BEAM: A Monte Carlo code to simulate radiotherapy treatment units

D. W. O. Rogers,[a] B. A. Faddegon,[b] G. X. Ding,[c] C.-M. Ma, and J. We[c]
Ionizing Radiation Standards, Institute for National Measurement Standards, National Research Council Canada, Ottawa K1A 0R6, Canada

T. R. Mackie
University of Wisconsin, Department of Medical Physics, Madison, Wisconsin

(Received 24 August 1994; accepted for publication 28 December 1994)

This paper describes BEAM, a general purpose Monte Carlo code to simulate the radiation beams from radiotherapy units including high-energy electron and photon beams, ^{60}Co beams and orthovoltage units. The code handles a variety of elementary geometric entities which the user puts together as needed (jaws, applicators, stacked cones, mirrors, etc.), thus allowing simulation of a wide variety of accelerators. The code is not restricted to cylindrical symmetry. It incorporates a variety of powerful variance reduction techniques such as range rejection, bremsstrahlung splitting and forcing photon interactions. The code allows direct calculation of charge in the monitor ion chamber. It has the capability of keeping track of each particle's history and using this information to score separate dose components (e.g., to determine the dose from electrons scattering off the applicator). The paper presents a variety of calculated results to demonstrate the code's capabilities. The calculated dose distributions in a water phantom irradiated by electron beams from the NRC 35 MeV research accelerator, a Varian Clinac 2100C, a Philips SL75-20, an AECL Therac 20 and a Scanditronix MM50 are all shown to be in good agreement with measurements at the 2 to 3% level. Eighteen electron spectra from four different commercial accelerators are presented and various aspects of the electron beams from a Clinac 2100C are discussed. Timing requirements and selection of parameters for the Monte Carlo calculations are discussed.

Dedication: This paper is dedicated to the memory of our friend and colleague, Jiansu Wei, who made a significant contribution to this project before he passed away on March 15, 1993.

TABLE OF CONTENTS

 0094-2405/95/22(5)/503/22/$6.00

I. INTRODUCTION

The OMEGA project is a collaborative project to develop a full 3-D electron beam treatment planning system based on Monte Carlo simulation techniques to calculate the dose to the patient.[1,2] The rationale is that analytic methods of calculating dose in an electron beam suffer from various short comings which leave residual uncertainties or errors of 10% or more using 3-D implementations of pencil beam algorithms and much worse using the more common 2-D implementations.[3,4] Advanced analytic models may remove some of these problems,[5–7] sometimes at the expense of very considerable computation time. However, in the end, the Monte Carlo technique for dose calculation is the gold standard, capable, in principle, of accurately computing the dose under almost all circumstances.[8–10] In particular Monte Carlo techniques can handle backscatter from prostheses and bone, or scatter perturbations by air cavities much more accurately than any of the other current models.

The major drawbacks of the Monte Carlo technique are computing time and the need for a detailed knowledge of the incident radiation beam. The actual code to compute the dose in the patient is relatively simple and a completely general code capable of using CT data in each voxel was written in 1987 and has been used as a standard timing benchmark for the EGS4 code on a large number of computers.[9,11] Using this code it was shown that a realistic treatment plan for a 20 MeV electron beam with ± 1% statistical uncertainty in voxels with 2.5 mm sides would take about 600 hours on a VAX11/780, although the actual timing was dependent on a large number of parameters such as field size, voxel size, beam energy and accuracy required.[9] In the mid-90s, computers in many radiotherapy centers are tens or even hundreds of times faster than the VAX11/780 and the Monte Carlo code can be run in parallel on several machines, thus reducing the brute force calculation to a few hours. We have also been investigating ways to speed up the calculation using variance reduction techniques such as correlated sampling,[12] which provides a definite speed up in most situations, or use of the macro Monte Carlo technique,[13] which is capable of increases in computing speed of about a factor of 10. With future increases in computing speed and the development of routine procedures for farming out calculations to large numbers of computers, the Monte Carlo calculation of dose delivered to a patient will be achievable in a few minutes. This will still be longer than required for advanced analytic techniques, but, at minimum, the Monte Carlo technique will be of direct clinical use for "final" dose calculations or verification of other, simpler techniques which will be used in optimization or other exploratory calculations.

The problem addressed in this paper is that of calculating accurate information about the incident electron beam in order to initiate the transport of particles in the patient model. In another paper we address the issue of how to characterize these beams with sufficient accuracy for routine use.[14] It has been shown that to match accurately the measured central-axis depth-dose curves in a water phantom irradiated by "clean" 10 and 20 MeV electron beams,[8] it is essential to include the complete energy distributions of the electron beam rather than just using the correct average energy from a point source.[9] Udale has further shown the need to include the correlation between energy and angle in the incident electron beam if good agreement is to be obtained on the central axis.[15–17] Our results confirm this result[14] although Udale's work concentrated on the 10 MeV beam from an SL75-20 which represents the worst case we have studied, and in general angular effects are small on the central axis (up to about 3% near the surface and less at depth). However, if the doses in the wings or penumbral regions are to be calculated accurately, we find one needs a detailed knowledge of the spatial distribution of electron fluence and to a lesser extent, the angular information about the beam. Other groups have demonstrated that Monte Carlo calculations can be done in 3-D patient models[18–20] but they have always used very simple models of the incident beam. In some cases this may be accurate, but in general it is not, and full models of the beam are necessary to assess the accuracy of the simple models.

The major emphasis in this project has been on simulating electron beams. These are more difficult to simulate accurately than photon beams. However, the code BEAM is equally applicable to photon beams from accelerators, ^{60}Co beams or even low-energy x-ray beams and a variety of capabilities related specifically to photon beams have been included. The code is also a very versatile, general purpose Monte Carlo transport package which can be used in a wide variety of applications other than simulating radiotherapy beams.

As well as being useful as the first step in accurate 3-D treatment planning,[21,22] detailed information about radiotherapy beams has a wide variety of applications in clinical physics and radiation dosimetry. For example, Andreo *et al.*[23] have shown that the energy and angular distributions of electron beams affect the stopping-power ratios needed for reference dosimetry. Detailed simulations using realistic beams allow a more systematic study of these effects.[24] Calculated information on photon beam characteristics can be very useful for a wide variety of photon dosimetry problems such as studies of electron contamination in photon beams.[25,26]

A. Previous Monte Carlo models of radiotherapy beams

Application of Monte Carlo techniques to calculations of radiotherapy beams has a long history. ICRU Report 18 made extensive use of various calculations of the spectra from ^{60}Co sources.[27] More recently various groups have done sophisticated simulations of both ^{60}Co units[28–30] and photon beams from linear accelerators[25,26,31–34] although in all cases these studies were restricted to models with cylindrical symmetry. Recent studies have shown that neither the EGS nor the ITS (ETRAN) codes can be used in default mode to calculate the angular distribution of bremsstrahlung photons near 0°.[35–37] Although both codes can give accurate results, the default versions have roughly 40% errors which are related to different problems (see section II.H). Thus any previous calculations must be repeated if the detail of these angular distributions plays a role.

The more difficult task of simulating electron beams has

been studied by Udale.[15–17] Using the EGS4 code, she did very extensive modelling which was no longer restricted to cylindrical symmetry. Udale studied 3 different accelerators and although her code was well structured, it could not be applied to other accelerators without a large effort. One interesting feature of her code's design was that it considers a series of more or less independent slabs, each slab containing a separate accelerator component. We have extended this approach.

B. Design philosophy

To meet the goals of the OMEGA project, the BEAM code has been designed to simulate the radiation beams from any radiotherapy source, including low-energy x-rays, ^{60}Co units and both electron and photon beams from accelerators.

BEAM is based on the PRESTA extension[38] of the EGS4 Monte Carlo system for simulating radiation transport.[39,40,9] This code system is extensively documented and benchmarked. EGS4 can be tailored to produce the outputs required and to implement a wide variety of variance reduction techniques. The code is written in MORTRAN3, a Fortran77 preprocessor which is used for the EGS4 system.[41] Extensive use is made of the powerful macro capability of this preprocessor. The BEAM system runs under the Unix operating system using nearly 2000 lines of scripts.

One element of the basic design philosophy is that the BEAM code produces a phase-space output of the beam (i.e., the position, energy, direction, charge and history tag for each particle) at any specified plane in the model. This phase-space file can either be re-used by the BEAM code itself (see below), used as an input to the patient dose computation algorithm directly, analyzed using various data analysis programs (such as the CERN package called PAW[42]) or used to characterize the beam in a more compact way.[14] This approach retains complete flexibility without requiring a definitive a priori model to characterize the beam, albeit at the expense of large data files ($\approx$100 Mbytes is common even in a compressed format). Once a beam characterization approach is developed, it may be added to the BEAM code so that the large data files may be avoided.

Another element of the design philosophy of BEAM is that the model is built up from a series of individual component modules (CMs), each of which operates completely independently of the other component models and occupies a slab at right angles to the beam axis. Each CM deals with a specific class of geometric shape within the horizontal band, e.g., a simple set of slabs or a complex set of nested truncated cones. A CM does not specify explicit values related to the geometric shape (such as the number of slabs, their thickness or the materials they are made with) but allows all such possibilities and handles input of all the necessary data for a specific run. The independence of each component of the geometry package allows each CM to be tested and debugged in isolation. It also allows many very different designs of machines to be simulated because the various CMs can be used in a wide variety of configurations.

The first step in modeling a particular machine is to "specify" the machine by selecting, in order, the CMs to be used and also associating a unique name with each so that a particular CM can be used more than once. For example, a pair of collimator jaws and a pair of electron applicator scrapers may have the identical geometric requirements and thus they can both use the same CM coding. However the coding is used twice with a unique name attached to each use. Once the machine is "specified," the Unix script "builds" the machine by pulling together all the relevant source code and automatically editing it appropriately to avoid redundant variable names and then compiling the code. Once compiled, this "machine" model can still model a wide variety of configurations since each simulation reads a user-supplied input file which contains all the specific information related to this run (e.g., the initial electron energy distribution, the dimensions and materials of all parts of the machine, the location of output planes, how to track a particle's history, which variance reduction techniques to apply, what transport parameters to use, etc.). This file can be interactively created. This approach was adopted because it maintains the ability to add new CMs and model new configurations easily, at the expense of having a very complex series of scripts to run the code. The other alternative was to write a single general purpose code which allowed any machine to be specified solely from a user-supplied input file. While this might be easier to use, it would inevitably require compromises in generality if unexpected components were required to simulate a new machine (a common experience) and it would be much harder to debug the geometry. Furthermore, by automating our scripts one step further (a straightforward task), we could achieve the goal of having a system driven by one input file. The machine specification and run-time input files would be concatenated and the accelerator code recompiled for each run. Figure 1 shows the overall structure of the BEAM system and the steps involved in using it.

Another element of the design philosophy was to make the code so it could be used accurately by others. We have made a significant effort to document the code well, especially in the source code. For example, all variables and input requirements are defined within the source code and scripts are available to extract stand-alone files containing these descriptions. This was essential since for each machine "specified" the input file could be different. The code also outputs a file which can be used to produce a 3-D display of the machine geometry and particle histories using EGS_windows, a standard graphics package for EGS4 simulations.[43] This feature is very important for verifying the input of any user-supplied geometry.

II. BEAM CAPABILITIES

A. General characteristics

The BEAM code models the therapy source with the z-axis taken as the beam-axis and usually the origin is defined as the center of the beam as it exits from the accelerator vacuum. The model consists of a series of component modules (CMs), each of which is contained between two planes which are perpendicular to the z axis and which can not overlap.

There can be an arbitrary number of scoring planes which are at the back plane of a CM and thus perpendicular to the

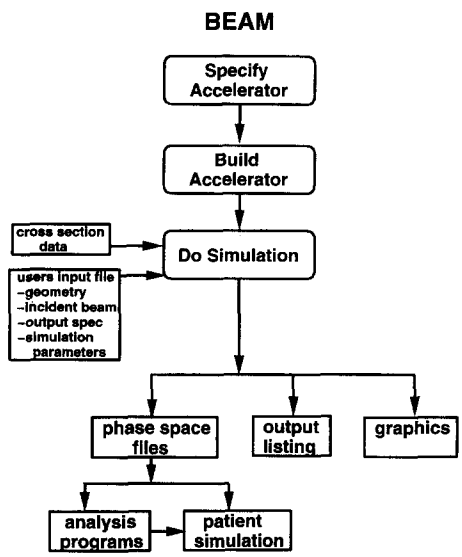

FIG. 1. The steps involved in using the BEAM system. To specify an accelerator means to define an ordered set of component modules (CMs) to be used in the simulation. To build an accelerator consists of gathering all the source code, automatically editing it to avoid duplicate names and compiling the resulting Mortran and fortran code. To do the simulation, BEAM reads in a cross section data set [prepared by the PEGS package from the EGS4 system (Ref. 39)] and an input file which contains all the information related to that specific run: the complete geometry; the initial beam (either from the accelerator vacuum or from a previous run); the output required from this run; and the simulation control parameters (cutoffs, forcing regions, etc.). BEAM outputs phase-space data and graphics if requested and produces an output listing. The phase-space data can be used in off-line analysis programs or fed directly into patient or other simulation codes.

z-axis. At these planes the particle fluences, average energies and average angles are scored in an arbitrary number of circular rings or concentric square regions of arbitrary width, and a phase-space data file can be created at each plane.

Dose can be scored in an arbitrary number of regions, and the user can arrange to average the dose in an arbitrary group of regions. These regions may span several CMs.

