

Monte Carlo simulation of electron beams from an accelerator head using PENELOPE

J Sempau¹, A Sánchez-Reyes^{2,3}, F Salvat⁴, H Oulad ben Tahar⁴,
S B Jiang⁵ and J M Fernández-Varea⁴

¹ Institut de Tècniques Energètiques, Universitat Politècnica de Catalunya, Diagonal 647, 08028 Barcelona, Spain

² Servei d'Oncologia Radioteràpica, Hospital Clínic de Barcelona, Villarroel 170, 08036 Barcelona, Spain

³ Institut d'Investigacions Biomèdiques August Pi i Sunyer (IDIBAPS), Universitat de Barcelona, Spain

⁴ Facultat de Física (ECM), Universitat de Barcelona, Societat Catalana de Física (IEC), Diagonal 647, 08028 Barcelona, Spain

⁵ Department of Radiation Oncology, Stanford University School of Medicine, 300 Pasteur Drive, Stanford, CA 94305-5304, USA

E-mail: josep.sempau@upc.es

Received 22 May 2000, in final form 1 February 2001

Abstract

The Monte Carlo code PENELOPE has been used to simulate electron beams from a Siemens Mevatron KDS linac with nominal energies of 6, 12 and 18 MeV. Owing to its accuracy, which stems from that of the underlying physical interaction models, PENELOPE is suitable for simulating problems of interest to the medical physics community. It includes a geometry package that allows the definition of complex quadric geometries, such as those of irradiation instruments, in a straightforward manner. Dose distributions in water simulated with PENELOPE agree well with experimental measurements using a silicon detector and a monitoring ionization chamber. Insertion of a lead slab in the incident beam at the surface of the water phantom produces sharp variations in the dose distributions, which are correctly reproduced by the simulation code. Results from PENELOPE are also compared with those of equivalent simulations with the EGS4-based user codes BEAM and DOSXYZ. Angular and energy distributions of electrons and photons in the phase-space plane (at the downstream end of the applicator) obtained from both simulation codes are similar, although significant differences do appear in some cases. These differences, however, are shown to have a negligible effect on the calculated dose distributions. Various practical aspects of the simulations, such as the calculation of statistical uncertainties and the effect of the 'latent' variance in the phase-space file, are discussed in detail.

1. Introduction

High-accuracy calculations of dose distributions inside tumours and the surrounding tissue are essential for assessing the effectiveness of a given treatment in terms of the probability of tumour control and of radiation-induced complications. Thus, the AAPM (1983) protocol indicates that relative errors as small as 5% can be decisive in evaluating whether a treatment is suitable.

Although Monte Carlo (MC) simulation is, at present, the most accurate method for describing the coupled transport of electrons and photons in complex geometries, residual inaccuracies may nevertheless exist. These can be caused by either statistical uncertainties, which are inherent in the method, or by the approximations used in the modelling of cross sections, in the space displacement algorithms employed for charged particles⁶, in the geometrical description of the material system, etc. Whereas statistical effects can be reduced to any predetermined level by simply simulating more particles, the use of approximate models for the physical processes is unavoidable and their effect on the final results is not always easily predictable.

An MC algorithm can be trusted only after these limitations have been proved to be irrelevant, ideally through extensive benchmarking with experimental results. However, comparison with experiments may be inconclusive in many cases for a number of reasons, for example because geometrical details of complex systems are not always unambiguously defined or because experimental results are sometimes affected by correction factors (such as those applied to ionization chamber readings in radiotherapy clinical practice) which are unknown or determined with the aid of an MC simulation. Experimental errors can also be the origin of discrepancies. An alternative and desirable validation practice would be to compare results from various independent MC codes, which employ different models and approximations. This should, at least, indicate how sensitive simulation results are to the differing features of the codes used. In spite of this, most developments in the area of medical applications, and more specifically on therapy machine heads, which present a considerable degree of geometrical complexity, rely almost exclusively on two well-known code systems, EGS4 (Nelson *et al* 1985) and ETRAN-ITS (Halbleib *et al* 1992). Other simulation packages, such as MCNP (Briesmeister 1997) are also employed by some authors (see for example Lewis *et al* (1999) or Schach von Wittenau *et al* (1999)). However, the comparison of results obtained with different algorithms is very rarely made—as an example, see Siebers *et al* (1999) for a comparison of photon beams with EGS4 and MCNP4b.

The history and achievements of the MC simulation of accelerator heads have been discussed in several reviews (Andreo 1991, Ma and Jiang 1999). Initially, most studies were devoted to photon beams (e.g. Petti *et al* 1983, Mohan *et al* 1985, Mohan 1988). In the late 1980s, the simulation of electron beams started to attract the interest of researchers (see for example Udale (1988), Udale-Smith (1992), Ding *et al* (1996), Ma *et al* (1997)). Many of these studies were performed with BEAM, developed by Rogers *et al* (1995) within the framework of the OMEGA project (Rogers *et al* 1990). BEAM, which is based on the EGS4 system, allows the user to simulate high-energy electron and photon beams by modelling the different hardware elements with a variety of geometrical entities, such as slabs, stacked cones, flattening filters, etc.

In the present paper we describe the MC simulation of electron beams from a Siemens Mevatron KDS linac operating at the Hospital Clínic de Barcelona. Two codes were employed: PENELOPE (Salvat *et al* 1996), a general-purpose package developed by some of the authors of this work, and BEAM. The aim is to enlarge the spectrum of codes available to the

⁶ In general, we shall use the term 'particle' to refer either to photons, electrons or positrons.

medical physics community by showing, with the aid of a practical example, the suitability of PENELOPE for dealing with complex radiation transport problems involving electrons and photons. We would like to stress that the present paper is primarily concerned with the accuracy of the calculation, since our aim is to determine the degree of agreement that can be expected between different MC algorithms and between these and the experimental data. Speed, in our approach, is a secondary issue.

The rest of the paper is organized as follows. In section 2, the main characteristics of the PENELOPE code system are described. The modelling of the accelerator geometry is presented in section 3. Section 4 is devoted to important features of the simulation of these type of systems, such as calculation of uncertainties, parallelization, variance reduction, etc. Simulation results from PENELOPE are compared with both measurements and EGS4 results in section 5, and some conclusions are drawn in section 6. In the appendices, we present a detailed discussion on practical features of the calculation of statistical uncertainties and on their propagation from the phase-space file.

2. The PENELOPE code system

PENELOPE (an acronym for PENetration and Energy LOSS of Positrons and Electrons in matter—photons were introduced later) is a general-purpose MC code system for the simulation of coupled electron–photon transport in arbitrary materials and complex geometries that has been developed at the Universitat de Barcelona over the last 10 years (Baró *et al* 1995, Salvat *et al* 1996, Sempau *et al* 1997). It is applicable to energies ranging from 1 keV up to 1 GeV, approximately. The core of the system is a FORTRAN 77 subroutine package that generates and follows particle trajectories inside material systems consisting of a number of homogeneous regions, called bodies. These subroutines are invoked from a simple steering main program, to be provided by the user, which controls the evolution of the tracks and keeps score of the quantities of interest, such as energy deposition, particle current, etc. A particle track is considered as a sequence of free flights or ‘jumps’. At the end of each jump, the particle suffers an interaction with the medium, a ‘knock’, where it loses energy, changes its direction of flight and, in certain cases, produces secondary particles. The actions to be taken in the main program are the same for all kinds of particles, thus simplifying its structure. The complete package (including FORTRAN source files, database, geometry viewer and debugger and documentation) is available from the NEA Data Bank⁷ and through the Radiation Shielding Information Center (Oak Ridge).

MC codes often use condensed schemes (or class I, according to Berger’s (1963) terminology) to simulate high-energy electron and positron transport. Thus, the particle track is divided into a large number of segments (steps) long enough to ensure that many interactions occur along each step. The energy loss and the global angular deflection of the particle due to the interactions in each step are sampled from suitable multiple-scattering theories (Goudsmit and Saunderson 1940, Lewis 1950, Molière 1948). The procedure is not completely satisfactory due to the fact that the distribution of space displacements is partially unknown. Moreover, when particles approach an interface, the step length must be progressively reduced to ensure that the complete segment is contained within the initial medium, which conflicts with the use of preselected step lengths and thus complicates the logic of the code.

