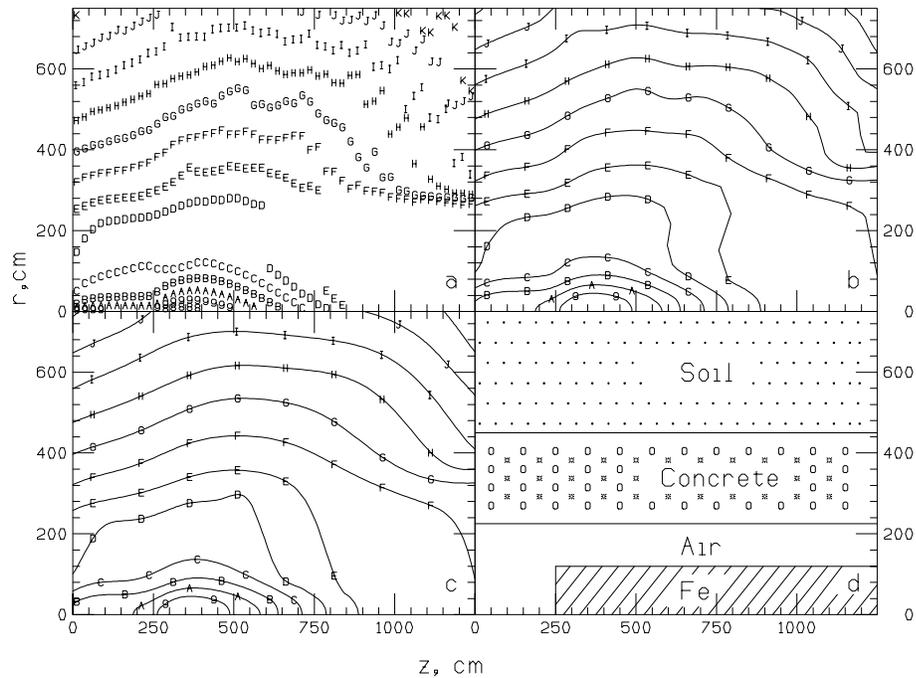


Figure 1: Contour plots for good statistics run of 800 GeV protons on geometry shown in (d). Plots are shown for (a) raw output, (b) level-1 smoothing, and (c) level-2 smoothing (c). A contour is labeled by $-\log_{10}$ of the dose it represents. Below 10^{-9} labeling continues with A, B, C, etc., representing 10^{-10} , 10^{-11} , 10^{-12} , etc.

Fig. 1 is the result of a run in which $3 \cdot 10^6$ 800 GeV protons are incident on a large shielding configuration—entirely fabricated for the purpose at hand. While it is a rather long run, it is representative of efforts made, e.g., in designing a beam dump. Fig. 1a shows contours as obtained without any smoothing, i.e., the symbol