As well as a selection of standard source routines (pencil beam, point source, circular or rectangular beams), the simulations can be started using a phase-space data file as input. This allows the user, e.g., to simulate the head of an accelerator down to the applicator and then reuse the same data with a variety of different applicators in place. When a phase-space file is re-used, to prevent double counting, particles which have crossed the scoring plane previously are not used. Re-use of the phase-space data also allows the simulation of a ^{60}Co capsule in a separate calculation and then the use of the phase-space data file as an input to a model of the ^{60}Co therapy unit.

BEAM has a restart facility which allows a run with inadequate statistics to be restarted, or a run which was terminated before completion for some reason, to have the data analyzed. This is done by writing all scoring results to disk at the end of each "batch" which is normally taken as 1/10-th of the run. As well, full information about the state of the random number generator is stored at this point. Based on these data, the run can be re-started or analyzed up to the end of the last completed batch. Information about how many records are written to the phase-space file is also written out so that the phase-space data can be properly added to from the end of the batch (for the case of an interrupted run).

For all scored quantities, a statistical estimate of the uncertainty on the quantity is given by breaking the calculation into an arbitrary number of separate runs (usually 10) and then calculating the mean and variance of the mean to estimate the uncertainty (the method is discussed in detail in Ref. 9). To allow for a restart facility, each batch in the run is broken into multiple separate calculations so that the statistical analysis is available after each batch.

The CMs are not allowed to overlap each other and any user request to do so causes an error message.

Each CM prepares output which can be used by the EGS_windows package to display the geometry. The actual display is set up to help understand an accelerator simulation rather than give an accurate representation of the physical structure. For example, when a primary collimator is simulated by a truncated cone, all that is displayed is the inner surface of the cone, not the exterior.

B. Geometric Considerations

A recurring problem in Monte Carlo codes is numerical roundoff problems as particles cross boundaries. The problem is that one must calculate the distance to the boundary and if numerical accuracy problems cause the particle to be transported just short of the boundary, or too far, the tracking algorithm can get confused. To avoid this, we have used various techniques such as: making obvious assumptions in some cases (e.g., when crossing a plane); or using "boundary buffers" which cause the boundary to always be slightly further away than it really is so that the calculation will always ensure the particle is just across the actual boundary. The tolerances are set to be reasonable for an accelerator simulation, but may fail if much smaller geometries are being simulated.

C. Component Modules (CMs) Available

The code has been written as a series of independent component modules. These are described in detail in a user's manual which is being developed with specific information about input requirements and details about restrictions. There is also a report being developed which describes the detailed specifications for writing a CM and doing the quality assurance on the CM. The following brief descriptions are meant to give an indication of the range of geometries which can be handled currently. For each CM, all the physical dimensions and materials involved are set by the user at run time. Thus, in one use of a given CM the central region may be air with lead shielding on the outside, e.g., to represent jaws, and in another application the central core may be lead and the outer region air to model a flattening filter or bolus. In most cases the CM extends to either an outer cylinder or a square box. The outer boundary is set separately for each CM and is constant for the CM. The CMs do not have to be contiguous, and any gaps between them are automatically filled with air or any other default medium.

The CMs have been named after the components they were originally written to model, but they can be applied to many more structures.

1. SLABS

This is the simplest CM but has many applications. It models parallel slabs in the $x-y$ plane. They can have different materials and be of arbitrary thickness. The outer boundary is square.

2. CONS3R

Figure 2(a) shows the component module CONS3R which models a stack of truncated cones (which can also be cylinders) treated as three regions. It is useful for modeling flattening filters where the inner region is a heavy material and the outer region is air, or conversely, for modeling a conical collimator where the central region is air. It can be used for many situations if there is cylindrical symmetry and if there are only two radial regions. If the cones are specified by a set of paired depth and radius values, the depth values are constrained to increase constantly, i.e., they can not "double" back on themselves in the Z direction but they may in the radial direction. The outer boundary is a cylinder.

Although the CONESTAK and FLATFILT CMs can handle this geometry as a subset of more general cases, this particular CM executes much more quickly because there are no internal boundaries within media for particles to cross and range rejection (see Sec. II.G.1) is more efficient in large regions.

3. CONESTAK

Figure 2(b) shows the CONESTAK CM which models a series of stacked truncated cones (which can be cylinders). A cylindrical wall of arbitrary material surrounds the entire CM, just inside its outer cylindrical boundary. The radii of the cones must not decrease as the depth increases. The inner and outer conical regions of each layer have their materials specified separately.

This CM can be used to model scattering foils, primary collimators and many other components. This CM runs more slowly than CON3R because there are far more boundaries for a complex case. However, it is applicable in many more situations than CONS3R. It is also a subset of FLATFILT, but the required inputs are somewhat simpler to handle.

4. FLATFILT

Figure 2(c) shows the FLATFILT CM which is the most complex CM and is designed to handle some of the very complex beam flattening filters required in photon beam simulations. It is basically a generalization of CONESTAK to an arbitrary number of truncated conical sections on each layer and with the removal of the restriction about radii increasing with Z. In particular, FLATFILT can handle a flattening filter inside a conical collimator and the central component of the flattening filter can consist of a different material from the outside component.

5. CHAMBER

This CM models an ionization chamber and many other structures with cylindrical symmetry and centered on the beam axis [Fig. 2(d)]. For a monitor chamber, both the front and back walls may have several layers of different materials. The side walls consist of three layers, a chamber wall, a gap and a container wall. The central part of the chamber may consist of several air cavities separated by layers of other materials (serving as electrodes). This CM is also useful for central-axis depth-dose calculations in a phantom. Because there are fewer boundaries for the particles to cross in the side walls of this CM (i.e., the regions adjacent to the dose-scoring regions) than in other cylindrical-planar or voxel geometries, the computing efficiency can be greatly improved by using the electron range rejection technique (see section II.G.1). A major feature of this CM is that it allows the scoring of dose which has been broken down into components based on the settings of LATCH (see Sec. II.D on LATCH and discussion of Fig. 12).

6. PYRAMIDS

Figure 2(e) shows the PYRAMIDS component module which models a series of stacked, truncated pyramids. In this figure, the central region is the air gap which is partially indicated by the dotted lines. The four faces of each layer are parallel to the x or y axes but can have arbitrary orientations specified by the positive and negative x and y values of the front and back edges of each face (XFN, YBP, etc.). The layers must have an air gap of at least 0.01 cm between them. Each layer can have different materials while the material inside the pyramid defaults to air.

The materials extend to the square outer boundary of the CM. This CM may be used for modeling rectangular jaws or collimators or beveled electron applicators.

7. APPSQ

Figure 2(f) shows the APPSQ component module which models a series of square applicators or scrapers. There can be air on the outside of the scrapers. The inner faces are all parallel to the z-axis as well as their respective x and y axes. The outer boundary of this CM is a square. There is an air gap of at least 0.01 cm required between each layer of applicator and surrounding materials.

This CM may be used for modeling the square scrapers found in electron beams although it does not allow for a beveled edge. This CM allows for leakage outside the scrapers, unlike the PYRAMIDS CM.

8. JAWS

Figure 2(g) shows the JAWS component module which models pairs of opposing flat surfaces which are perpendicular to the x or y axis. The X-jaws are defined as those which move in the direction of the x-axis. The jaw surfaces may be at an arbitrary angle with respect to the z-axis. This CM is primarily intended to model the movable collimator jaws in an accelerator. One simplification is that the upper and lower surfaces are perpendicular to the axis of the beam whereas in some clinical accelerators the jaws rotate about the source to

maintain a sharp edge to the beam. This simplification is not expected to have any effect since the front face can be properly rotated. If found to be critical, thin JAWS above and below each main jaw could be used to model these surfaces accurately .

The outer boundary of the accelerator is square for the JAWS component module. It can not have jaws on all 4 sides at once, but can be set up with multiple sets of jaws in the x or y direction, one set behind the next. Each side of the jaw can be positioned independently and hence can be used to model a wedge.

9. MIRROR

Figure 2(h) shows the MIRROR CM used to simulate mirrors in the accelerator. The mirror is assumed to be one or more flat layers of different materials. It is taken to be uniform in the y-direction and varies with x (i.e., it is like the face of an x-jaw with multiple layers). Although accelerator mirrors are usually centered in the beam, the flat surface can be positioned arbitrarily in the x direction. It always extends to the outer y-boundary. The outer boundary is square.

10. XTUBE

This CM is needed to model an x-ray tube as the photon source and thus must be the first CM in any given simulation configuration. It simulates an x-ray target (possibly several layers) in air or vacuum backed by a target holder (backing material). The target angle is fixed by the target surface and z-axis [see Fig. 2(i)]. Particles reaching the boundaries of the CM are assumed to be absorbed by the shielding except for

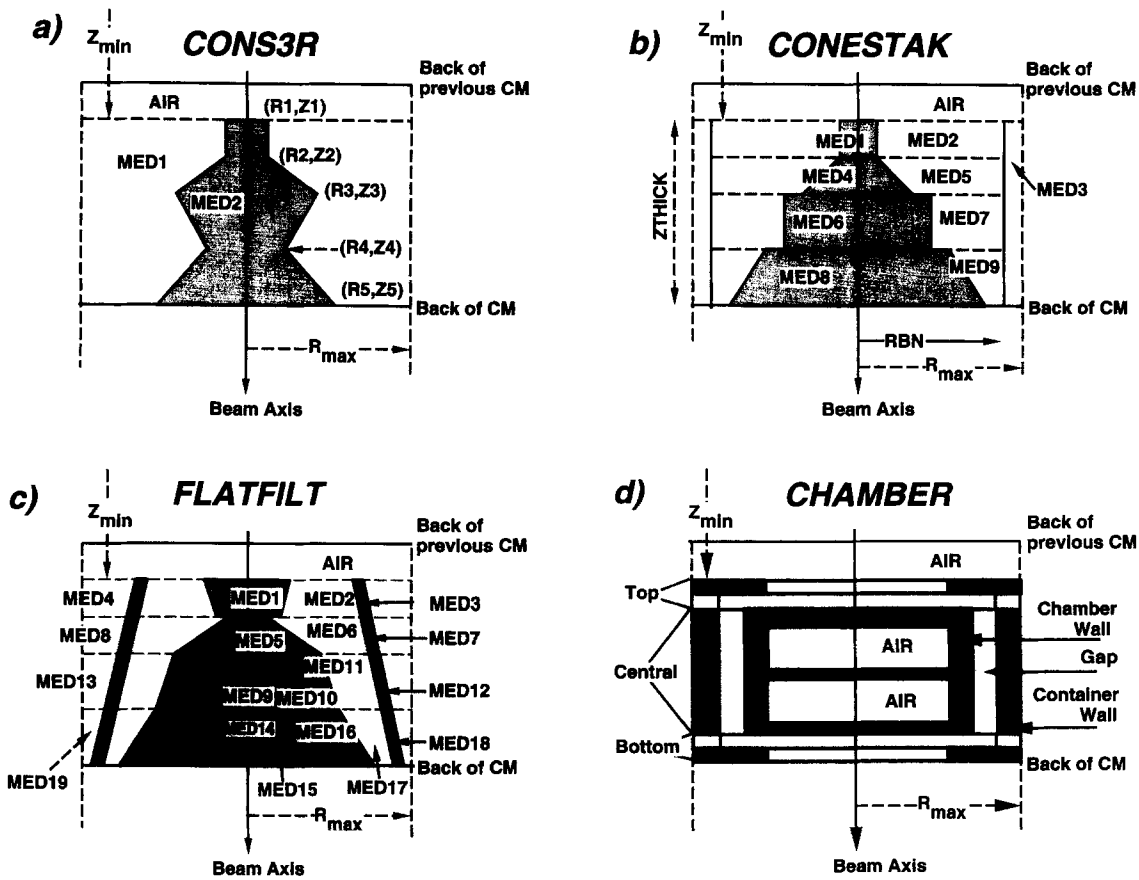

FIG. 2. (a) Component module CONS3R defined by 5 node points. Z_{MIN} is the distance from the reference plane ($Z=0$) to the front plane of the CM (excluding the air gap). (b) Component module CONESTAK shown with 4 layers of stacked coaxial truncated cones, one pair per layer, with a cylindrical wall surrounding. (c) Component module FLATFILT shown with an outer conical shell and 4 layers of stacked truncated coaxial cones, with an arbitrary number on each level. The number of coaxial cones on a given layer and the material in each cone need not be the same. Both the number of cones and the radii of cones in each layer can be specified independently. (d) CHAMBER component module for an ionization chamber or any symmetric cylindrical-planar geometries. The CM consists of up to three parts. Both the top and the bottom part may have several flat layers of different materials, two are shown; the materials may be different for the inner (circular) region and the outer (annular) region in the same layer. The central part is for the chamber, which may have several air cavities separated by layers of different materials and encircled by a chamber wall, a gap and a container wall. (e) The PYRAMIDS component module shown with two stacked, truncated pyramids. The positive and negative x and y values of the front and back edges of each face are defined by XFP, XFN, XBP, XBN and the same for Y, etc. The layers must have an air gap of at least 0.01 cm between them. (f) The APPSQ component module shown with two layers of truncated square columns or scrapers. Each scraper may have specified material and dimension. The minimum space between two layers of scraper is 0.01 cm of air. (g) The JAWS component module for a case of paired x and y jaws which can be asymmetric (hence the need to specify the XFP, XBN, etc., Front and Back x values in the Positive and Negative directions). The distances to the top and bottom surfaces of each pair of jaws are given relative to the reference plane ($Z=0$). The CM can model multiple sets of x or y jaws and/or wedges. (h) The MIRROR component module shown with two layers of material which are perpendicular to the z axis, but rotated about a line parallel to the y axis and truncated by two planes perpendicular to the z axis. The material in each layer may be specified by the user. (i) The XTUBE component module models an x-ray tube. The target may consist of several flat layers of different materials backed by a target holder (backing material). The target angle is defined by the target surface and the z-axis.