PENELOPE has recourse to a systematic use of mixed simulation (class II schemes in Berger’s terminology) for electrons and positrons. A mixed simulation algorithm requires the

⁷ OECD Nuclear Energy Agency Data Bank. Le Seine Saint-Germain, 12 Boulevard des Iles, 92130 Issy-les-Moulineaux, France. E-mail: NEA@db.nea.fr; <http://www.nea.fr>

specification of cut-off values for the energy loss and the polar scattering angle. Individual interactions with energy losses or angular deflections larger than these cut-offs are termed hard (or catastrophic) and simulated in a detailed way. The global effect of the remaining interactions, called soft, that occur between each pair of successive hard interactions is simulated as a single 'artificial' event, where the particle loses energy and changes its direction of movement according to the corresponding multiple scattering theories. Space displacements are simulated by using a simple mechanism, the so-called 'random hinge'⁸, that, owing to the relatively gentle effect of soft events, is particularly accurate and applicable even in the vicinity of interfaces so that no step-size adaptive algorithm is needed. Thus, as opposed to EGS4 when the PRESTA (Bielajew and Rogers 1987) module is active (as is the case for all the EGS4 simulations reported in this paper), PENELOPE does not require knowledge of the *minimum* perpendicular distance between the charged particle and the boundaries of the material region where this particle moves. Since the calculation of this minimum distance implies, in the general case of an arbitrary quadric surface, solving the roots of a sixth-order polynomial equation (see Bielajew (1995) for a detailed discussion on the subject), the application of PRESTA is, in practice, limited to a subset of quadrics for which the minimum distance can be obtained with a modest computational effort (i.e. planes, cylinders along the *z*-axis, spheres, etc). Fortunately, many real-life problems can be described with this subset of quadrics, as happens in the case of an accelerator head. Hence, the combined use of class II schemes for electrons and positrons and the random hinge mechanism not only improves the accuracy of the simulation but also simplifies the operation of the code.

We would like to point out that the multiple-scattering formalism adopted in PENELOPE does have a limitation. This limitation is related to the requirement that scattering in the medium has to be multiple, that is, at least about 20 artificial interactions have to take place. To guarantee the fulfillment of this condition the parameter HFPMAX, which sets an upper bound for the allowed values of the mean free path between hard interactions, can be defined by the user for each material. In our simulations we set HFPMAX equal to $\simeq 1/20$ of the characteristic width of each homogeneous region. This is particularly important in the case of an electron accelerator head due to the presence of one or more very thin vacuum window and scattering foils in the beam path (see the description of the geometry below).

Two different energy loss cut-offs for charged particles, WCC and WCR, are defined for the collisional and radiative (bremsstrahlung) losses respectively. The angular cut-off, in turn, is indirectly defined through the mean free path $\lambda_{el}^{(h)}$ that results for hard elastic collisions. The parameter C1, set by the user, determines $\lambda_{el}^{(h)}$ by means of the relation

$$C1 \equiv \frac{\lambda_{el}^{(h)}}{\lambda_1} \simeq 1 - \langle \cos \theta \rangle^{(h)} \quad (1)$$

where λ_1 is the so-called first transport mean free path. The last term in equation (1) evidences the approximate relation between C1 and the mean angular deflection $\langle \cos \theta \rangle^{(h)}$ suffered by the electron along a hard mean free path. In PENELOPE, C1 can take values from 0 to 0.2. When set to zero, the simulation of angular deflections is detailed, i.e. no multiple scattering approach is used. The higher its value, the larger the cut-off and, hence, the steps taken, thus speeding up the simulation at the cost of reducing somewhat the accuracy of the simulated transport process. An additional restriction on the angular cut-off is imposed through the parameter C2, defined as the maximum fractional kinetic energy that can be lost in a single step. Its value can be set from 0 (detailed simulation) up to 0.1. The user also defines three absorption energies, EABS(1, 2, 3) for electrons, photons and positrons respectively; when a particle reaches the

⁸ The term was coined by Dr R P Kensek of Sandia National Laboratories.

corresponding absorption energy, its simulation is discontinued and its energy is assumed to be locally deposited.

In spite of the fact that PENELOPE and EGS4 use quite different physical models, both codes have recourse to mixed simulation to account for energy losses of electrons and positrons. As a result, some of the PENELOPE user-defined parameters find their counterpart in EGS4. Thus, WCC and WCR are equivalent to AE and AP respectively, except that in EGS4 the rest mass is included in all definitions affecting electrons or positrons. Furthermore, C2 is analogous to ESTEPE, when the latter is active. The set of parameters EABS are equivalent to ECUT and PCUT. In contrast, there is no analogy for C1, since EGS4 does not employ mixed simulation for the angular deflections. This observation is, perhaps, the most relevant feature that distinguishes PENELOPE from EGS4.

Photon transport is performed by using the conventional detailed scheme, i.e. all the interactions are simulated in chronological succession. PENELOPE also considers the emission of characteristic x-rays and Auger electrons that result from vacancies produced in K shells by photon interactions. Radiation resulting from the filling of vacancies in outer shells is not followed; in these cases, the excitation energy is taken out by the emitted Compton electron or photoelectron. This limits the validity of the simulation to photons with energies larger than the L-shell ionization energy of the heaviest element in the medium, that is, about 15 keV for the accelerator head under study. The probability of inner-shell ionization by electron impact is very small. When it occurs, it is assumed that the excitation energy of the residual ion is locally deposited.

The code system also includes a subroutine package, called PENGEO2, for automatic tracking of the particles in quadric geometries, i.e. material systems consisting of bodies limited by quadric surfaces. Bodies can be grouped into modules that, in turn, can form part of larger modules. This hierarchic modular structure allows a large simplification of the geometry calculations. Moreover, individual modules can be independently translated and rotated with all their submodules. This makes the configuration of a given accelerator easy to modify (e.g. to open or close the jaws). The auxiliary program GVIEW, which operates under DOS, displays 2D cuts of the geometry on the computer screen and is useful for debugging the geometry definition file.

The geometry definition logic adopted in PENGEO2 is similar to the 'Constructive Quadric Geometry' used in the POV-Ray⁹ ray-tracing software. Recently, a FORTRAN translator of the PENGEO2 definition files to POV-Ray format has been developed to generate true 3D pictures of the geometry on a variety of platforms. It will be made available in the future.

3. The geometry of the Siemens Mevatron KDS accelerator

The Siemens Mevatron KDS linac is a standing-wave linear accelerator with a 270° deflection chamber that provides electron beams with nominal energies of 6, 8, 10, 12, 15 and 18 MeV. The elements of the head assembly have been modelled, both in PENELOPE and BEAM, following the specifications provided by the manufacturer. A Cartesian reference frame is defined such that the z -axis coincides with the symmetry axis of the primary collimator (the third element in the list below), pointing along the direction of motion of the beam, and with the origin of coordinates located at the entrance side of the vacuum window (first element in the list). The elements considered are:

⁹ Persistence of Vision Ray Tracer (POV-RayTM) version 3.1. This is a powerful tool for producing high-quality computer graphics. POV-Ray is copyrighted freeware and can be downloaded from the Internet at <http://www.povray.org>

- (a) The vacuum window, at the exit of the deflection chamber, consisting of two tungsten foils with an inner water flux for cooling.
- (b) A primary scattering foil of gold or iron. Its thickness and composition depend on the energy setting.
- (c) A cylindrical tungsten primary collimator with a conical aperture.
- (d) An aluminium secondary foil. Homogeneity of the electron field is achieved through the combined effect of scattering and absorption in the primary and secondary foils, respectively.
- (e) A cylindrical monitoring ionization chamber consisting of three mylar foils, separating two volumes of air. For simplicity, this element has been modelled as a single layer of mylar of equivalent thickness. Following the manufacturer's indications, the conductor of the chamber is ignored.
- (f) The mirror of the Siemens KDS is automatically removed when the accelerator is operated in the electron beam mode. Therefore this element has not been included in the simulation.
- (g) Two pairs of tungsten jaws that limit the radiation field in the x and y directions. For electron beams, as opposed to photon beams, their aperture, set so as to produce a $25 \times 25 \text{ cm}^2$ field at 100 cm from the source ($\text{SSD} = 100 \text{ cm}$), is fixed irrespective of the beam energy or the actual field size at the patient surface, which is determined by the applicator—see below.
- (h) A shielding cover and an accessory holder, both modelled as iron plates with a square opening to let the beam pass through. The arms of the accessory holder, which are near the beam border, have been modelled as four slanted iron cylinders. Since BEAM does not allow the modelling of these latter structures they were the only element not considered in the EGS4 simulations, an approximation that is well justified due to their relatively small contribution to the quantities of interest studied here.
- (i) A variable electron applicator consisting of two pairs of small lead jaws that collimate the field in the x and y directions. Its aperture can be varied to produce rectangular fields at the patient surface with sides ranging from 6 to 20 cm. In the present study, only a $10 \times 10 \text{ cm}^2$ field size has been considered. In order to reduce the number of geometrical calculations, only those parts of the applicator that are close to the beam axis have been included. These are modelled as a square lead frame and four cylindrical aluminium bars connecting the frame with the small jaws. The total distance from the vacuum window to the last small jaw of the applicator is 95 cm.