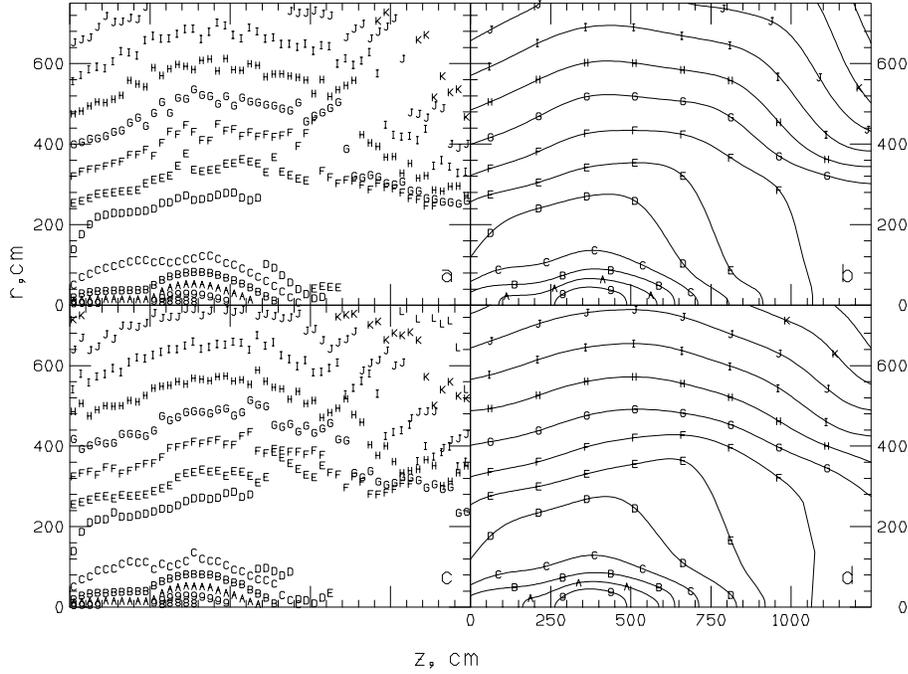


Figure 2: Contour plots for the same problem as fig. 1 but with decreased statistics. Plots are shown for (a) raw output, (b) level-2 for medium statistics, and (c) raw output, (d) level-2 for poor statistics. A contour is labeled by $-\log_{10}$ of the dose it represents. Below 10^{-9} labeling continues with A, B, C, etc., representing 10^{-10} , 10^{-11} , 10^{-12} , etc.

is printed at the dose level crossing as per the raw CASIM results. Figs. 1b and 1c show level-1 and level-2 smoothing, respectively, while 1d presents the geometry of the shielding. For a long run like this, it can be seen that any of the three ways predicts about the same result almost anywhere within the shielding, but both level 1 and 2 results are perhaps somewhat better suited for reading a dose off the graph and for presentation. In contrast, fig. 2 shows some shorter runs with the same geometry as in fig. 1. Figs. 2a and 2b show, respectively, the contours of the raw distribution and its level-2-smoothed version for $3 \cdot 10^5$ protons on the geometry of fig. 1d. It can be seen that while the smoothing definitely works harder here some problems arise in the form of underestimation at large depths and radii. Figs. 2c and 2d are the corresponding results for $3 \cdot 10^4$ protons on the same geometry. Serious underestimation occurs, especially at large radii where it approaches a factor of 10, yet the graph in fig. 2d appears reasonably smooth. This underlines the caveat that the smoothing procedures presented here are no substitute for CPU time.

For energy deposition, the example studied is that of an 800 GeV proton beam, $\sigma_x = \sigma_y = 0.1$ cm, incident on a homogeneous iron cylinder 1.5 m in length. Results of nine runs are analyzed which represent both the effects of statistics ($3 \cdot 10^4$, 10^5 ,

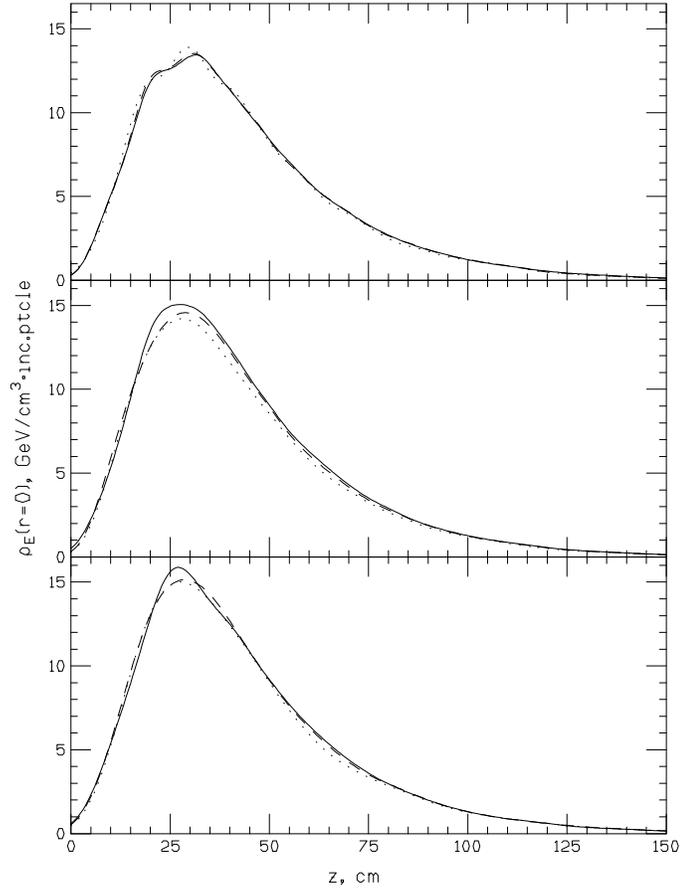


Figure 3: On axis energy density for 800 GeV proton beam, $\sigma_x = \sigma_y = 0.1$ cm, incident on a homogeneous iron cylinder. The curves connect $\rho_E(Z, 0)$ at 251 Z -points as calculated via eqs. (5) and (6). Solid-, dashed-, and dotted curves correspond to $3 \cdot 10^5$, 10^5 , and $3 \cdot 10^4$ incidents, respectively. **Top:** Target radius, $R = 50$ cm, **middle:** $R = 25$ cm, **bottom:** $R = 10$ cm.