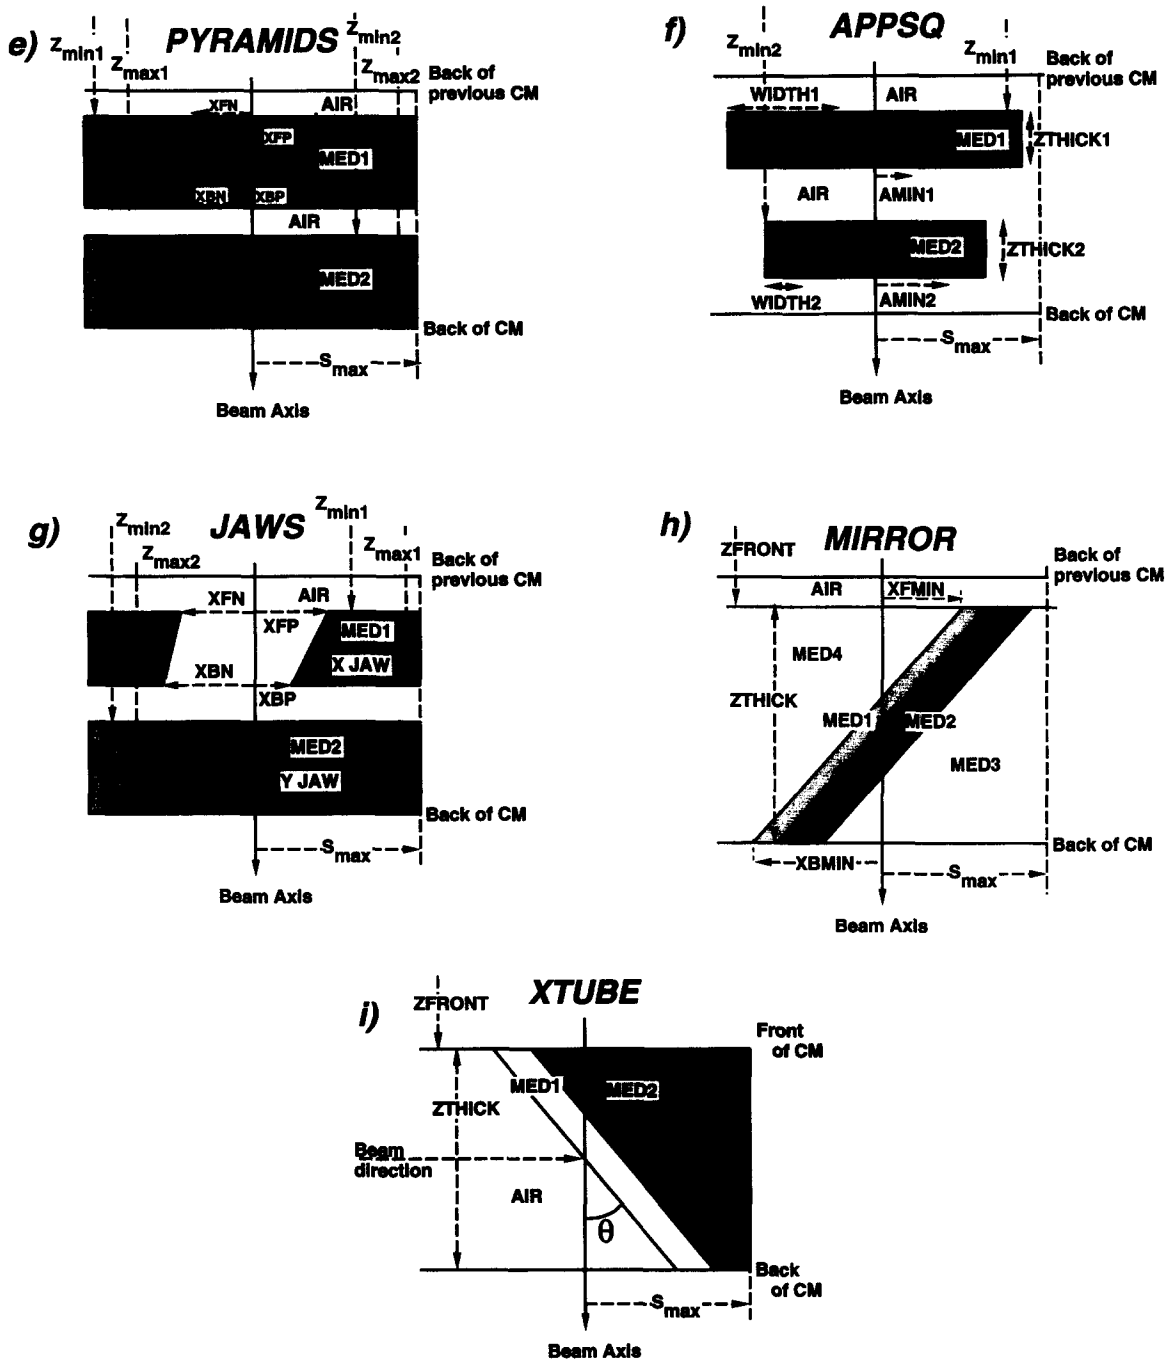

FIG. 2. (*Continued.*)

those entering the second CM. Special source configurations have been written for this CM with the incident rectangular or circular electron beam coming in normal to the z-axis. It should be emphasized that although this allows an x-ray tube to be simulated, the EGS4 code does not handle impact ionization and thus is expected to underestimate the fluorescent lines from the target.[44]

11. OTHER CMs

The above CMs are capable of simulating most standard accelerators, including wedges. However, other CMs will need to be developed for applications such as multi-leaf collimators or non-rectangular cutouts or beam blocks. For example, a CM modeling a wire mesh has been written and applied in one study.[45] It is also possible to include bolus and other similar modifiers in the patient simulation code, DOSXYZ.

D. Outputs from BEAM

BEAM has three major forms of output. The first, the listing file, is always present and optionally a phase-space data file can be produced and/or a graphics file can be produced which shows both the accelerator geometry in a stylized manner and every step in the particle's history.

1. Listing File

The listing file is the primary record of a simulation and is designed to make it possible to reproduce a given run by completely specifying all the input information plus an accurate specification of exactly what accelerator model was used. In particular, the entire system is handled under the Unix sccs library system and the exact version number of each piece of code used is echoed to the output listing and can be recovered using the sccs system if needed. In addition, all the other input control and geometric parameters are listed. There is a summary of the accelerator regions which includes absolute region numbers for every region in the accelerator plus the local region numbers of each region in its component module since these are more easily understood. In this same table the mappings from geometric regions into dose zones are defined (i.e., the geometric regions belonging to each dose scoring zone are specified) along with the mapping between geometric regions and bits set in LATCH (see section II.E).

The listing file contains several standard outputs. For each scoring plane defined in the accelerator model, a variety of fluence related parameters are scored in an arbitrary number of nested squares or circles. The number and widths of the scoring regions are selected on input, as are the locations of the scoring planes. For each particle type (electrons, positrons and photons) the total number of particles, the proper average particle fluence (cm^{-2}), the average energy and the average angle with respect to the z-axis are scored along with the estimated uncertainties on each. The particle fluence is scored by summing the inverse of $\cos(\theta)$ where θ is the angle with respect to the z-axis as the particle crosses the plane (capped at 85°, see Ref. 9). This method of scoring fluence has some limitations in the case of low-energy electrons especially and is only meant to give a reasonable estimate. The final output is the dose and energy deposited in an arbitrary number of dose zones (which can include multiple regions).

2. Phase-space data file

The phase-space data files contain the data about each individual particle crossing the scoring planes. These are usually the most important output. The exact contents of the file depend on the application, but a general purpose Fortran subroutine for reading and writing the phase-space files has been written and can be used by other programs. The standard files contain the variables LATCH, E,X,Y,U,V, SIGN(W), WT, IQ, and NPASS. LATCH is the history variable discussed below, E is the particle's energy, X and Y are the position coordinates in the scoring plane (the Z value is a constant and not stored), U and V are the direction cosines with respect to the X and Y axes, respectively, SIGN(W) is the sign of W, the direction cosine of the angle with respect to the Z axis [where the magnitude of W is recovered from the fact that $W=\sqrt{1-(U^2+V^2)}$], WT is the weight of the particle, IQ is the charge of the particle (-1, 0 or $+1$) and NPASS is a variable specifying whether this is the first or later time that the particle has crossed this scoring plane. A slight compression is achieved by storing SIGN(W) and NPASS in 1 bit each and IQ in two bits. Each record contains 28 bytes and thus a file with 1 million particles requires 28 Mbytes of disk space. Rather surprisingly, these files do not grow much as lower-energy electrons are tracked in an electron beam because there are so few of these electrons(see section IV).

3. Graphics output file

The component modules are set up to output a simplified representation of the geometry for input to the EGS_windows graphics package.[43] On request the user can output a complete history file which tracks every step in the simulation or output a data file which can be read by EGS_windows and displayed in three dimensional space. This feature can be very instructive and useful in debugging. The data files can become very large and only a few histories at a time should be followed.

Figure 3 shows representations of four of the accelerators studied in this work. Even though the representations are somewhat stylized, these figures immediately give considerable insight into the models being used. When sitting at a colour terminal with the capability of rotating and zooming the image, this tool becomes invaluable.

E. Tracking each particle's history

One of the major advantages of the Monte Carlo technique is that it allows detailed information about each particle's history to be known. To allow flexible access to this information, BEAM includes a general technique built upon the LATCH feature of EGS4.[39] LATCH is a variable which is passed on to a particle's descendants and which the user can interrogate or modify. In BEAM, each bit of LATCH is manipulated separately. Bits 1 to 23 are set whenever a particle interacts in a region (the region-to-bit correspondence is set by the user in the individual input files). For an electron, passing through a region constitutes interacting there whereas for a photon there are two options. The photon is either tagged the same way as an electron or only when an interaction actually occurs. The LATCH technique allows one to separate the effects of particles which have hit the applicator, or the jaws, etc.

Bit zero is set whenever a bremsstrahlung photon is created. Thus it is possible to interrogate any particle about whether there was a bremsstrahlung photon anywhere in its history.

The last piece of information which is stored is the region number in which the particle was created (bits 24–28). This can tell about the sources of secondary electrons in an electron beam and is useful for determining the region in which bremsstrahlung photons are created.

The LATCH feature is an option and requires setting up a bit-to-region correspondence in the input file. Although the feature requires some extra computing time, in practice the extra time is negligible ($<1\%$).

F. Calculation of charge from ion chambers

One of the features of BEAM, which no other accelerator model has included, is the direct calculation of the dose de-