The whole system is immersed in dry air. In addition to the elements described above, vertical walls have also been defined in order to limit the lateral dimensions of the simulation 'universe'. This limitation fulfils two aims: firstly, it facilitates the visualization of the geometry with rendering tools such as POV-Ray; secondly, since tracks that reach the outer vacuum are no longer simulated, time is saved by discarding those particles that are too far from the beam axis to reach the phantom surface. Figure 1 shows a 3D rendering, generated with POV-Ray, of the geometry of the accelerator head.

4. The simulation

The process of calculating dose distributions within a patient is usually split into at least two steps. First, the particles are transported through the treatment head and beam modifiers down to a predefined scoring plane, just before the patient surface, and their final state is stored in a phase-space file, PSF for short. Alternatively, some authors determine only the energy and/or angular distributions of the radiation field at the phantom surface, although this




Figure 1. Siemens Mevatron KDS accelerator head as modelled for this study. In the BEAM simulations the four slanted cylinders in the accessory holder were not considered. This 3D view was generated by POV-Ray—see text for details.

procedure disregards correlations that, in principle, could be relevant for an accurate description of dose distributions in the patient. The PSF contains the type of particle, its position, direction of motion, energy and statistical weight (only necessary if variance-reduction—VR hereafter—techniques are employed) for each particle that reaches the scoring plane. To evaluate statistical uncertainties properly, it is also necessary to store information about the primary ancestor of the particle, as will be discussed below. Some authors consider a further division of this first step into two substeps so as to separate the simulation of the patient-independent portion of the head geometry from that of the beam modifiers (if any) which, in principle, can depend on the particular treatment being considered and, therefore, on the patient.

In our simulations, electrons coming out of the deflection chamber and entering the vacuum window are modelled as a point-like, monoenergetic and monodirectional source. Although some more refined source models, which take into account spatial, angular and energy spreads, can be included in PENELOPE in a straightforward way, our simulation results showed little sensitivity to these details of the initial beam in comparison to other features such as the geometry of the treatment head and the accuracy of the description of the transport process. The trajectory of each particle, either primary or secondary, is followed until its energy drops below the absorption energy, EABS, or until it leaves the system. The scoring plane is defined perpendicular to the z -axis and located 5 cm below the last small applicator jaw, so as to make the source-to-surface distance (or SSD) equal to 100 cm.

In a second stage, the particles contained in the PSF are transported through the irradiated object, which in our case is a water phantom. In the PENELOPE simulations it has been modelled as a semi-infinite homogeneous body divided into spatial bins where energy depositions are scored. The EGS4-based user code DOSXYZ (Ma *et al* 1995), which considers the phantom divided in a large number of small volume elements, or voxels, has also been employed and its results compared to those obtained with PENELOPE.

Nominal energies of 6, 12 and 18 MeV are considered in this work. Since the average kinetic energy of primary electrons just before entering the vacuum window does not

necessarily coincide with the nominal energy specified by the manufacturer, some preliminary inspection is required to determine the value to be used in the simulation. The procedure that we adopted was to progressively modify the energy until the experimental and simulated values of the depth R_{50} at which the central axis depth-dose reaches 50% of its maximum agreed reasonably well. Thus, the initial energies were set to 6.80, 12.75 and 18.87 MeV respectively for the three nominal energies mentioned above. The same initial energies were used for both PENELOPE and BEAM. An alternative approach could have been to determine the initial energies for the two codes by matching the results from each code with those from the experiments. We have deliberately avoided this method in order not to introduce possible sources of discrepancy between the two sets of simulation results. As a consequence, only comparisons between PENELOPE and experiments, on the one hand, and between PENELOPE and EGS4, on the other, are really meaningful and hence comparisons between EGS4 and experiments will not be considered.

Because of continued development, accelerator units of the same series may differ in geometrical details, which can be difficult to elucidate from the commercial specifications. In this respect, special care must be taken to make sure that all the relevant elements have been correctly modelled. In our case, some preliminary calculations put into evidence that an important element was missing in the description of the accelerator. It turned out that the vacuum window in our unit incorporates a cooling system with two tungsten foils (see the geometry description above) which was not included in the first simulations.

With regards to user-defined parameters, all PENELOPE simulations were carried out with absorption energies EABS equal to 100 keV and 10 keV for electrons (or positrons) and photons respectively. Mixed simulation cutoffs WCC and WCR were set to 5 keV and 1 keV respectively. Angular deflection cut-offs are determined by the values of the parameters C1 and C2, which were chosen to be 0.05. EGS4 simulations, in turn, were performed with ECUT, PCUT, AE and AP set to a kinetic energy of 10 keV. It is interesting to point out that EGS4 makes use of the continuous slowing down approximation (CSDA) to describe energy losses *below* AE and AP. As is well-known, CSDA does not account for energy loss straggling and, in consequence, EGS4 has to be forced to simulate most inelastic interactions in a detailed way to obtain accurate electron energy distributions (i.e. AE and AP have to be set to a value smaller than the width of the energy distribution). This slows down the simulation slightly. On the contrary, PENELOPE implements a model for energy losses below WCC and WCR that preserves the first *two* moments of the energy loss distributions, namely the average energy loss (which would be equivalent to CSDA) and the energy straggling. Hence, higher cut-offs could, in principle, be selected, although we preferred not to take advantage of this feature—actually the PENELOPE cut-offs were set to values even lower than those selected for the EGS4 simulations—so as to minimize the effect of condensed simulation over the MC results.

4.1. The phase-space file

The variables stored in the PSFs generated by PENELOPE are the particle type, its position (X and Y coordinates), direction cosines corresponding to the x and y directions (U and V , respectively), energy, statistical weight—which is irrelevant if the simulation is analogue (i.e. when no VR is applied)—and an integer label (HIST), whose use will be explained in full detail later. To avoid filling up the PSF with superfluous information, for the $10 \times 10 \text{ cm}^2$ field studied here, only those particles that reach the scoring plane with lateral coordinates $|x|$ and $|y|$ smaller than 10 cm are stored. This defines the ‘universe’ of the PSF as a $20 \times 20 \text{ cm}^2$ square centred on the z -axis.

Since $z = 100$ cm for particles in the PSF, there is no need to store this coordinate. Furthermore, the direction cosine corresponding to the z -axis, W , can be obtained as

$$W = +\sqrt{1 - U^2 - V^2}. \quad (2)$$

If particles carry unit weight, all the required information for a single particle (in REAL*4 precision) can be stored in 24 bytes or, in other words, every million particles occupy less than 22.9 MB of disk space. An adequate binary representation of the pertinent state variables permits the storing of a PSF containing 10^8 particles in 2.24 GB of disk space. The use of compression techniques would somewhat reduce the required volume even further. Since this PSF would be large enough for virtually all present-day practical purposes, one might think that the availability of low-cost, portable mass storage devices might render the use of semiempirical models of the PSF, and the corresponding approximations introduced by them, unnecessary. However, these models continue to be useful for a number of reasons, such as facilitating the tasks involved in the commissioning of the machine, an example of which is given by Jiang *et al* (2000).

4.2. Parallelization and variance reduction

Even with the use of fast single-processor computers, the high-accuracy MC simulation of an accelerator head is a task that may take several days. Simulation efficiencies can be significantly improved by using conventional VR techniques (e.g. range rejection, splitting and Russian roulette, etc) or/and by parallelizing the computer program. Nowadays, parallel computers, essentially formed by many processors sharing a common random access memory, and task farms consisting of many interconnected stand-alone computers, are becoming widely available to the medical and academic communities. In both cases, the basic idea is to split a lengthy calculation into smaller parts which are then run simultaneously by different processors, thus effectively increasing the speed by a factor which is roughly equal to the number of processors involved. In order to take advantage of these architectures, the computer code must be modified so as to define how to split the job to be done into smaller tasks and how to gather the results obtained by the various processors.