and $3 \cdot 10^5$ incidents) and target radius (10, 25, and 50 cm). Note that these target radii correspond to $\Delta R = \sigma/5$, $\sigma/2$, and σ , respectively, which represent essentially the resolution of each case, where ΔR is the bin-width of the inner 50×10 array. *A priori*, only $\Delta R = \sigma/5$ is expected to provide adequate sampling of the ‘beam region’. Fig. 3 shows the energy deposition along the axis for the nine cases and it can be seen that while gross shape gets reproduced throughout, more detailed behavior, e.g., near the maximum in energy deposition, suffers considerably when either statistics or resolution is lacking. Fig. 4 shows the radial dependence near the axis, computed on axis and at another ten points out to ΔR , at the z -location of the on-axis maximum. Fig. 4 shows more quantitatively the need for both good statistics and, especially, sufficient radial resolution when exploring the region near

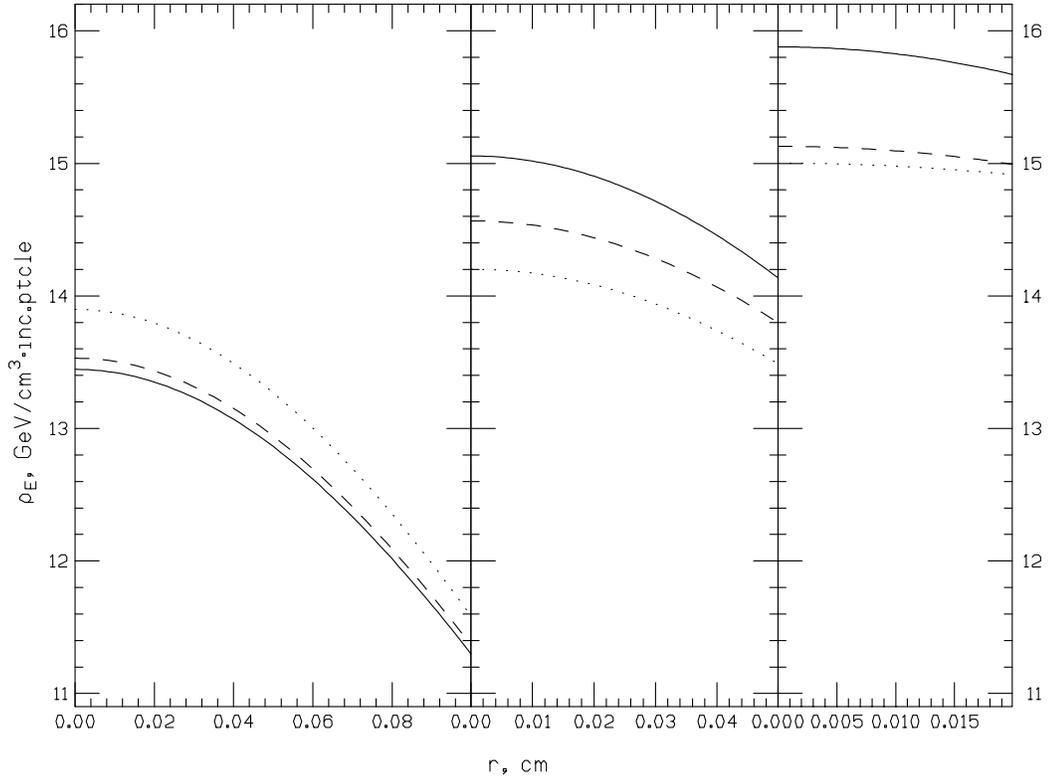


Figure 4: Radial dependence of energy density close to axis for the same cases as in fig. 3. The curves connect $\rho_E(Z, R)$ at 11 R-points as calculated via eq. (7) at Z corresponding to the maximum of $\rho_E(Z, 0)$. Solid-, dashed-, and dotted curves correspond to $3 \cdot 10^5$, 10^5 , and $3 \cdot 10^4$ incidents, respectively. **Left:** target radius, $R = 50$ cm, **middle:** $R = 25$ cm, **right:** $R = 10$ cm.

the beam. Reflection about the z-axis plus the variation diminishing property of the Bernstein polynomials(see [1]) are likely the reasons why there appears to be consistent underestimation of ρ_E^{max} at low resolution [11].