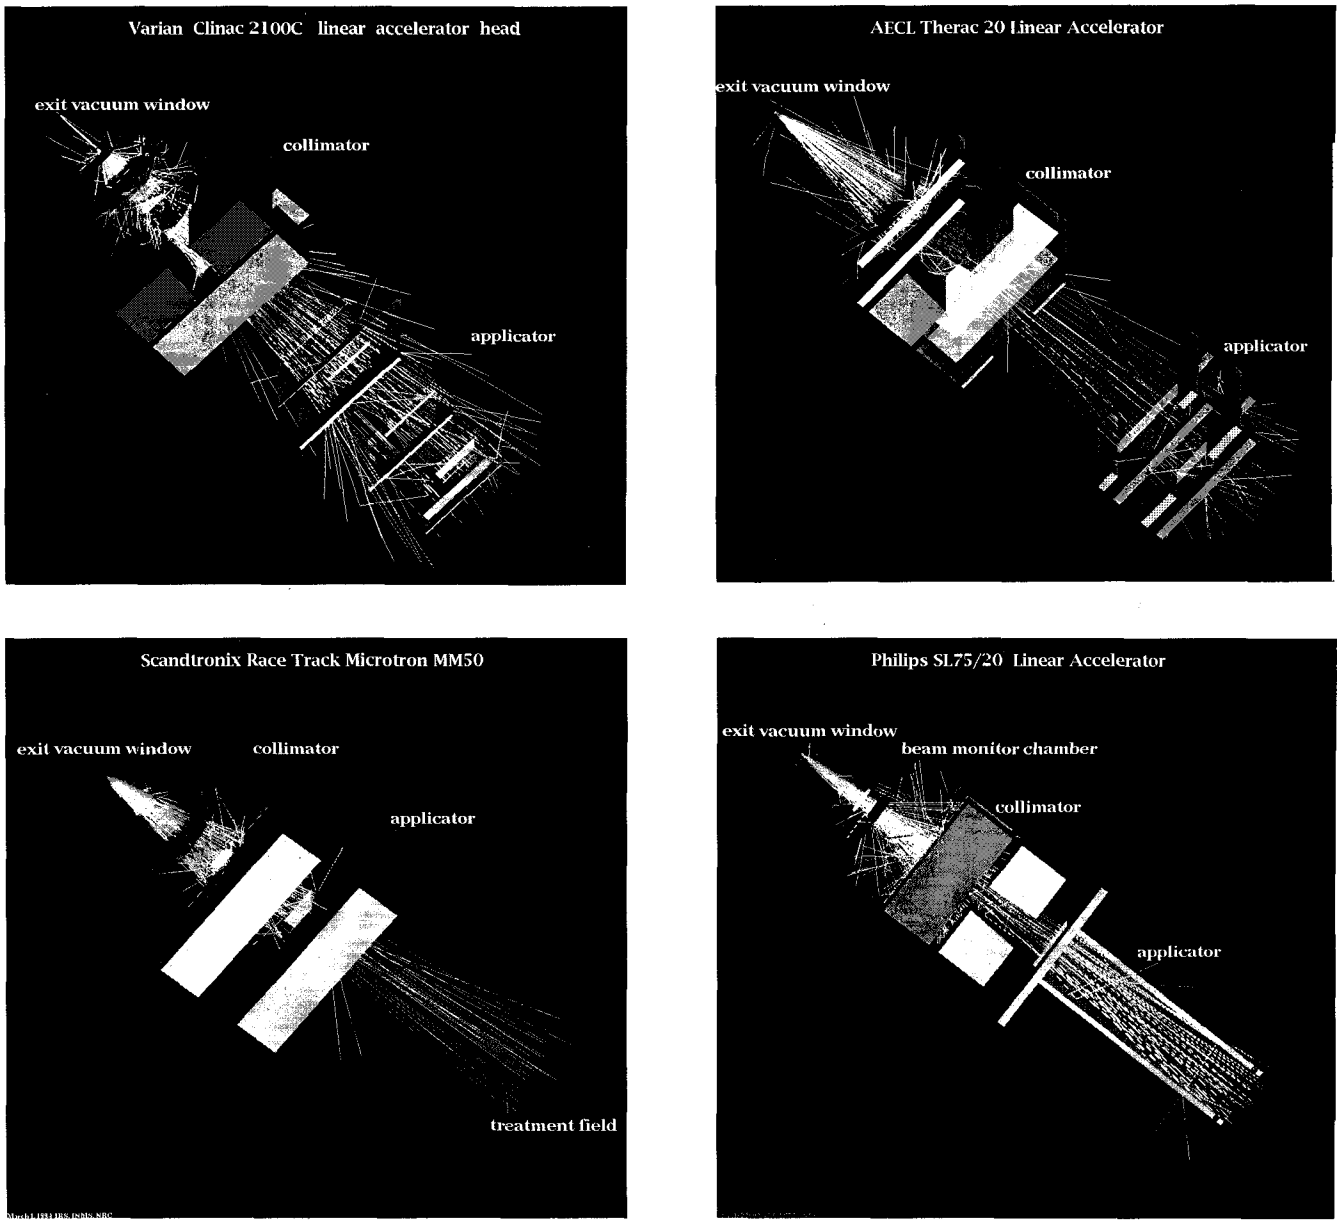

FIG. 3. Four commercial electron accelerators (Varian Clinac 2100C, AECL Therac 20, Philips SL75-20 and Scanditronix MM50) as shown by EGS_windows (Ref. 43) with about 100 histories. For full colour versions on the WWW, see http://www.irs.inms.nrc.ca/inms/irs/papers/BEAM95/beam.html.

posited in the air in the internal monitor chamber. This feature will allow studies to be made of the variation in output factors with field size and energy. The accurate calculation of absolute ion chamber response is perhaps the most difficult problem one can attempt with Monte Carlo simulations in the sense that at the 1% level, present calculation techniques do not work very well.[9,46] This is not a problem for radiotherapy beam calculations, but it can be critical for ion chamber response calculations. However, even if the absolute accuracy is somewhat uncertain, relative currents (absorbed dose to ion chamber air) should be calculated accurately. This feature has not been fully exploited yet, but is based on use of the CHAMBER component module which allows separation of the dose to the air into a variety of components caused by particles which have followed different types of histories. For example, studies have been done which show that the amount of backscatter into the current monitor decreases substantially as the electron beam energy is increased but is independent of the jaw settings for a Clinac 2100C.[47]

G. Variance reduction features

1. Range rejection

Range rejection can save significant quantities of computing time for electron transport calculations.[9] The basic method is to calculate the residual range of a charged particle and terminate its history if it cannot escape from the current region. The BEAM code includes a subroutine, MXRNGE, which is used to pre-compute the residual ranges to the threshold energy (AE MeV, total energy) in each medium as

a function of electron energy. From these tables, the residual range to the lowest energy for which an electron is transported (ECUT MeV, total energy) can be worked out since ECUT values may vary in different regions. These residual ranges represent the pathlength traveled by electrons slowing to ECUT if they undergo no discrete interactions and thus are calculated using the restricted stopping powers, i.e., they represent the maximum possible pathlength the particle could go in the simulation, and are, in general, larger than the residual CSDA ranges calculated with the larger unrestricted stopping powers. This is a conservative approach since very few particles actually go anywhere near these maximum possible residual ranges, both because of the discrete interactions, which also cause energy loss, and because of the multiple scattering of the electrons.

Range rejection involves an approximation, viz., any bremsstrahlung photons that would have been generated by the electrons as they slow down, are assumed not to escape from the region. The effects of this approximation can be minimized by setting a maximum energy for which a history can be terminated (called ESAVE in BEAM) and thereby allowing the higher energy electrons a possibility of creating bremsstrahlung which escapes from the region, even if the electron itself could not escape from the region. Another option is not to do range rejection in a region in which there is a lot of bremsstrahlung production (such as the bremsstrahlung target) but do it everywhere else.

The appropriate selection of ESAVE is situation dependent. For initial electrons of 20 MeV incident on W, simple radiative yield arguments show that about 11% of the bremsstrahlung energy fluence is created by electrons with energies below 5 MeV and only 2% by electrons below 2 MeV. In this case ESAVE of 2 MeV is acceptable, especially since attenuation of low-energy photons will reduce the amount of bremsstrahlung from low-energy electrons even further. For initial electrons of 10 MeV, the corresponding 5 and 2 MeV values are 32% and 7% and thus one would need to use an ESAVE value of 1 MeV to reduce the error to 2% or less.

The above range rejection approach is very general but in BEAM we can save further time by terminating any history which cannot reach a scoring region. This is hard to specify in general but as a partial answer to the problem, a threshold energy, called ECUTRR, is automatically calculated for each region such that an electron must be able to escape the region with energy greater than ECUTRR. If its energy is less than ECUTRR as it leaves the region, the particle cannot reach the scoring region of interest. For example, say we are interested in all electrons with energies down to 10 keV at the base of the accelerator. As an electron leaves a scattering foil 50 cm from the base, assuming only air in the way, an electron needs a kinetic energy of at least 200 keV to get to the base of the accelerator. Thus, there is no need to track electrons inside the scattering foil unless they will leave it with at least 200 keV. Thus when doing range rejection in the foil, BEAM terminates the history if the range of the electron as it slows down to 200 keV is less than the distance to the nearest boundary. Taking into account even thin monitor chambers and other materials in the beam saves time if electrons are being tracked to low energies. This is because, to first order, it takes as long to track an electron from 200 keV to 100 keV as from 20 MeV to 10 MeV. BEAM can automatically compute the minimum energy needed by electrons as they leave each CM if they are to reach the base of the accelerator and this can be the energy used in the range rejection routines. However, this option must not be used if, for example, one wants to calculate the charge created in the machine's monitor chamber and thus needs to track low-energy electrons inside the accelerator.

Table I shows a variety of timing results for simulations of an 18 MeV beam from the Clinac 2100C studied extensively in this paper. If we take case 1 as the base case, a variety of comments are in order. Firstly, these times are for an SGI Indigo with R4400 cpu but when run on a SUN Sparcstation 2, the case 1 calculation takes exactly 5 times as long. Thus, for this very different type of EGS4 simulation, we find that timing comparisons based on the standard EGS4 benchmark code[11] are applicable here. However, absolute timing comparisons only apply for calculations done on otherwise un-used machines and we find variations in cpu time per history of up to 50% depending on the computer's load. Case 2 *vs* case 1 shows that for ECUT=0.700 MeV (total energy), use of the automatically increased values of the low-energy cutoff as one moves further away from the scoring plane (as in case 1) saves only a few percent compared to leaving the low-energy cutoff fixed. Case 3 *vs* case 1 shows that allowing electrons with energies between 5 and 2 MeV to be transported, even when they cannot escape from their current regions, adds very little to the computing time but it also slightly increases the overall number of bremsstrahlung photons in the beam, reflecting that these electrons between 5 and 2 MeV are producing some photons but far fewer than the 9% one might expect based on the simple arguments given above. As expected, the number of electrons in the beam is unchanged. Case 4 shows that turning off range rejection altogether (by setting ESAVE=0.0) increases the computing time by about a factor of 2, with virtually no change in the number of photons or electrons in the beam compared to case 3. Cases 7, 8 and 9 indicate a similar trend if electrons are being tracked to even lower energies except that in this case (0.521 MeV instead of 0.700 MeV), range rejection improves the efficiency by a factor of 6 and tracking electrons from 5 to 2 MeV adds an extra 17% to the already longer calculation times. Comparison of case 7 and case 7a (where the values of the low-energy cutoff are not automatically increased) shows that this feature can lead to considerable improvement (30%) in computing speed when ECUT is low.

When simulating photon beams we find that large numbers of bremsstrahlung photons are generated in the target and subsequently absorbed deep in the primary collimator. Very few of these escape except those "near" the beam. To avoid senselessly tracking these photons in the collimator, we find it useful to introduce a very thin region on the upstream side of the collimator. This region prevents all particles from entering the top of the collimator except for those actually in the beam or nearby. This form of range rejection is very crude but can save up to 20% in computing time for photon beam calculations.

TABLE I. Timing variations for simulation of the 18 MeV electron beam from a Clinac 2100C. The beam is 10 $\times$ 10 cm^2 at an SSD of 100 cm but the particles are scored to a radius of 14 cm. Times are for an SGI Indigo with an R4400 cpu which runs at 55 VUPs (i.e., 55 times faster than a VAX11/780) for the EGS4 benchmark code (which can be compared to many other machines via the data in Ref. 11). Values of ECUT are automatically increased further from the scoring plane to improve range rejection efficiency, except where noted. AE is the threshold for the production of secondary electrons (total energy, kinetic energy=AE−0.511 MeV), a lower value implies more accurate energy-loss straggling. Electron histories are followed down to a total energy of ECUT(MeV). Electrons below energy ESAVE (total) are subject to range rejection, 5 MeV implies more than at 2 MeV, and 0 MeV implies no range rejection is used. The "total to file" is the number of particles written to file per incident particle ($\pm$ 0.5%). The last two columns show the breakdown between electrons and photons written to file($\pm$ 1%, except where noted). LATCH values are scored for all cases and data are written to a phase space file. Default PRESTA is used except where noted.

Case	AE (MeV)	ECUT (MeV)	ESAVE (MeV)	cpu s per history	total to file/inc e^-	e^- per 100 inc e^-	γ per 100 inc e^-
1	0.700	0.700	5.0	0.0124	0.417	9.15	32.4
2	0.700[a]	0.700	5.0	0.0127	0.411	9.09	31.9
3	0.700	0.700	2.0	0.0125	0.420	9.07	32.8
4	0.700	0.700	0.0	0.0249	0.421	9.03	33.0
5	0.700	0.900	5.0	0.0118	0.417	9.09	32.5
6	0.700	1.100	5.0	0.0083	0.411	9.13	31.9
7	0.521	0.521	5.0	0.0538	0.414	8.79	32.6
7a	0.521[a,b]	0.521	5.0	0.0735	0.413	8.51	32.7
8	0.521	0.521	2.0	0.0631	0.416	8.83	32.7
9	0.521	0.521	0.0	0.300	0.421	8.94	33.1
10	0.521	0.700	5.0	0.0178	0.416	9.09	32.4
11	0.700[c]	0.700	5.0	0.0172	0.414	9.01	32.3
12	0.700[d]	0.700	5.0	0.0120	0.413	9.23	32.0

[a]Automated variation of ECUT away from scoring plane not done.
[b]Ten times fewer histories run, thus statistical uncertainties are 3 times worse.
[c]ESTEPE=1%.
[d]ESTEPE=4%.

2. Bremsstrahlung splitting

For the simulation of clinical photon beams from a therapy machine, a variance reduction technique called "particle splitting"[35,9,48] can be used in BEAM to increase the number of bremsstrahlung photons created in the simulation geometry. At each bremsstrahlung interaction site, N photons are sampled individually and their weights are reduced by a factor of $1/N$. The statistical uncertainty in the photon energy fluence spectrum scored for a given number of incident electrons can be reduced dramatically since most of the time in these simulations is taken up tracking electron histories. Since one is normally interested in photon spectra when using bremsstrahlung splitting, BEAM can use Russian Roulette[49,9] on the electrons created by the multiplicity of bremsstrahlung photons. In essence, Russian Roulette means the number of electrons followed is restricted to the same number as would occur if bremsstrahlung splitting was not used (i.e., if N photons are created in each bremsstrahlung event, BEAM kills off N-1 out of N secondary electrons created by these photons and increases the weight of the survivors by a factor of N).

The increase in computing efficiency obtained using bremsstrahlung splitting is very dependent on the parameters of interest, as is the choice of an optimal splitting number. In realistic photon beam calculations we find it useful to use a splitting of 10 to 20 and can see increases in efficiency of up to a factor of 10.

3. Forced interactions

One variance reduction technique, forcing photon interactions, has been shown to be very helpful in some cases.[50,9] It is particularly helpful when studying electron contamination in accelerator beams since virtually no photons interact in the air and yet that is a major source of the electron contamination.[51,28] The way around this is to force photons to interact in the geometry. This can be very efficiently implemented in any EGS4 user code because of the general structure of the code.[52,9] However, one must take care to ensure that after the photon is forced to interact the original photon continues with a reduced weight so that the fluence of particles is not distorted anywhere in the geometry. In the implementation in BEAM, photons can be forced to interact in any subset of component modules.