In this work, and for comparison purposes, we have separately employed both VR and parallelization. On the one hand, several VR techniques are already embedded in BEAM, among which range rejection is the most relevant for our study. This technique, consisting of discarding those electrons with energies too low to reach the scoring plane, has the effect of ignoring any bremsstrahlung photon that would have been produced by the discarded particles. BEAM minimizes these effects in different ways (Rogers *et al* 1995) and takes full advantage of range rejection. This allowed us to run the code on a modest Pentium Pro 200 MHz personal computer.

On the other hand, bearing in mind the fact that our primary concern was accuracy, not speed, we opted not to use any VR technique in the PENELOPE user code developed in this work to obtain the PSFs. To compensate for the much lower simulation efficiency that this implied, the code was parallelized and run on eight processors of an SGI Origin 2000, a parallel machine located at the Centre de Supercomputació de Catalunya (CESCA). The parallelization was carried out with the aid of the message passing interface (MPI) library, version 1.2, which is freely available and can be used on several platforms. The necessary routines can be invoked from the main program, which is unique for all processors. Each processor, or group of them, is then identified by an index and the execution of the code is forked accordingly.

Parallel programming has been employed to reduce computing times in the past (Ma 1992). One of the features where care must be exercised refers to the generation of random numbers.

The main concern here is to make sure that the contributions to the quantity of interest from particles simulated by different processors are uncorrelated. This can be accomplished by using special random number generators (RNGs) that were developed to produce multiple, ‘high-quality’ sequences in parallel machines (see, for example, Mascagni *et al* (1995)). Another approach, which we have followed in this study, consists of using the same RNG implemented in the single-processor MC code. PENELOPE, for example, combines two different pseudo-random sequences according to the scheme proposed by l’Ecuyer (see Salvat *et al* (1996) for a detailed description of this implementation), resulting in an RNG of the congruential type with a period of the order of 10^{18} . It has been claimed that one can safely use this approach provided that different seeds are introduced in the various processors. The lack of correlation is reasonably guaranteed by the fact that the simulation of each particle involves a random number of calls to the RNG, thus making any ‘synchronism’ between different processors very unlikely. The results obtained in this study seem to corroborate the validity of this assumption.

4.3. Uncertainties

After completing a number N of histories¹⁰, the expected value $\langle q \rangle$ of a certain quantity of interest q (e.g. the energy deposited in a certain voxel) is estimated by means of the average \bar{q} , namely

$$\bar{q} \equiv \frac{1}{N} \sum_{i=1}^N q_i \quad (3)$$

where q_i is the contribution to q due to the i th history. By virtue of the central limit theorem (Gnedenko 1975), in the limit $N \rightarrow \infty$ the probability distribution of \bar{q} is Gaussian. In this limit

$$s^2 \equiv \frac{1}{N} \left(\frac{\sum_i q_i^2}{N} - \bar{q}^2 \right) \quad (4)$$

is an unbiased estimator of the variance of \bar{q} . Simulation results are commonly expressed in the form $\bar{q} \pm ks$. In this work all uncertainties have been obtained using $k = 3$ so that the probability that the true value $\langle q \rangle$ lies within the error bar is 99.7%.

It is worth stressing the fact that q_i refers to the sum of contributions of a primary particle (say q_{i1}) and all the secondaries originated by it (q_{ij} with $j = 2, \dots$). This can be expressed by writing $q_i = \sum_j q_{ij}$. As a consequence, the average of squares in equation (4) cannot be obtained by summing over contributions from individual particles, i.e. $\sum_i q_i^2 \neq \sum_{ij} q_{ij}^2$. Owing to the fact that the simulation is split into two parts, as explained above, the former observation has a practical consequence on the structure of the PSF, because the intermediate information contained in it ought to make the identification of all the secondary particles with the same primary ancestor possible. To this end, each record output by PENELOPE carries an integer label, HIST, whose purpose is to tag the current particle with the index of the primary particle that it originated from.

A few words regarding programming details are in order here. Consider, for instance, the simulation of the 3D dose distribution in a phantom involving a large number of voxels, of the order of one million. Each voxel has an associated random variable q . After completion of a history, a literal application of formula (4) would imply the updating of the counters storing the energy deposited in each voxel and its square for *all* voxels. This approach, which represents

¹⁰ Hereafter the term ‘history’, when used in this context, must be understood as the process of simulating a primary particle and all the secondaries produced by it.

the calculation of millions of additional sums and products each time a history is finished, would slow down the calculation unacceptably and should therefore be avoided.

To overcome this difficulty, most EGS4 user codes adopt a different approach, which explains the absence of a label equivalent to HIST in their PSFs. The procedure consists of splitting the simulation into a number N_B (≥ 10) of independent 'batches'. This scheme is not completely satisfactory, because the reliability of the estimated variance depends on the batch size. Indeed, it can be proved that smaller sizes yield more accurate variances, i.e. the variance of the estimated variance decreases. In fact, the optimum batch size is one history, which would make this method totally equivalent to the approach represented by equation (4) provided a label like HIST were added to the BEAM PSFs. Fortunately, the overload produced by the naive application of equation (4) can be easily avoided so that there is no need to rely on the concept of batches. A computationally inexpensive algorithm that implements equations (3) and (4) for voxel-like distributions is presented in detail in appendix A.

Another interesting subtlety originated by the division of the simulation into two parts is related with the statistical uncertainty with which q can be known for a given PSF, which represents the output of the first part. In appendix B it is proven that the use of VR techniques during the second part of the process, that is, for the transport within the phantom of the particles that form the PSF, cannot reduce the variance of \bar{q} below a certain limit which shall be called the latent variance. In the appendix we give a closed expression for this quantity, as well as a practical method to determine it. It is important to bear in mind that the latent variance is an intrinsic property of a PSF and of the nature of the quantity of interest q . It can be used to measure the 'quality' of a PSF for a particular calculation q and as a guide to optimize its size. It is also useful to measure the increase in efficiency achieved by application of a certain VR technique during the calculation of the PSF in question.

5. Results

Dose measurements were carried out in a phantom consisting of a water-filled tank with a scanning volume of $49.5 \times 49.5 \times 49.5 \text{ cm}^3$ (RFA 300 Scanditronix) and with the electron beam impinging normally on the water surface. Two detectors, a Scanditronix p-type silicon diode detector and an NACP chamber operating at a polarization potential of -300 V connected to a Farmer dosimeter 2570/1 from Nuclear Enterprise were employed. The effective point of measurement of the chamber is immediately inside its front window. The detectors were always moved from the bottom to the surface of the phantom, to avoid backlash and water surface problems. Data were collected at depth increments of 1 mm and each measurement was repeated several times. Depth doses were evaluated from depth-ionization data obtained with the ionization chamber by using stopping power ratios from the IAEA (1987) protocol, whereas the silicon diode detector was used directly, without any depth-dependent correction. The ionization chamber was used as the reference instrument and the silicon detector, once proved to yield results comparable to those obtained with the chamber in a limited depth interval, was employed instead to take advantage of its simpler operational requirements.

Simulations of 3D dose distributions in water were performed using previously generated PSFs. Central-axis depth doses were tallied within a cylindrical region of 1.2 cm in diameter, centred on the z -axis and split into bins 2 mm deep. Lateral profiles were tallied using cubic bins 2 mm wide. As the treatment head presents a nearly perfect fourfold symmetry with respect to 90° rotations around the z -axis, in the PENELOPE simulations it was assumed that the flux entering the phantom had this symmetry too and, as a result, PENELOPE's lateral profiles are symmetric with respect to their origin. By taking advantage of this fact, the standard deviation of the dose off-axis is reduced by a factor roughly equal to 2 at no additional computational cost.