7 Concluding Remarks

Briefly, the smoothing algorithms introduced in [1] have been applied to problems of smoothing dose at large distances from the beam and to problems of trying to predict energy density at and near its maximum in the target. Options allow for two levels of smoothing of dose as a function of location: the higher level smooths dose contours obtained at the lower level. Another option enforces dose or energy density to decline monotonically with r .

As yet there has been no extensive experience with either algorithm variant in the CASIM context. In some limited tests with CASIM it seems to work quite

well. In what has been prepared for this note, and a few other examples, ZMOOZ works as expected. Broader experience may well lead to some revision, e.g., of the specific form and constants adopted for the stop criteria. These constants are identified in the program listing and are thus at the easy disposal of the user. The time required to run the smoothing procedure is small—most often negligible—compared with that of a typical Monte Carlo. Running times are significant only when large inequalities between neighboring bins are present in the results with small statistical error—which is likely due to poor choice of bin size. Even then CPU time for smoothing remains a small fraction of that of the Monte Carlo. But it is well to emphasize again here that smooth does not mean accurate and therefore that smoothing is not a substitute for CPU time.

One possible further development would be to combine smoothing and statistical analysis. Presently in CASIM statistics are gathered over the course of 10 sub-runs of the Monte Carlo. It might make sense to smooth each sub-run to provide a statistical error of the (final) smoothed value. But such an error would be unambiguous only if the final result is derived as the average of the smoothed sub-runs—which would likely be inferior to smoothing based directly on the entire Monte Carlo run. To use the same type of stop criterion, one must then collect statistics on each sub-run by further dividing it into (perhaps another 10) sub-sub-runs.

The goal of the present effort is to help interpret raw Monte Carlo results much like a careful and unbiased user might do—though the algorithm is likely to make less use of *a priori* information. Only frequent usage over a large range of problems can decide how well that goal is achieved. The smoothing techniques used here are obviously not limited to CASIM but can be applied to similar programs and, with extra labor, to entirely different applications as well. A more portable and adaptable version of ZMOOZ with variable array and patch size would certainly be more useful than the present, more limited, effort. Such a flexible version could also include—as user options—a variety of boundary and other conditions. The number and type of such conditions along with the method by which they are introduced may amount to a large coding effort. Such an undertaking best awaits some experience with ZMOOZ on realistic problems.

My thanks to N. Mokhov for a careful reading of the manuscript and for many useful comments.

References

- [1] A. Van Ginneken, Nucl. Inst. Meth. **A305**, 453 (1991) of which Fermilab-Pub-90/197 is an expanded version.
- [2] A. Van Ginneken, Fermilab-FN-272 (1975).

- [3] For large scale shielding, variable spacing can be ignored to first approximation, i.e., simply relabel the bin coordinates to correspond to equispacing, perform the smoothing, then convert back to the correct variable spacing. More finesse will be required when high resolution is needed, such as for determining maximum energy density, though for such problems variable spacing offers relatively little advantage.
- [4] Already in plain CASIM smoothness is no guarantee of low statistical error. Because of the weighting techniques introduced in CASIM, there is less of the ‘graininess’ of analog Monte Carlo and, often prematurely, results leave an impression of being statistically precise.
- [5] This may well put the outer limits of the 50×50 array in the wrong place. However, their precise location will have little influence on the central energy deposition and at worst another run is needed to cover the problem at larger radii.
- [6] This yields a valid estimate of dose equivalent only at large distances and outside of a sufficient thickness of soil or concrete since otherwise the conversion factor (which assumes the star density is produced by an ‘equilibrium spectrum’ in soil) is not valid. It is assumed that bin borders coincide with the materials boundaries of the problem, otherwise results lack clear interpretation even in the raw output.
- [7] Unlike in analog Monte Carlo ‘recording stars’ in CASIM are calculated at equispaced points along a hadron’s trajectory in all media, except in vacuum. To pick a number: a ‘vacuum density’ of 10^{-10} g/cm³ is practically invisible to the cascades, yet large enough to avoid underflows, etc.
- [8] Instead of negative r one may envision ρ as a function of, e.g., x or y —in which case a negative coordinate has a more physical meaning.
- [9] The first step of which is likely to keep a separate tally of energy deposited by hadrons and by electrons and photons.
- [10] This is done only for energy density smoothing. For stars, alternate smoothing of the 2-D array and the marginal distributions appears to prohibit this type of slow convergence. For energy deposition marginal smoothing may work in the early stages but tends to ‘oversmooth’ when convergence slows.
- [11] Other techniques which do not symmetrize ρ_E but force a maximum (or saddle point) at $r = 0$ by fixing ρ_E at negative r are found, in many instances, to predict unphysical behavior of the near-axis radial distribution.