H. Bremsstrahlung angular distributions

The default version of EGS4 deflects a bremsstrahlung photon by a fixed angle with respect to the electron generating the photon. This angle is the mean deflection angle and is about 1 to 2° for 15 MeV electrons. Since multiple scattering of electrons in high-Z materials is generally much larger than this, the default version of EGS4 assumes that the small effect of the bremsstrahlung angular distribution is washed out by the electron multiple scattering. It is, except on the first step. This first step can cause significant problems for the

angular distribution of bremsstrahlung photons from clinical accelerators. An extension to EGS4 which removes this problem has been utilized here.[35] This extension has up to a 40% effect near the central axis. The angular distribution of bremsstrahlung photons from high and low-Z materials has been shown to be in excellent agreement with measured angular distributions, but only if the extension is used.[36]

It may be worth noting that a problem of a similar magnitude has been identified with the default version of the ITS system and great care must be taken with that system as well to get the angular distribution of bremsstrahlung photons correct near the central axis.[37]

I. Incident beams

BEAM can handle input beams of electrons, photons or positrons. The selection of energies and geometries of the sources is generally decoupled except when a phase-space data file is being used as an input. These phase-space data files can be read in at any plane between CMs. In general this is to allow restarting calculations part way though the machine. This facility could also be used to input an absolutely arbitrary source by creating a phase-space file. The user also has the option of using a generalized source routine by writing a subroutine (with well defined requirements) which is used by BEAM. This option is used to allow characterized beams from earlier runs to be used as input sources.[14]

If phase-space data or the generalized source routine are not used, the user has the option of specifying an arbitrary energy spectrum for the incident beam.

BEAM handles a wide variety of input source geometries, most of which are assumed to start at the top of the accelerator. The most straightforward is a pencil beam of particles (usually electrons) incident along the z-axis at the origin. A slight extension is a parallel circular or rectangular beam incident on the first plane at an arbitrary angle. Another possibility is a point source on the accelerator axis at some arbitrary distance, uniformly irradiating a circle on the first plane.

There are three source geometries which simulate moving electron beams. The first models a beam which is used at NRCC to achieve flat photon beams without the need for a field flattening filter. The electron beam is incident on the first plane of the accelerator in a small circle centered on the axis and is swept to generate an incident cone of electrons. The geometry is defined by the radius of the circle and the half-angle of the cone. In the absence of any scatter in the accelerator, the electrons would sweep out a narrow ring about the beam axis. The second moving beam simulates the scanned beam from an AECL Therac 20 in which the beam is swept in a saw-tooth manner across the field of interest. The pattern is defined by the size of the region scanned at 100 cm from the accelerator exit window and the ratio of x and y scanning frequencies. Figure 4 shows the fundamental scanning pattern, once with no air present so that the underlying swept-beam characteristics are clear, and also with 100 cm of air present to demonstrate that just the air scattering in this swept 10 MeV beam is adequate to give a flat beam. The model simplifies reality because all the angles for sweeping the beam are selected prior to the beam entering the accelerator whereas in reality the beam is swept after it has passed through the vacuum exit window. The third moving beam is for modeling the scanned electron beam from MM50 accelerator and is arranged so it would sweep out a uniformly irradiated circle of arbitrary radius 100 cm from the start of the accelerator if there were no materials in the way.

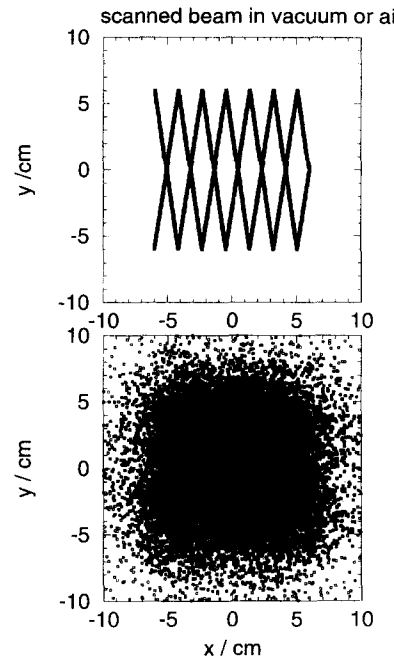

FIG. 4. Scanned beam pattern for 10 MeV electrons at 100 cm from the accelerator exit window. In the upper figure the electrons pass through vacuum after being scanned and in the lower figure they are scattered by 100 cm of air. The field size is artificially clipped near the top of the model to give a 12×12 cm^2 field.

For completeness the BEAM code has a source routine for simulating an ortho-voltage x-ray unit which is described in section II.C.10.

III. SELECTION OF RUN PARAMETERS

Selection of various run-time parameters for any Monte Carlo calculation can be very complex[9,39,53,38] and often depends on what aspect of the results are most important and how many resources are to be consumed by the calculation.

The BEAM code uses the PRESTA algorithm for electron transport.[38] For calculations down to ECUT=700 keV (total energy), comparison of cases 1, 11 and 12 in Table I suggests that utilizing the default PRESTA algorithm produces comparable results (to the 1% uncertainties) to calculations with steps sizes also restricted to ESTEPE values of 4% and 1%. For other calculations it may be important to use ESTEPE of 1% or even less[46] but at the expense of 35% or more computing time.

Another parameter of interest is AE, the lowest total energy for the production of secondary electrons. Lower values simulate more interactions and create more realistic energy-loss straggling distributions. If studying energy spectra per se, it is important to use a low value of AE (0.521 MeV), but if calculating depth-dose curves, much higher values are completely accurate.[52] Case 10 compared to case 1 in Table I

shows that using an AE value of 0.521 MeV instead of 0.700 MeV when using a value of ECUT = 0.700 MeV has virtually no effect on the results shown, but increases the computing time by 43%. It also significantly affects the high-energy electron fluence spectrum and thus the lower value should only be used if energy spectra are of interest.

Comparison of cases 1, 5 and 6 for AE=0.700 MeV and cases 7 and 10 for AE of 0.521 MeV show that increasing ECUT, the low-energy threshold for electron transport can reduce computing times substantially, especially from lower values (e.g., going from ECUT=0.521 to 0.700 MeV with AE = 0.521 MeV, reduces the computing time by a factor of about 3). However the rather surprising result is that the number of particles exiting from the accelerator actually decreases somewhat for the lower values of ECUT whereas one would expect the inclusion of the small number of low-energy electrons to increase the total number of electrons. Although the effect is just barely statistically significant, the unexpected result must come via a very weak dependence of the PRESTA algorithm on the value of ECUT.[38] We have chosen to do our calculations for ECUT = 0.700 MeV, but based on this table, it might often be wise to use a much higher value unless effects on the patient skin are an important part of the study.

IV. BENCHMARKING THE CODE FOR ELECTRON BEAMS

The problem of checking a code of this complexity and size is formidable. Just the Unix scripts to run the system are longer than the original EGS4 source code and with 20 000 lines of fortran (with all comments removed), the code to simulate a typical accelerator is bound to have many mistakes in it. The issue is to demonstrate that they are not critical. This has been done by extensively testing each individual CM in a wide variety of both typical and highly unusual situations against the results of widely used NRCC EGS4 user codes (DOSRZ, FLURZ and DOSXYZ) which have been extensively benchmarked.[9] The first two codes are restricted to cases with cylindrical symmetry and the latter to rectilinear voxel geometries. These obviously do not test all possible situations, but because each component module can be tested separately, one can cover a wide variety of cases. Much of the code has also been checked by simply having more than one person read the code and its documentation. The final test is to use BEAM to simulate the beams from a wide variety of accelerators and calculate the dose distributions in a homogeneous water phantom. We then compare calculated results to measurements. Comparisons to measured data in phantoms with heterogeneities are being presented elsewhere.[54]

Comparison to experiment is complicated by the fact that clinical accelerators have many physical parameters which users usually do not know. The most important of these may be the energy spectrum and direction of the electrons leaving the accelerator vacuum. This is discussed in section VII and for the moment, all beams in vacuum are assumed to be monoenergetic and on-axis. For many machines this is a good approximation. For example the Clinac 2100C spec sheets indicate it delivers a beam with a maximum energy width of ±3%. The incident beam energy for simulations of clinical electron beams are usually selected by matching the value of the calculated and measured values of R_{50}, the depth at which the dose falls to 50% of the dose maximum. Another option would be to match R_p, the projected range. We have calculated depth-dose curves for electron beams with a mean energy of 20 MeV and symmetric energy widths varying from 5 to 20%. For widths up to 10%, the energy width had little effect on the values of R_{50} or R_p but above that width, the value of R_p started to change since it has some sensitivity to the highest energy in the spectrum. In other words, the final depth-dose curve contains little usable information about the details of the *incident* electron spectrum unless the spectrum is very broad. We thus chose to match R_{50} since this is the easiest procedure in practice. It is interesting to note that we get very good agreement with the measured depth-dose curves by using this procedure (see below) but this may only reflect the insensitivity of the final depth-dose curves to the *incident* spectrum. In the case of our worst agreement (the SL75-20, 20 MeV beam, see below) we calculate a depth-dose curve which falls off too slowly past the dose maximum and a broad incident spectrum would only make it fall-off more slowly. Thus the discrepancy is not caused by the width of the incident spectrum although Deasy has shown that electron energy spectra from these machines do have a considerable width.[55]

We have obtained detailed information about a variety of accelerators. Often, when we failed to get agreement between the experiments and calculations we found that we were given incorrect information about a given accelerator because not all accelerators are standard, even for the same models. Thus the data calculated here apply to the specific machines discussed. Many accelerators are adjusted in the clinic to match previous accelerator beams (by selection of scattering foils or tuning of the waveguide) and manufacturers keep improving the hardware (e.g., developing thinner monitor chambers).

In the following sections we will describe the simulation of electron beams from 5 widely differing electron accelerators and comparison of the measured dose distributions in homogeneous water phantoms with those calculated using the NRCC EGS4 user-code DOSXYZ[9,11] which has been modified to use the full phase-space data file from BEAM as an input. DOSXYZ is a very simple Monte Carlo code which calculates the dose in rectangular 3-D voxels when irradiated by a beam described by the phase-space output from BEAM. One option which is available in DOSXYZ is to "re-use" the phase-space data as often as required to improve the statistical precision of calculated depth-dose curves. The phase-space file always contains enough particles (several million) to ensure accurate representation of the incident beam. While re-using the beam seems counterproductive, we find it leads to a distinct improvement in the precision of the calculated dose distributions. The random numbers ensure each particle follows a different path in the phantom. This should be compared to simpler situations in which incident particles have identical values for some or all parameters (e.g., when calculating the dose in a monoenergetic broad parallel beam). In DOSXYZ, if the incident beam is assumed to be symmetric

about the x and y axes (a common situation), then when re-using the phase space the particles are moved into the 4 symmetric positions [(x,y) to $(x,-y)$ with (u,v) to $(u,-v)$, etc.].

It should be emphasized that these comparisons require no approximations or models of the output obtained from the BEAM code. For disk space and other reasons, most previous comparisons of electron depth-dose curves broke the calculations into two steps.[8,16,17,56] The first step was to simulate the accelerator and to score the energy, spatial and sometimes angular distributions of the electrons from the accelerator. The second step was to calculate the central-axis depth-dose curve in a phantom using various approximations to generate the electron beam incident on the phantom, such as independence of energy, angular and/or spatial distributions. Udale[16,17] examined the effects of using different levels of complexity and Andreo and Fransson[56] demonstrated explicitly that ignoring correlations between energy and angular distributions leads to inaccuracies in the reconstructed incident beam. All these concerns are overcome in our case by using a brute force approach since the calculation retains all the phase-space information for every particle. In related work, the issue of accurate representation of these beams is studied in detail.[14]

These comparisons are restricted to electron beams since they are generally more difficult to simulate accurately than photon beams. Furthermore, various previous studies have demonstrated that EGS4 simulations predict accurate dose distributions for photon beams, even with relatively simple models.[31,28,29,32,25,30]

A. NRCC accelerator

The NRCC accelerator is a Vickers research accelerator with no head assembly. The energy of the electron beam exiting the vacuum window is known independently with an accuracy of $\pm 1\%$.[57] The accelerator has a 0.127 mm titanium exit window, a 0.092 mm thick tungsten scattering foil, a very thin mylar-walled monitor ion chamber, square steel jaws which define a fieldsize of 8×8 cm^2 on the phantom surface at an SSD of 96 cm. The water phantom has a 9 mm wall of PMMA and the calculations explicitly account for this. The measurements have been done with a small diode detector, a parallel plate ion chamber and a cylindrical farmer chamber. After correction of the ion chamber results for stopping-power ratios, all three measurement techniques gave the same results well within 1%.