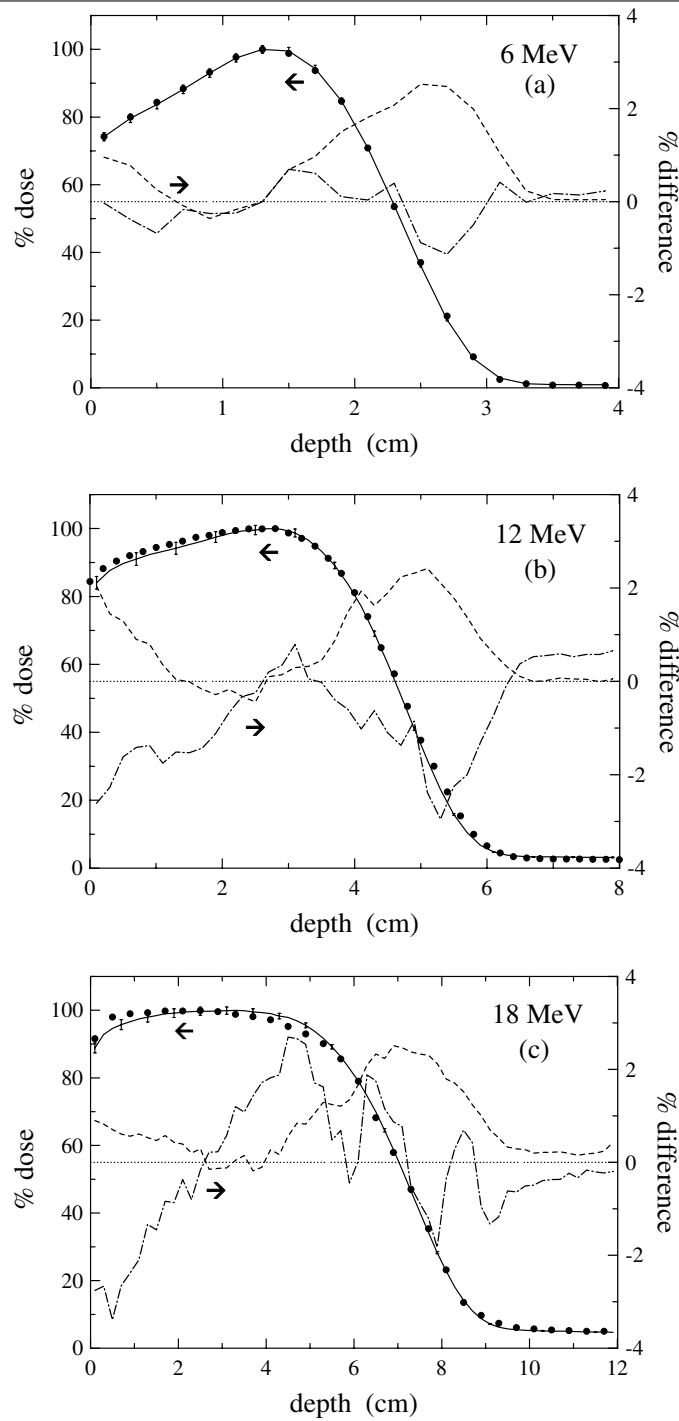


Figure 2. Central-axis depth dose in a water phantom relative to the dose maximum for the 6, 12 and 18 MeV beams. The symbols represent experimental measurements (D_{exp}) and the full curves correspond to PENELOPE results (D_{pen}). Superimposed, the differences of relative doses between PENELOPE and experiments ($D_{pen} - D_{exp}$, dot-dashed curves) and between PENELOPE and EGS4 ($D_{pen} - D_{egs}$, dashed curves) are also plotted. Error bars (at 3σ), when shown, correspond to the statistical uncertainty in the PENELOPE data.

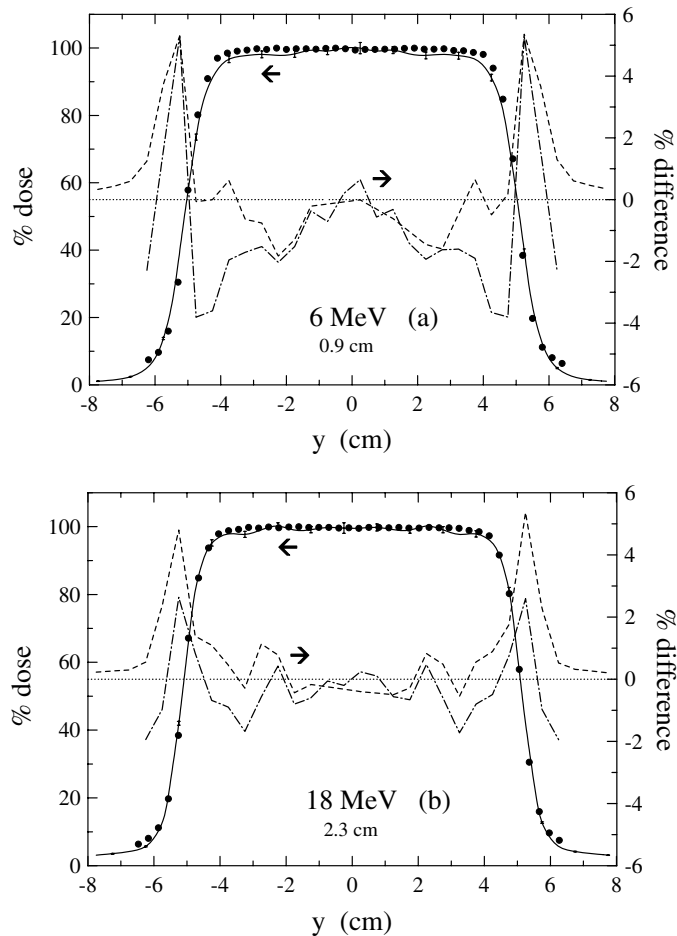


Figure 3. Lateral profiles, relative to the dose maximum, for the 6 and 18 MeV beams at depths below the water surface indicated in the insets. The values displayed refer to a line along the y -axis with $x = 0$. Symbols represent experimental measurements (D_{exp}) and the full curves correspond to PENELOPE results (D_{pen}). The differences of relative doses between PENELOPE and experiments ($D_{\text{pen}} - D_{\text{exp}}$, dot-dashed curves) and between PENELOPE and EGS4 ($D_{\text{pen}} - D_{\text{egs}}$, dashed curves) are also shown. Error bars (at 3σ), when shown, correspond to the statistical uncertainty in the PENELOPE data.

For the sake of clarity, only PENELOPE's error bars, at the 3σ level, are displayed in most figures. Since the statistical uncertainties associated with the results presented for the EGS4 calculations are smaller than those of PENELOPE, the latter represent an upper bound in most cases. This is so because, firstly, BEAM has recourse to VR techniques which were deliberately avoided in PENELOPE so that the accuracy was not compromised in any way (see section 4.2); secondly, mixed simulation cut-offs were chosen to be smaller in PENELOPE for the same reason; and finally, in some cases the EGS4 bin sizes were larger.

As explained in section 4, comparisons between PENELOPE and experiments and between PENELOPE and EGS4 will be presented separately. In figure 2, central-axis depth doses for the nominal energies of 6, 12 and 18 MeV are given. Overall, the agreement between the three sets of data is good, the differences being usually below the 2% level. Some departures from

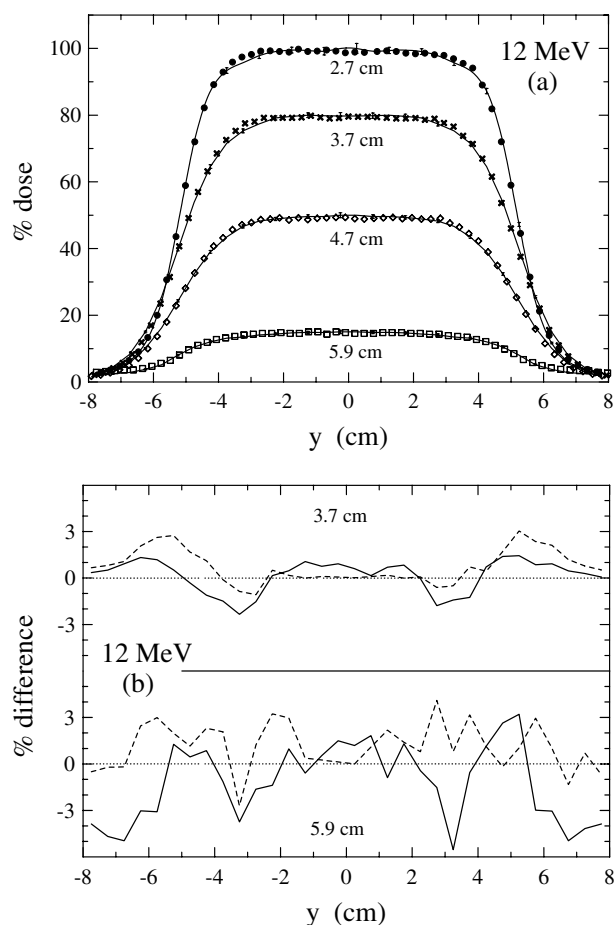


Figure 4. Lateral profiles, relative to the dose maximum, for the 12 MeV beam at depths below the water surface indicated in the inset. The values displayed refer to a line along the y -axis with $x = 0$. In graph (a), which shows the relative doses, symbols represent experimental measurements (D_{exp}) and the full curves correspond to PENELOPE results (D_{pen}). The lateral profiles at 3.7, 4.7 and 5.9 cm have been multiplied by factors 0.8, 0.5 and 0.15 respectively to avoid overlap of the different curves. Again, error bars correspond to the statistical uncertainty in the PENELOPE data. In graph (b), on the other hand, the differences of relative doses between PENELOPE and experiments ($D_{pen} - D_{exp}$, full curves) and between PENELOPE and EGS4 ($D_{pen} - D_{egs}$, dashed curves) are presented.

this value are found especially at shallow depths and for the largest energies, an effect that is probably due to the failure of the detectors to yield accurate readouts when there is a lack of electronic equilibrium. Lateral profiles for the same beams are presented in figures 3 and 4. The differences fall below 3% in most cases, being occasionally larger in the shoulders of the curves.

For the sake of completeness, it is also interesting to compare diverse features of the PSFs generated with PENELOPE and BEAM. To this end, energy and angular distributions of electrons and photons are displayed in figures 5 to 8. It is interesting to point out that whereas photon energy distributions show a very good agreement PENELOPE's electron energy distributions are significantly wider in the peak region. This difference is mainly caused by the neglect in the EGS4 code of the straggling due to energy losses below AE, the threshold




Figure 5. PSF electron energy distributions for the 6 and 18 MeV beams. Full and dashed curves represent PENELOPE and BEAM results, respectively. All curves have been normalized to unit area and error bars refer to the PENELOPE curves only.

for production of delta rays. Indeed, the agreement between the two codes can be largely improved by convolving the BEAM energy distributions with a Gaussian function whose width is determined by the straggling generated below AE and AP. Regarding angular distributions, a significant difference is also observed here, the PENELOPE results implying that fewer electrons are travelling along the forward direction, or in other words, that they experience a stronger scattering power. This difference stems from the fact that PENELOPE uses scattering powers obtained from partial wave analysis, while EGS4 is based on the Molière (1948) theory. As shown by Fernández-Varea *et al* (1993), the latter approach underestimates the scattering power for medium and large atomic numbers—such as the ones encountered in the scattering foils employed for electron beams. The angular distribution of photons, in turn, mimics that of electrons, since bremsstrahlung production is strongly correlated with the direction of flight of the emitting particle.

Of course, the importance of these divergences, which are caused by differences in the physical models implemented in both codes, have to be put in perspective by considering the relative contribution that a variation in the dynamical variable under analysis may have




Figure 6. PSF photon energy distributions for the 6 and 18 MeV beams. Full and dashed curves correspond to PENELOPE and BEAM results, respectively. All curves have been normalized to unit area.

on the actual quantity of interest, that is, on the dose. The agreement found for the latter corroborates the reliability of both codes and, at the same time, shows the limitations that the use of approximate models imposes on the degree of detail with which the transport process can be described.

An interesting feature of the photon spectra in figure 6 is the presence of several prominent peaks. The one at the lowest energy (between 50 and 100 keV) corresponds to the x-rays produced in the parts of the accelerator containing W and Au (vacuum window, scattering foil for the 18 MeV case, collimator, etc). The peak at the highest energy (around 511 keV) is caused by annihilation photons produced after the generation of an electron–positron pair by a bremsstrahlung photon. More interestingly, there is a third skewed peak between the other two in the 18 MeV case, extending slightly farther to the higher energies than to the lower ones. It corresponds to annihilation photons initially emitted upstream that interact with some element of the accelerator head and they are backscattered downstream, thus arriving at the scoring plane with an energy roughly equal to $511/(2 - \cos \pi) = 170$ keV. In the 6 MeV beam there are far fewer annihilation photons and therefore this backscattering peak is not seen.




Figure 7. Angular distributions of electrons in the PSF for the 6 and 18 MeV beams. Full and dashed curves are PENELOPE and BEAM data respectively. Distributions have been normalized to unit area when represented versus $\cos \theta$, which gives more weight to larger angles. This explains the seemingly different areas enclosed by the two displayed curves.

Experimental measurements of dose distributions with half the beam obstructed by a semi-infinite lead foil placed right under and parallel to the water surface were also carried out with the Scanditronix silicon detector. The different densities and scattering powers of the metal and the water give rise to striking dose inhomogeneities which represent a stringent test on the multiple scattering scheme adopted in the simulation algorithm. In fact, this experimental set-up is similar, although geometrically simpler, to the benchmark test described by Shortt *et al* (1986). As shown in figure 9, a very good agreement is found between the three sets of results, the difference usually being below the 3% level.

6. Conclusions

The MC code PENELOPE has been used to simulate electron beams from a Siemens Mevatron KDS linac for three nominal energies. Central-axis depth doses and lateral profiles inside a water phantom have been obtained and compared with experimental results and with




Figure 8. Angular distributions of photons in the PSF for the 6 and 18 MeV beams. Full and dashed curves are PENELOPE and BEAM results, respectively. As in the previous figure, distributions have been normalized to unit area when they are represented versus $\cos \theta$.

calculations performed with the well-known MC code EGS4 (more specifically, with user codes BEAM and DOSXYZ). Generally speaking, the agreement is good. Very close agreement has also been found between PENELOPE and EGS4, on the one hand, and between PENELOPE and measurements, on the other, when the beam was partially obstructed by a lead slab put on the water surface, a stringent test of the angular scattering mechanisms implemented in these codes.

Phase space distributions at the scoring plane for electrons and photons have been calculated with both PENELOPE and BEAM. Small but significant differences, probably caused by the way in which condensed simulation physics are implemented in the two codes, were found in the energy distribution of electrons and the angular spread of particles. These differences proved not to be very significant regarding the calculation of dose distributions inside the water phantom. Ongoing work on the extension of this conclusion to the inhomogeneous geometry of an actual patient, obtained by means of a CT scan, will be presented elsewhere.

The motivation for this study was to show that PENELOPE is well suited to describing the transport of radiation in complex geometries which consist of multiple elements with relatively

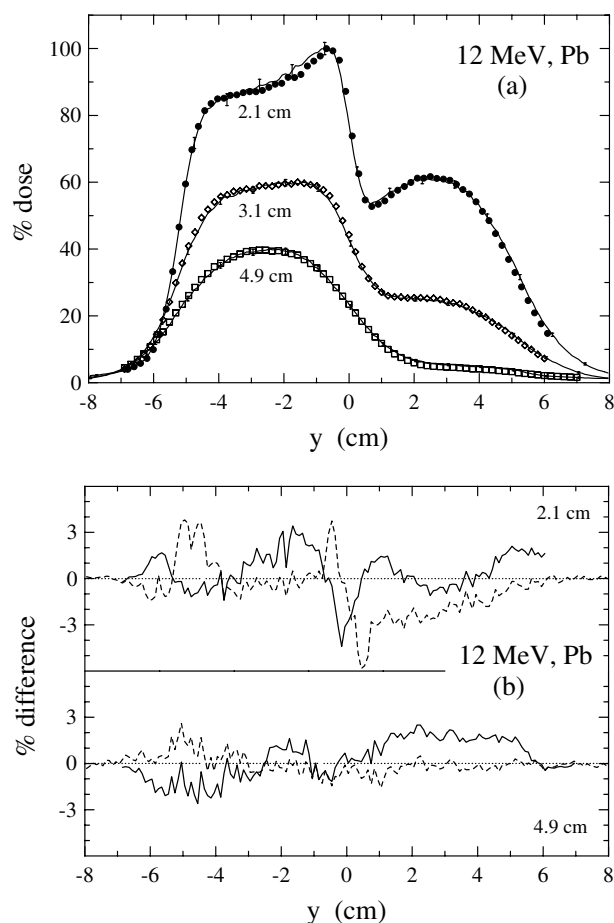


Figure 9. Lateral dose profiles relative to the dose maximum in a water phantom for the 12 MeV beam at the depths indicated in the insets. The water surface was partially covered with a lead slab—see text for details. In graph (a), which shows the relative doses, symbols represent experimental measurements (D_{exp}) and the full curves correspond to PENELOPE results (D_{pen}). The lateral profiles at 3.1 and 4.9 cm have been multiplied by factors 0.6 and 0.4, respectively, to avoid overlap. In graph (b), on the other hand, the differences of relative doses between PENELOPE and experiments ($D_{pen} - D_{exp}$, full curves) and between PENELOPE and EGS4 ($D_{pen} - D_{egs}$, dashed curves) are presented. PENELOPE's statistical uncertainties were about or less than 2% of the dose maximum in all cases.