The upper panel of Fig. 5 shows a comparison between the calculated and measured depth-dose curves when the incident accelerator energy in the simulation is that determined independently (viz., 20 MeV). There is a slight discrepancy in the values of R_{50} but otherwise the agreement is excellent. The figure also shows a similar comparison with the incident beam energy increased in the simulation by 1.5%. In this case even the value of R_{50} is in close agreement with experiment. The required incident beam energy is only 1.5 standard deviations from the measured value and thus the measured energy may be the problem. Alternatively, a 1.5% error in the water electron stopping power would not be too surprising and this would affect the calculated depth-dose curve, although it would be somewhat surprising if the stopping power is 1.5% too small at all energies.[58]

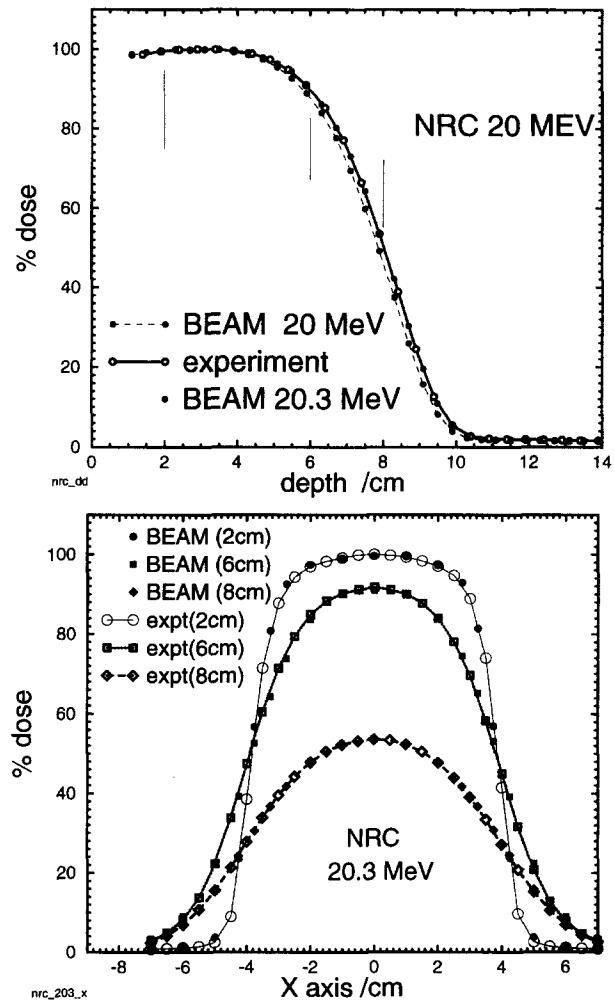

FIG. 5. Comparison of measured and Monte Carlo calculated depth-dose curves on the NRCC research linac with the incident energy in the simulations taken as the independently known energy of the electron beam leaving the accelerator vacuum (20.0 MeV) or increased by 1.5% to match the experimental data. The depth-dose curves are measured with diodes and ion-chambers. The ion-chamber readings are corrected with stopping-power ratios from the AAPM TG-21 protocol. The lower panel compares BEAM-calculated dose profiles and those measured with diodes at the depths shown by the vertical lines in the upper panel. Energy used in the simulation is 20.3 MeV. The first 9 mm of the phantom are PMMA.

Note the close agreement for the bremsstrahlung tail.

The lower panel of Fig. 5 also presents a comparison of the calculated and measured dose profiles in the NRCC beam measured with a diode detector. All data are normalized to the peak of the depth-dose curve. The agreement is remarkable. To our knowledge this represents the first *ab initio* calculation which properly reproduces the dose profiles in a realistic electron beam.

B. Varian Clinac 2100C accelerator

Figure 6 compares the calculated and measured central-axis depth-dose curves for the 6 and 18 MeV beams from a Clinac 2100C at UW. These machines have dual scattering foil systems and this particular accelerator has the scattering

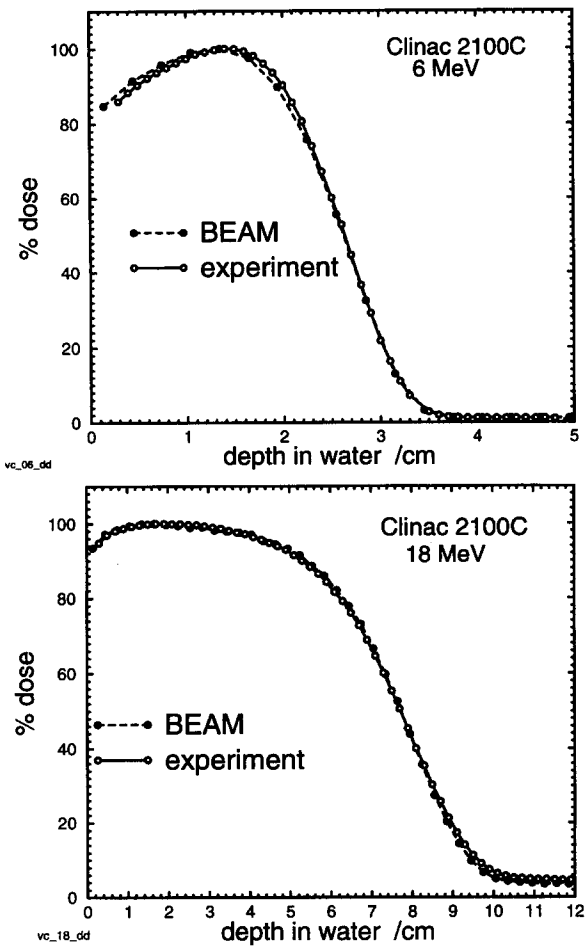

FIG. 6. Comparison of measured and BEAM/DOSXYZ calculated central-axis depth-dose curve for 6 (upper) and 18 (lower) MeV electron beams from the Varian Clinac 2100C accelerator at UW. The field size is 10×10 cm^2 at an SSD of 100 cm. The measurements are made with a diode detector. This accelerator has dual scattering foils, the monitor chamber is made with gold plated mica and the applicator has 4 sets of scraper bars.

foils of a Clinac 1800 in order to match the depth-dose curves of an earlier machine. It also has the relatively thick walled gold-plated mica ion chamber and a Type II applicator cone with 4 relatively thin scrapers. In these calculations the incident electron energies in the simulation are adjusted to give the measured values of R_{50}. The agreement is very good, well within 3% of D_{max} everywhere. At 6 MeV the calculated surface dose appears to be somewhat high and at 18 MeV the calculated dose fall-off may be somewhat too steep and the bremsstrahlung tail slightly underestimated. The 18 MeV results suggest there is something in the beam which is not included in our model or the scattering foils are slightly thicker than their nominal values since more material would increase the bremsstrahlung tail and increase the straggle which would make the fall-off slower. A comparison of measured and calculated dose profiles (in this case at $d_{max}=1.9$ cm in the 18 MeV beam) is shown in Fig. 7. The calculated and measured data are in good agreement, even outside the beam (see section VI for a discussion of what causes the dose outside the beam).

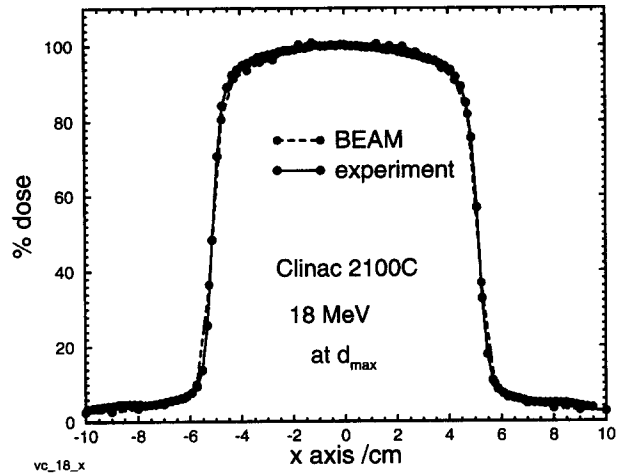

FIG. 7. Comparison of Monte Carlo calculated and measured dose profile at d_{max} ($=1.9$ cm) for the 18 MeV electron beam from the Varian Clinac 2100C accelerator at UW. The field size is 10×10 cm^2 at an SSD of 100 cm. The measurements are made with a diode detector.

C. AECL Therac 20 accelerator

Figure 8 compares the calculated and measured central-axis depth-dose curves for the 9 and 20 MeV beams from the AECL Therac 20 at the Ottawa Regional Cancer Centre. This accelerator scans the electron beams to achieve beam flatness and, as can be seen from Fig. 3, has much thicker applicator scrapers than the Clinac 2100C. Both these features lead to a much cleaner beam and very different central-axis depth-dose curves. The agreement between calculation and measurement is excellent. This indicates that BEAM is capable of accurately modeling clinical accelerators as different as the Clinac and Therac. Note that for the Therac 20 the bremsstrahlung tail is much less than for the Clinac and other scattering foil accelerators although the calculations also underestimate it slightly in the 20 MeV beam.

D. Philips SL75-20 accelerator

Figure 9 compares the calculated and measured central-axis depth-dose curves for the 10 and 20 MeV electron beams from a Philips SL75-20 accelerator which uses only a single set of scattering foils at any particular energy (measured values and machine specifications from Udale's thesis[15]). The applicator on this accelerator has solid flat walls compared to the series of scrapers in the applicators of the AECL Therac 20 and Clinac 2100C. This accelerator is again very different from the previous ones but the agreement between the calculations and measurements is very good. Udale also obtained good agreement between calculations and experiment for the 10 MeV beam.[17] For these beams the bremsstrahlung tail is very accurately predicted although, as mentioned above, the calculated fall-off in the 20 MeV beam is slower than experiment.

E. Racetrack MM50 accelerator

Figure 10 compares the calculated and measured central-axis depth-dose curves for the 25 and 50 MeV beams from the Scanditronix MM50 at the Memorial Sloan Kettering

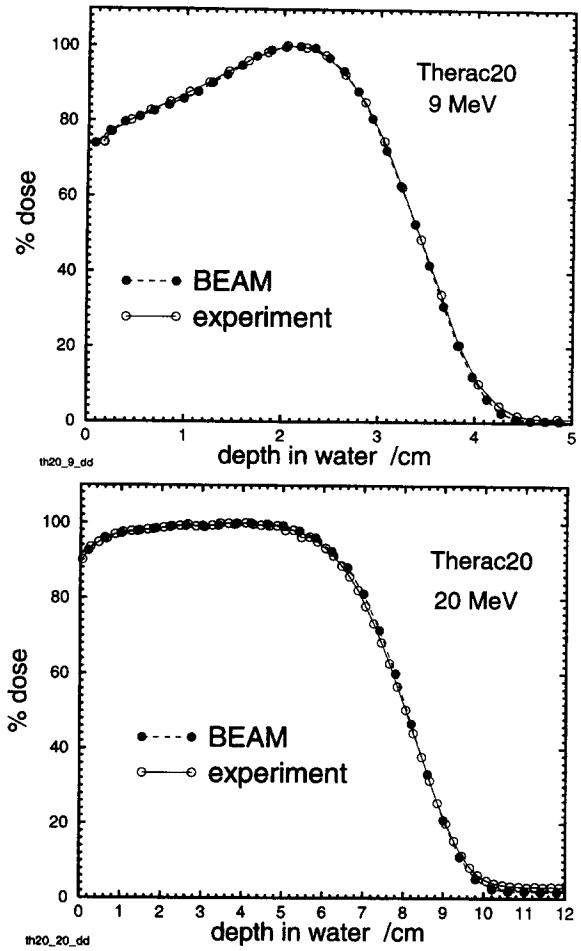

FIG. 8. Comparison of measured and BEAM/DOSXYZ calculated central-axis depth-dose curve for 9 (upper) and 20 (lower) MeV electron beams from an AECL Therac 20 accelerator. The field size is 10×10 cm^2 at an SSD of 100 cm. This accelerator has a scanned electron beam (with a 40 $\times$40 cm^2 scan pattern), the monitor chamber is made with aluminized Kapton and the applicator has thick scraper bars. Measured data are from Joanna Cygler of the Ottawa Regional Cancer Centre.

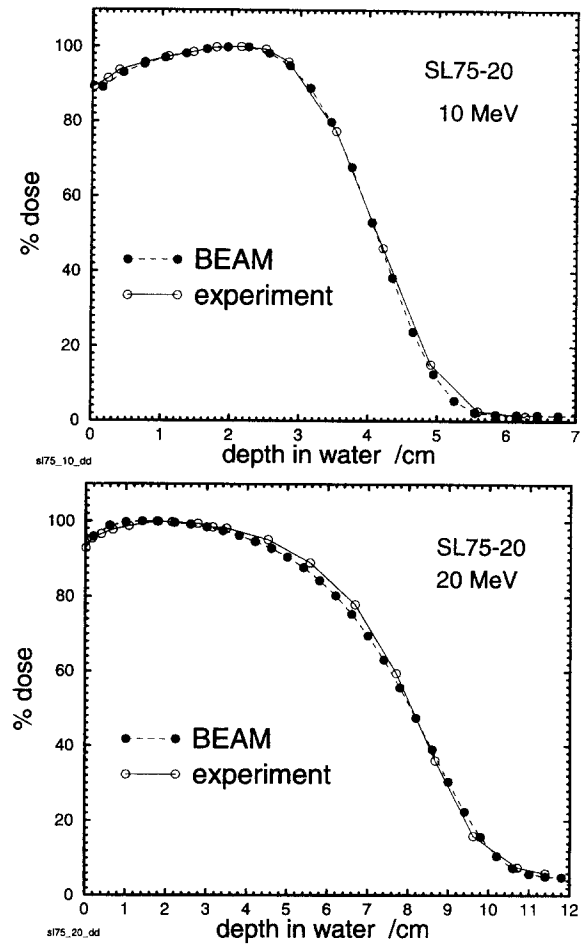

FIG. 9. Comparison of measured and BEAM/DOSXYZ calculated central-axis depth-dose curve for 10 (upper) and 20 (lower) MeV electron beams from a Philips SL75-20 accelerator. This accelerator has a single set of scattering foils and the applicator has solid walls. The field size is 10×10 cm^2 at an SSD of 100 cm. Measured data and information on the accelerator are from Udale (Ref. 15).

Cancer Center. The agreement is excellent. The MM50 uses both beam scanning over an 18 cm radius circle and a very thin scattering foil to give a uniform beam distribution. The beam sizes are larger than for the lower-energy beams to ensure broad-beam conditions. Unlike other accelerators, this machine does not have electron applicators at these beam energies. The beam is shaped by a primary collimator and a multileaf collimator which for the present case is set for a square. The beams from this accelerator are close to monoenergetic beams because of two factors: the scanned beam and the large air space between last part of the beam defining system and the phantom surface. This air gap means many of the electrons scattered from the beam defining system leave the field.

F. Summary of Experimental Comparisons

The above comparisons show that the BEAM code is capable of calculating the electron beams from a wide variety of accelerators very accurately. The depth-dose curves can be calculated with an accuracy which is certainly better than 2% of D_{max} on average. To achieve a tighter comparison would require more accurate measurements since most measurements ignore changes with depth of various factors and these may have a 1% or 2% effect. It is also worth noting that our calculations appear more accurate for the scanned beams. This may just reflect that they are much simpler beams or the fact that it is easier to get complete information about the accelerator because there is so much less information required.

V. ELECTRON BEAM SPECTRA AND OTHER CHARACTERISTICS

BEAM allows the energy distributions of electrons and photons to be studied in great detail. Variations in spectra based on location in the beam or last point of scatter can be important and are being investigated.[14]

Figure 11 presents the energy spectra for the standard electron beam energies for each of 4 accelerators studied in detail (a non-standard Clinac 2100C with a Type II applicator, SL75-20 with a solid walled applicator and the scanned beams from the Therac 20 and MM50 Racetrack Microtron).

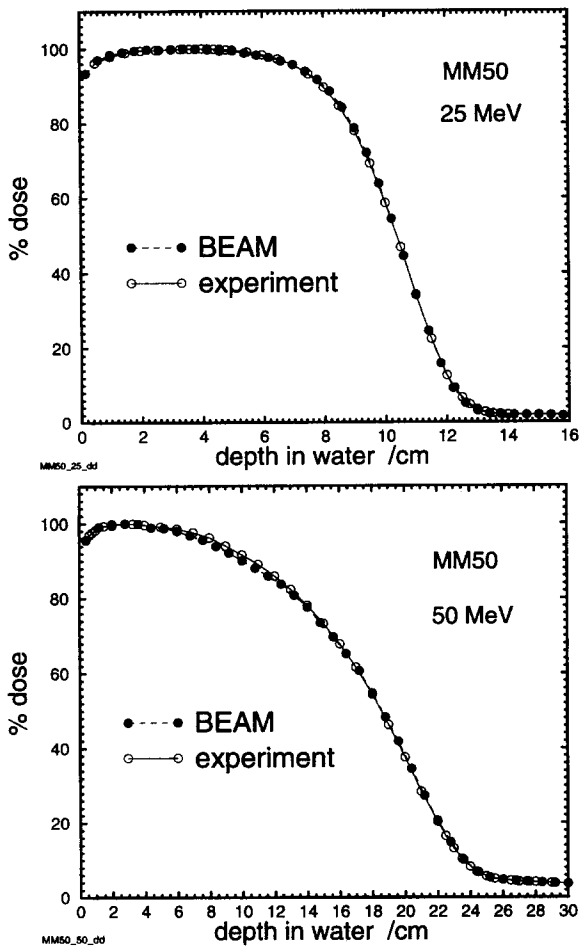

FIG. 10. Comparison of measured and BEAM/DOSXYZ calculated central-axis depth-dose curve for a 25 MeV (upper) and 50 MeV (lower) electron beam from the Racetrack MM50 accelerator. This accelerator has a scanned beam (which would uniformly irradiate a circle of radius 18 cm at the phantom surface) with a very thin scattering foil and no scrapers. The field sizes are 25×25 cm^{2}and 20×20 cm^2 at SSDs of 100 and 110 cm for the 25 and 50 MeV beams respectively. The measured dose curves come from measured ionization curves which are converted to dose using stopping power ratios calculated using the same BEAM simulation (Ref. 24). The measured ionization data are from Chen Chui, Memorial Sloan Kettering Cancer Center.

These are planar fluence spectra averaged over the open beam which is usually 10×10 cm^2 and at an SSD of 100 cm (details in previous section). The plots are shown on a logarithmic scale to display the data at low energies but it must be emphasized that on a linear plot the scanned beams appear to be mono-energetic whereas the machines with either single or dual scattering foils exhibit significant low-energy tails.

When calculating energy spectra, it is essential to use a lower threshold for the production of secondary electrons (AE = 0.521 MeV) since this ensures adequate energy-loss straggling. If a threshold of 700 keV is used, as done for the depth-dose calculations, there is a distinct artifact in the energy spectrum (see Ref. 53 for a complete discussion). In the Clinac 2100C spectra, the lower-energy peak about 2 MeV below the main peak is caused by electrons passing through one applicator scraper and remaining in the beam. This is discussed further below. For the SL75-20, there is a distinct double peak for the higher energy beams. The scattering foils consist of circular foils of varying radii and the first of these does not intercept the entire beam. In a separate run we have established that the upper, smaller peak comes from those electrons which "miss" this initial foil and thus lose less energy.[63]

As well as comparing our results to experiment, we have compared our results for the SL75-20 to those of Udale.[16,17] When we compare mean energies at the phantom surface for 10×10cm^2 applicators, we get reasonable agreement at 5 MeV but differ by nearly 300 keV at 10 MeV. Our calculations include electrons inside the field down to 189 keV kinetic energy whereas hers appear to include electrons on the entire phantom surface and include all electrons down to 50 keV. While these differences may explain the differences in calculated mean energies, they are not important when we compare the electron energy spectra near the peak energy. Here we see significant differences which may be related to her use of an AE value of 50 keV (kinetic energy) compared to our use of 10 keV which gives more accurate energy-loss straggling in the BEAM calculations. We have done calculations with AE = 189 and 10 keV and find significant differences due to the artifacts mentioned in the previous paragraph, but it is not clear that these would explain all of our differences with Udale's results.

Table II presents a variety of data related to the 8 accelerator beams for which calculated and measured depth-dose curves have been compared. The third column relates the number of electrons hitting the patient per cm^2 to the number of incident electrons. For fields of 100 cm^2 these numbers give the percentage of electrons getting to the patient. The percentages vary greatly, from around 2% for the low-energy beam from the Clinac to around 13% for the high-energy beams from the MM50. These variations are exactly what is expected since electrons in machines with scattering foils and lower-energy electrons scatter more.

The number of photons per cm^2 given in column 4 is seen to vary even more, being highest for high-energy beams from machines with scattering foils and being lowest for scanned beams from machines which do not define the beam with scrapers (viz., the MM50).

Column 5 shows the cpu time on a Sun Sparcstation 2 per (electron/cm^2) at the patient surface since this takes into account different field sizes and the fact that the precision of calculations for dose distributions depends directly on the incident fluence. These timing figures vary considerably and are very sensitive to the details of each calculation (see the discussion related to Table I) and the exact machine model being used. Several general observations can be made. A major determinant of computing time is the number of electrons reaching the patient per particle incident (see col 2). The 6 MeV beam from the Clinac 2100C requires the overall longest computing time. In contrast, the 50 MeV beam from the MM50 also requires a long time, mostly because the high-energy electrons which get stopped in the beam defining system take a very long time to track (their long ranges make the range rejection algorithms ineffective).

Columns 6 and 7 give the mean kinetic energies of the electrons (above the cutoff of 189 keV kinetic energy) and

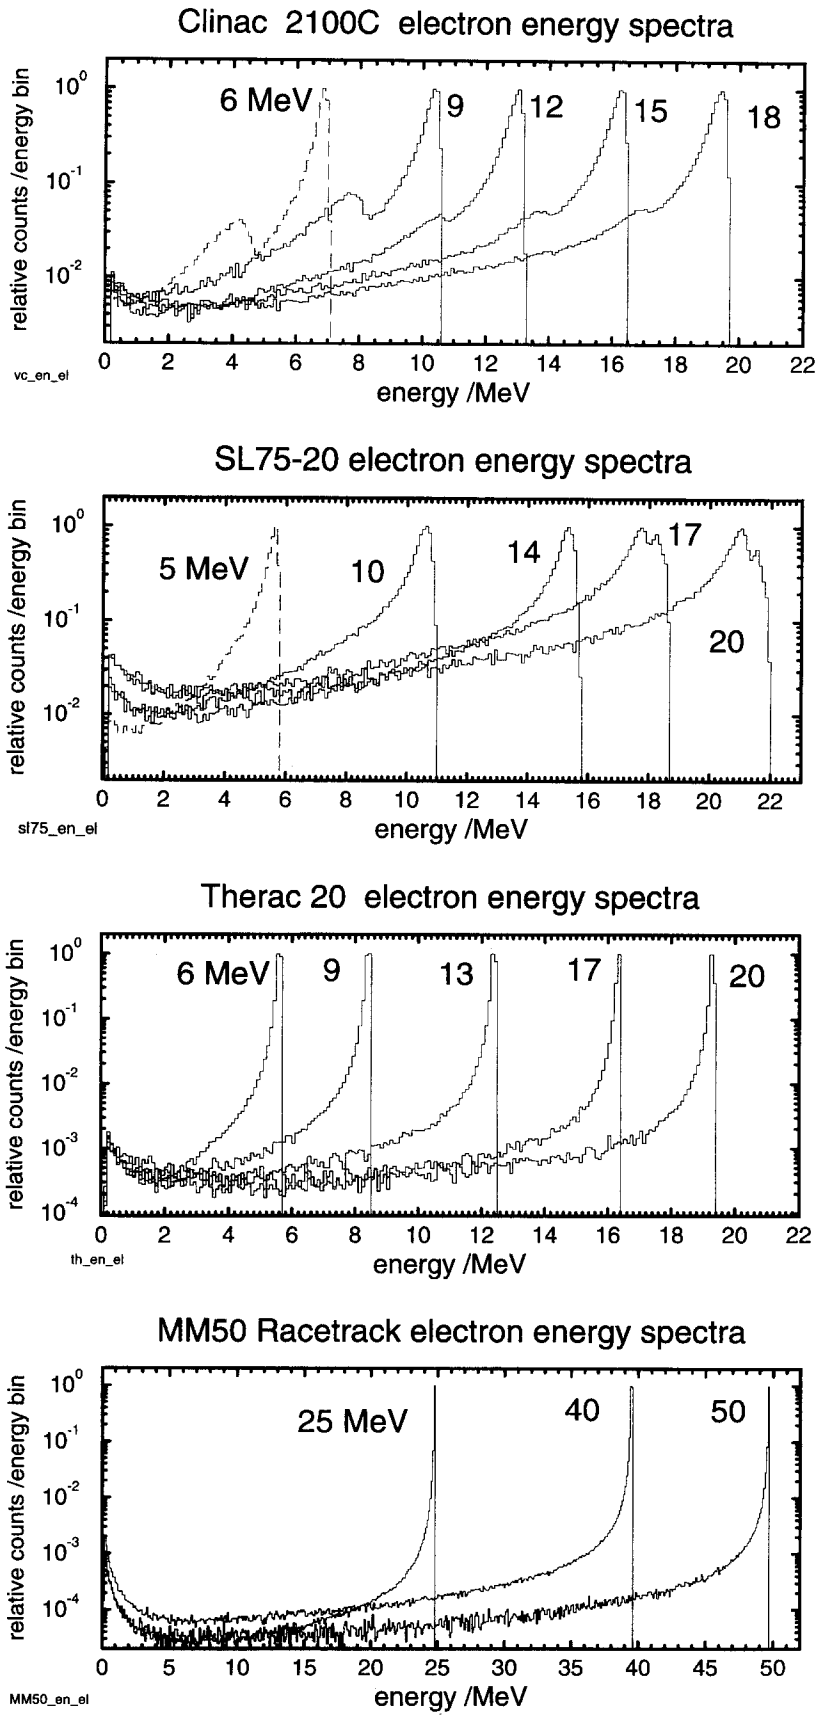

FIG. 11. Energy spectra from the 4 specific clinical accelerators studied in detail here. The spectra are for planar fluence averaged over the entire beam (usually 10×10 cm^2 at an SSD of 100 cm). Note that the y-axes are logarithmic. On a linear plot the scanned beams are virtually monoenergetic and the scattered beams contain a considerable low-energy tail. These spectra are for incident monoenergetic beams and any width in the incident beam from the accelerator vacuum would add to the breadth of the final spectra. All spectra are calculated with ECUT=0.700 MeV and AE=0.521 MeV (total energy) to ensure adequate energy-loss straggling.

TABLE II. Various data associated with electron beam simulations for 4 accelerators. Most particle fluences and average kinetic energies are number averages inside a 10×10 cm^2 field at an SSD of 100 cm. For the Therac 20 the scan pattern was set for a 40×40 cm^2 field but the scored field was 10×10 cm^2 while for the MM50 the scan pattern was set for a circle of 18 cm radius and a field size of 25×25 cm^2(25 MeV) and 20×20 cm^2(50 MeV). The EGS4 parameters AE, ECUT, AP, and ESTEPE are 0.521 (total), 0.700 (total), and 0.010 MeV and 1%, respectively. Direct particles are those which do not scatter from anything outside the beam. Times are for a SUN Sparcstation 2 which runs at 11 VUPs for the EGS4 benchmark code (which can be compared to many other machines via the data in Ref. 11).