simple shapes, usually limited by quadric surfaces, such as those found in medical problems involving accelerator heads, ^{60}Co units, brachytherapy units etc. With the aid of the debugger GVIEW, the coding of these geometries according to the PENGEOM2 format is nothing more than routine editing work. On the other hand, the code philosophy is very similar to that of EGS4, namely the subroutines of the package are called from a simple steering main program provided by the user, thus allowing full flexibility in the definition of the quantities of interest and in the application of variance-reduction techniques. The reliability of the simulation results and the confidence of the medical community on the MC method should benefit from the widening of the spectrum of available codes.

We have also introduced a robust and cost-effective method of programming the calculation of accurate statistical uncertainties for CT geometries. Additionally, it has been shown that

PSFs can be characterized by their 'latent' variance, the minimum attainable variance for a given quantity of interest. An explicit expression and a practical procedure to estimate its value are presented.

To reduce the time needed to complete the calculation, PENELOPE has been parallelized using the freely available MPI library. The speed-up with respect to the scalar code is roughly linear with the number of processors. The use of the concept of task farms, which can be based on fast and affordable state-of-the-art personal computers, puts this possibility within the reach of most medical and academic centres.

Acknowledgments

We would like to thank N Leinfellner for helping us in parallelizing PENELOPE with MPI and C-M Ma for enlightening discussions. We are also grateful to two anonymous referees for their pertinent comments and helpful suggestions and to Siemens for providing information on the structure of the accelerator.

The calculations reported in this paper have been performed using the computational facilities of the Centre de Supercomputació de Catalunya (CESCA). This work has been possible, in part, thanks to the financial support provided by the Fondo de Investigación Sanitaria (Spain), contract no 97/2111. One of the authors (SBJ) was financially supported by a grant from the US Army Breast Cancer Research Program (BC990213) and JS also acknowledges the financial support of the Fondo de Investigación Sanitaria de la Seguridad Social (Spain), contract no 98/0047-01.

Appendix A. Fast uncertainty calculation

The simultaneous evaluation, by using equations (3) and (4), of the means \bar{q} and variances s^2 associated with a large number of MC mean value estimators can be readily implemented in a computer program as follows.

For each quantity of interest q , e.g. the energy deposited in a voxel, we define three real arrays $Q(K)$, $Q2(K)$ and $QTEMP(K)$ and an integer label $LASTHI(K)$, where K is the voxel index. All these quantities are initially set to zero. The partial score q_{ij} due to the j th scoring event occurred during the simulation of the i th history is accumulated in the partial counter $QTEMP$, whereas the global contribution q_i and its square are accumulated in Q and $Q2$, respectively. Here, we recall that by 'history' it is meant a primary particle *and* all its secondaries. Each particle is assigned a label, $HIST$, equal to its history's order number i . Its value is transferred to $LASTHI$ the first time that the history contributes to $QTEMP$. In the course of the simulation, the value of $QTEMP$ is transferred to the global counters Q and $Q2$ only when it is necessary to store a contribution from a new history. Explicitly, the FORTRAN code for scoring a contribution $DELTAQ \equiv q_{ij}$ to the voxel K is

```

IF (HIST.NE.LASTHI(K)) THEN
  Q(K) = Q(K)+QTEMP(K)
  Q2(K) = Q2(K)+QTEMP(K)**2
  QTEMP(K) = DELTAQ
  LASTHI(K) = HIST
ELSE
  QTEMP(K) = QTEMP(K)+DELTAQ
ENDIF

```

At the end of the simulation, the residual contents of $QTEMP(K)$ must also be transferred to the global counters $Q(K)$ and $Q2(K)$ for all K .

Note that this algorithm has a very modest computational burden, since it only implies an additional IF each time a scoring event takes place and the introduction of the integer array LASTHI(K) for each existing counter Q.

Appendix B. Latent variance

In this appendix we address the problem of computing the statistical uncertainties that originate from the fact that the PSF has a limited size. For the sake of brevity, the equations that follow are given without proof. A paper describing these and other aspects of this concept is under preparation and will be published elsewhere.

The simulation process we deal with can be stated in the following general terms. The calculation of some quantity of interest q , e.g. the dose deposited in a voxel, is split into two parts. First, the trajectory of each particle is simulated until it reaches a predefined scoring plane. This part of the process gives rise to the PSF. Second, the transport of each particle in the PSF continues until it has been completed and a contribution to q is obtained.

Let us assume that the phase space of particles on the scoring plane is divided into a set of infinitely small portions, which we shall call cells. Each cell can be thought of as defining whether the particle is an electron, a photon or a positron, and an interval for its dynamical variables, i.e. its energy, position and direction of flight. As a result of the completion of a history (i.e. the simulation of a primary and all its secondaries) up to the scoring plane, a random variable n_b can be assigned to cell b ($b = 1, \dots$) such that, for analogue simulations, its value is the number of particles that happen to end up in that cell. In general, when VR techniques are employed, n_b is the sum of statistical weights of the particles occupying the cell, if any. We introduce a second random variable, q_b , defined as the contribution to the quantity of interest produced by simulating a particle with a set of dynamical variables defined by cell b and with unit statistical weight.

It can be proved that, owing to the mutual independence of the q_b s and the independence of q_b with respect to n_b , the variance of the estimator \bar{q} , given by equation (3), can be written as

$$\sigma^2(\bar{q}) = \frac{1}{N}(A + B) \quad (\text{B.1})$$

where N is the number of simulated histories and the quantities A and B are given by

$$A = \sum_b \langle q_b \rangle^2 \langle n_b^2 \rangle + \sum_{a \neq b} \langle q_a \rangle \langle q_b \rangle \langle n_a n_b \rangle - \langle q \rangle^2 \quad (\text{B.2})$$

and

$$B = \sum_b \sigma^2(q_b) \langle n_b^2 \rangle \quad (\text{B.3})$$

respectively. The factor $\sigma^2(q_b)$ is the variance associated with q_b . The key point here is that the use of VR techniques *downstream* of the scoring plane can reduce $\sigma^2(q_b)$, and hence B , to virtually zero. However, neither $\langle q_b \rangle$ nor n_b can be altered during this part of the simulation. Therefore, according to expressions (B.1) and (B.2), for a given PSF A/N is the *minimum* variance with which \bar{q} can be obtained; hereafter and in the text, this quantity will be referred to as the latent variance.

The latent variance can be numerically evaluated by means of a simple trick. It is easy to see that a K -fold splitting of the particles in the PSF, weighting them with a factor K^{-1} so as to keep the final scores unbiased, modifies the variances $\sigma^2(q_b)$ and B by the same factor K^{-1} . Thus, the variance $\sigma_K^2(\bar{q})$ of the simulation result becomes

$$\sigma_K^2(\bar{q}) = \frac{1}{N}(A + BK^{-1}). \quad (\text{B.4})$$




Figure 10. Variance of the average dose produced by the 12 MeV beam in a water phantom at depths between 2.7 and 2.8 cm, where the dose maximum lies. MC variances using various splitting factors K are displayed as dots. The straight line is a fit performed with the last four data points on the right. The PSF contained 20 000 particles corresponding to $N = 224\,948$ primaries. Error bars correspond to the uncertainty of the variance.

This expression can be used to determine the latent variance as follows. We use particle splitting to do the simulation after the scoring plane of the particles in the PSF. Various simulations are performed with increasingly larger splitting factors K . Then, the variance of \bar{q} is represented versus K^{-1} and fit to a straight line. The crossing of this line with the axis $K^{-1} = 0$ defines the latent variance. Owing to the fact that A is an intrinsic property of the system (and of the simulation algorithm used), this procedure can be applied to only a fraction of the PSF and the result scaled according to the actual value of N .

The results of the application of this method to a practical case are displayed in figure 10. The considered quantity q is the average dose, at depths between 2.7 and 2.8 cm, caused by the 12 MeV electron beam from our accelerator in a semi-infinite water phantom. The dose was calculated from a deliberately limited PSF with 20 000 particles, which correspond to $N = 224\,948$ simulated primary electrons (i.e. histories). The ‘measured’ MC variances $\sigma_K^2(\bar{q})$ are obtained from simulations using particle splitting with different splitting factors K . When plotted as a function of $1/K$, the MC variances do define a straight line, which confirms the adequacy of equation (B.4). The straight line shown in the figure was obtained by least squares fitting of the MC variances for $K = 1, 2, 4$ and 8 . The latent standard deviation obtained from this fit was

$$\left(\frac{A}{N}\right)^{1/2} = 0.75 \pm 0.05 \text{ keV cm}^2 \text{ g}^{-1} \text{ per history.}$$

This result can now be scaled to the actual PSF size, say $N' = 10^8$ histories, yielding

$$\left(\frac{A}{N'}\right)^{1/2} = 0.036 \pm 0.002 \text{ keV cm}^2 \text{ g}^{-1} \text{ per history.}$$

which is less than 0.1% of the dose at the considered depth. Other quantities of interest involving a lower interaction probability per history, such as the dose to a single voxel, are bound to have larger latent variances.

References

- AAPM TG-21 1983 A protocol for the determination of absorbed dose from high-energy photon and electron beams *Med. Phys.* **10** 741–71
- Andreo P 1991 Monte Carlo techniques in medical radiation physics *Phys. Med. Biol.* **36** 861–920
- Baró J, Sempau J, Fernández-Varea J M and Salvat F 1995 PENELOPE: an algorithm for Monte Carlo simulation of the penetration and energy loss of electrons and positrons in matter *Nucl. Instrum. Methods B* **100** 31–46
- Berger M J 1963 Monte Carlo calculation of the penetration and diffusion of fast charged particles *Methods in Computational Physics* vol 1, ed B Alder, S Fernbach and M Rotenberg (New York: Academic) pp 135–215
- Bielajew A F 1995 HOWFAR and HOWNEAR: geometry modelling for Monte Carlo particle transport *Technical Report PIRS-0341* (Ottawa: National Research Council of Canada)
- Bielajew A F and Rogers D W O 1987 PRESTA: The parameter reduced electron-step transport algorithm for electron Monte Carlo transport *Nucl. Instrum. Methods B* **18** 165–81
- Briesmeister J F 1997 MCNP—a general Monte Carlo *N*-particle transport code *Technical Report LA-12625-M Version 4B* (Los Alamos, NM: Los Alamos National Laboratory)
- Ding G X, Rogers D W O and Mackie T R 1996 Mean energy, energy-range relationship and depth-scaling factors for clinical electron beams *Med. Phys.* **23** 361–76
- Fernández-Varea J M, Mayol R, Baró J and Salvat F 1993 On the theory and simulation of multiple elastic scattering of electrons *Nucl. Instrum. Methods B* **73** 447–73
- Gnedenko B V 1975 *The Theory of Probability* (Moscow: Mir)
- Goudsmit S and Saunderson J L 1940 Multiple scattering of electrons *Phys. Rev.* **57** 24–9
- Halbleib J A, Kensek R P, Mehlhorn T A, Valdez G D, Seltzer S M and Berger M J 1992 ITS version 3.0: the integrated TIGER series of coupled electron/photon Monte Carlo transport codes *Technical Report SAND91-1634* (Sandia National Laboratories)
- IAEA 1987 Absorbed dose determination in photon and electron beams; an international code of practice *Technical Report 277* (Vienna: IAEA)
- Jiang S B, Kapur A and Ma C-M 2000 Electron beam modelling and commissioning for Monte Carlo treatment planning *Med. Phys.* **27** 180–91
- Lewis H W 1950 Multiple scattering in an infinite medium *Phys. Rev.* **78** 526–9
- Lewis R D, Ryde S J S, Hancock D A and Evans C J 1999 An MCNP-based model of a linear accelerator x-ray beam *Phys. Med. Biol.* **44** 1219–30
- Ma C-M 1992 Monte Carlo simulation of dosimeter response using transputers *PhD Thesis* Institute of Cancer Research and Royal Marsden Hospital
- Ma C-M, Faddegon B A, Rogers D W O and Mackie T R Accurate characterization of Monte Carlo calculated electron beams for radiotherapy *Med. Phys.* **24** 401–16
- Ma C-M and Jiang S B 1999 Monte Carlo modelling of electron beams from medical accelerators *Phys. Med. Biol.* **44** R157–R189
- Ma C-M, Reckwerdt P, Holmes M, Rogers D W O and Geiser B 1995 DOSXYZ users manual *Technical Report PIRS 509b* (Ottawa: National Research Council of Canada)
- Mascagni M, Cuccaro S A, Pryor D V and Robinson M L 1995 A fast, high quality, and reproducible lagged-Fibonacci pseudorandom number generator *J. Comput. Phys.* **15** 211–19
- Mohan R 1988 Monte Carlo simulation of radiation treatment machine heads *Monte Carlo Transport of Electrons and Photons* ed T M Jenkins, W R Nelson and A Rindi (New York: Plenum) pp 453–68
- Mohan R, Chui C S and Lidofsky L 1985 Energy and angular distributions of photons from medical linear accelerators *Med. Phys.* **12** 592–7
- Molière G 1948 Theorie der Streuung schneller geladener Teilchen II: Mehrfach- und Vielfachstreuung *Z. Naturforsch.* **3a** 78–97
- Nelson W R, Hirayama H and Rogers D W O 1985 The EGS4 code system *Technical Report SLAC-265* (Stanford, CA: Stanford Linear Accelerator Center)
- Petti P L, Goodman M S, Sisterson J M, Biggs P J, Gabriel T A and Mohan R 1983 Sources of electron contamination for the Clinac-35 25-MV photon beam *Med. Phys.* **10** 856–61
- Rogers D W O, Bielajew A F, Mackie T R and Kubsad S S 1990 The OMEGA project: treatment planning for electron-beam radiotherapy using Monte Carlo techniques *Phys. Med. Biol.* **35** 285–6
- Rogers D W O, Faddegon B A, Ding G X, Ma C-M, Wei J and Mackie T R 1995 BEAM: a Monte Carlo code to simulate radiotherapy treatment units *Med. Phys.* **22** 503–24
- Salvat F, Fernández-Varea J M, Baró J and Sempau J 1996 PENELOPE, an algorithm and computer code for Monte Carlo simulation of electron-photon showers *Technical Report 799* (Madrid: Ciemat)

- Schach von Wittenau A E, Cox L J, Bergstrom P M Jr, Chandler W P, Hartmann Siantar C L and Mohan R 1999 Correlated histogram representation of Monte Carlo derived medical accelerator photon-output phase space *Med. Phys.* **26** 1196–211
- Sempau J, Acosta E, Baró J, Fernández-Varea J M and Salvat F 1997 An algorithm for Monte Carlo simulation of coupled electron–photon transport *Nucl. Instrum. Methods B* **132** 377–90
- Shortt K R, Ross C K, Bielajew A F and Rogers D W O 1986 Electron beam dose distributions near standard inhomogeneities *Phys. Med. Biol.* **31** 235–49
- Siebers J V, Keall P J, Libby B and Mohan R 1999 Comparison of EGS4 and MCNP4b Monte Carlo codes for generation of photon phase space distributions for a Varian 2100C *Phys. Med. Biol.* **44** 3009–26
- Udale M 1988 A Monte Carlo investigation of surface doses for broad electron beams *Phys. Med. Biol.* **33** 939–54
- Udale-Smith M 1992 Monte Carlo calculations of electron beam parameters for three Philips linear accelerators *Phys. Med. Biol.* **37** 85–105