Machine	Nominal energy (MeV)	e^-/cm^2 /10^4 inc e^-	γ/cm^2 /10^4 inc e^-	cpu s/(e^-/cm^2) on patient	$\bar{E}$ (MeV)		% direct	% direct
					e^-	γ	e^-	γ
Clinac	6	2.12	2.08	250	6.11	1.36	78.9	69.6
2100C	18	7.56	12.47	66	17.35	2.81	80.8	63.7
Therac 20	9	5.64	1.24	50	8.25	1.30	95.9	16.3
	20	6.35	4.63	50	18.74	2.17	94.8	6.0
SL75-20	10	4.14	4.07	96	9.34	1.96	66.8	75.6
	20	6.94	13.60	68	17.76	3.35	69.5	78.5
MM50	25	13.6	0.80	21	24.53	3.50	99.6	50.9
	50	12.7	2.63	105	48.90	5.59	99.1	42.0

contaminant photons in the beams. These and similar data form the basis of an extensive study of electron energies in clinical beams.[59] The final two columns show the percentages of the electrons and photons which are "direct" in the sense that they have not scattered off the scrapers or jaws but have come directly through the beam region. Both the scanned beams have a large proportion of direct electrons. This is an indication of the effectiveness of their collimation systems at removing electrons completely, rather than scatter them since, e.g., the Therac 20, has roughly the same overall transmission as the scattering foil accelerators. The converse of this is that nearly all of the photon contamination in the Therac accelerator comes from the scrapers which stop most of the unused beam, and hence these photons are not "direct." For the scattering-foil accelerators, most of the photon contamination come from the scattering foils and are thus direct.

A report is available which presents complete spectral and other data for electrons and photons for up to 5 beam energies from each of the electron accelerators discussed in this paper.[63] It must be emphasized that data for one particular accelerator cannot necessarily be applied to another accelerator of the same model.

VI. ELECTRON BEAMS FROM A CLINAC 2100C

In sections IV and V we presented results for a variety of accelerators to demonstrate the ability of the code to model them accurately and as a source of useful information. In this section we present a more detailed analysis of the beams from a Clinac 2100C, both to demonstrate the capabilities of BEAM and to investigate the physics of electron beams.

Figure 12 shows the 18 MeV depth-dose curve from the Clinac 2100C and breaks the curve into various components. Electrons which do not scatter from any part of the accelerator outside the beam are referred to as direct. It can be seen that these produce somewhat more than 70% of the dose and define the fall-off region of the curve. Electrons which scatter from the applicator scrapers or main jaws both contribute roughly 10% to the dose near the surface, but are clearly of lower energy. The electrons from the applicator are the least penetrating since some of them go through the scrapers. The photon tail can be seen to be primarily generated by contaminant photons incident on the phantom as opposed to bremsstrahlung photons generated in the phantom. The figure also shows the dose from incident secondary electrons which are defined as those knocked on by any of the others. These low-energy electrons contribute very little to the dose (2% or less) and play no role past a few mm.

Figure 13 shows the dose components in beam profiles at depths of 1 and 10 cm (which is just past R_p). Near the surface it can be seen that the major contribution to the dose outside the beam comes from electrons passing through the applicator along with photon contamination to a lesser extent. At depth it is seen that on-axis about 60% of the dose comes from the contaminant photons generated in the accelerator (mostly the scattering foil). The contaminant photons also extend into the wings whereas the dose from the bremsstrahlung generated in the phantom by the direct electrons is relatively flat but mostly within the beam.

Figure 14 shows the components of the energy spectra in the 9 MeV beam from the same Clinac 2100C fitted with either the older Type II applicator or the new Type III Varian applicator design as described by Kassaee *et al.*[60] who have recently published a study including Monte Carlo calculations using the cylindrically symmetric code CYLTRAN.[61] Consistent with Kassaee *et al.*'s results, our spectra show clearly that the new applicator design, which gives a significantly cleaner beam, avoids generating electrons from the applicator. The low-energy peak in the old design is just over 2 MeV below the main peak and this corresponds to electrons which pass through a 6 mm aluminum scraper which defines the beam about 31 cm from the patient (this was shown in a separate run which separated the spectra from each scraper component). By calculation it was found that the peak could be made to disappear using an additional piece of lead on the top of this part of the applicator.[63] The

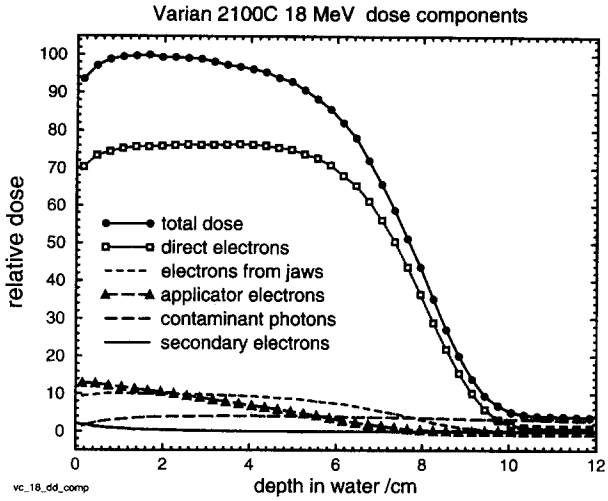

FIG. 12. Calculated components of depth-dose curve for Clinac 2100C with beam energy 18 MeV, field size=10×10 cm^2 and SSD=100 cm. The y-axis is offset from zero for clarity. Direct electrons are those which have hit no walls or the applicator. Electrons from the jaws or applicators are those which have interacted in the respective components at some point (and hence include some double counting). The contaminant photons and secondary electrons are those incident on the surface of the phantom. The central-axis scoring region is 2×3 cm^2. Thresholds of AE=ECUT=0.700 MeV are used.

calculated depth-dose curve with the lead in place was virtually identical to that of the new design applicator. We have been advised that Varian sells an upgrade kit which performs the same function.

VII. PROBLEMS

Despite the abilities of the Monte Carlo technique to do very accurate simulations, there are a variety of practical problems involved.

The most important problem is getting adequate information about the specific accelerator to be modeled. Manufacturers are reticent to provide the necessary level of detail because of the obvious commercial value. Furthermore, accelerator designs are constantly improving, even within a given model and thus one set of information will not cover, e.g., all Clinac 2100Cs. Finally, the machines are often adjusted for individual purchasers, often to match a previous machine's characteristics.

Another area of potential concern is adequate knowledge of the phase space of the electron beam as it leaves the accelerator vacuum. Our method of varying the energy of a mono-energetic incident beam to match R_{50} allows for accurate calculation of dose distributions, but there may be other, more subtle effects, for which this approximation fails. Thus our calculated energy spectra can be thought of as a smearing function for the accelerator and not necessarily an accurate representation of the final spectra. Nonetheless, as discussed in section IV, this uncertainty is not thought to affect dose-distributions significantly.

Another short-coming of the present technique is the large data files involved with phase-space files. This problem

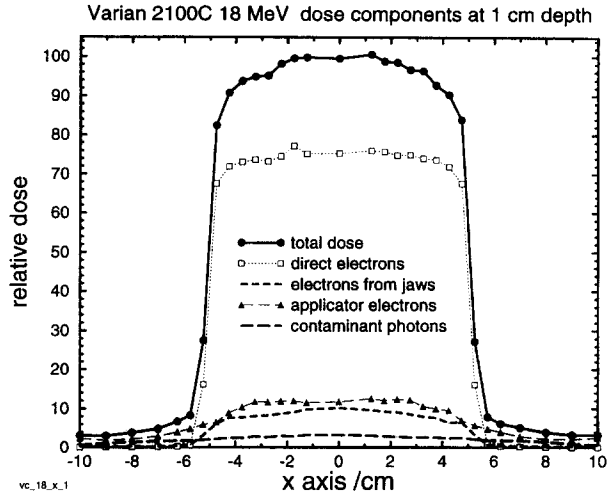

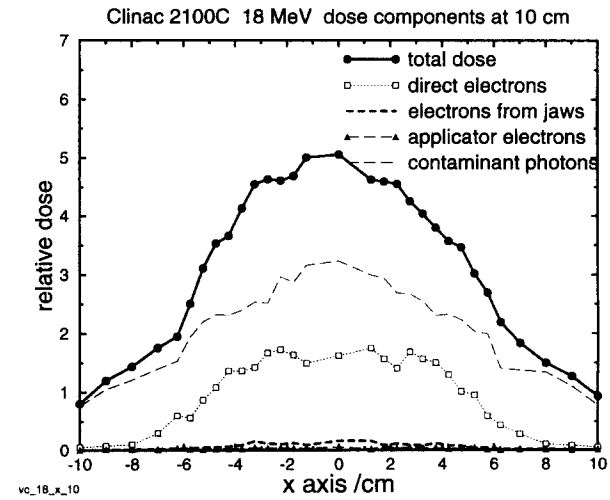

FIG. 13. Components of dose for dose profiles for an 18 MeV electron beam at depths of 1 and 10 cm for a Varian Clinac 2100C, field size= 10×10 cm^2 and SSD=100 cm. Components have the same meaning as in Fig. 12. The profile at 10 cm depth is past the range of the incident electrons but the bremsstrahlung associated with a given class of electrons is included in that component. The fluctuations are statistical.

should be solved by detailed models of the beam which will allow much smaller data files to represent the beams accurately.[14]

VIII. SUMMARY AND CONCLUSIONS

The code BEAM has been described in some detail and extensive comparisons to measured dose distributions in water phantoms have shown that it can reproduce these well within 3% of D_{max} at all depths. Extensive data concerning the characteristics of electron beams from 4 commercial accelerators of very different designs have been presented. Similar work is underway concerning photon beams from commercial accelerators.[62] The ability of the code to help design and modify accelerator equipment is obvious.

The major strength of the BEAM code is its flexibility and extensibility. It is possible to put together, with little effort, models which will cover very different accelerators. In fact,

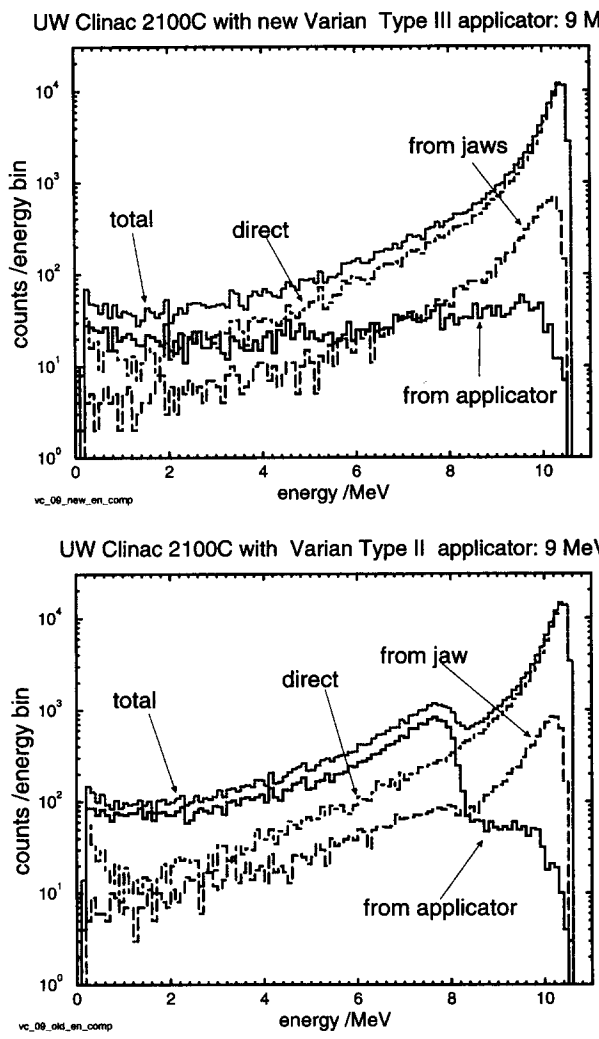

FIG. 14. Calculated energy spectra from UW's Clinac 2100C with old (Type II) and new (Type III) designs of applicators. The spectra for electrons from various components as defined in previous figures are shown. The peak at 7.8 MeV in the lower panel for the older, Type II design applicator can be seen to come from electrons which hit the applicator (and mostly come through the third scraper up, as shown in other runs) and the other components change very little.

the system represents a powerful tool for building a variety of complex geometries for simulating many situations which are not accelerators.

The code will be made generally available to the medical physics community as a research and educational tool. The current version requires a Unix operating system, lots of disk space (>1Gbyte), especially if generating phase-space files, and the NRC Unix-based EGS4 system must be in place with the "Big" Mortran pre-processor and a Fortran compiler.

ACKNOWLEDGMENTS

We would like to thank the following individuals for supplying measured data related to various accelerator beams: Carl Ross of NRC for data related to the NRCC accelerator; Joanna Cygler of the Ottawa Cancer Clinic for data related to the Therac 20; Chen Chui of Memorial Sloan Kettering Cancer Center for data related to the Racetrack Microtron; Carol Wells of the University of Wisconsin for data related to the Clinac 2100C. We wish to thank Alex Bielajew of NRC for extensive efforts to keep the EGS4 system running effectively as computing systems constantly change, and for his role as guru concerning EGS4, Mortran and EGS_windows. The assistance of Dr. Michel Proulx of NRC in keeping the IRS computing system functional is greatly appreciated. We wish to acknowledge the help of Blake Walters, Daryoush Sheikh-Bagheri and Geoff Zhang, all at NRC, who have contributed to the development of the BEAM code by using it extensively. We also acknowledge the efforts of Dr. J. Sun, who got the PAW system working at NRC, taught us how to use it and did a variety of quality assurance checks on the input/output routines for the CMs. We thank Stan Zurawski of the NRCC Computing Group for his work to re-write EGS_windows for the Silicon Graphics machines. We thank Cal Huntzinger of Varian for useful comments on the manuscript. We thank Mary Udale for many stimulating discussions about modelling accelerators and for installing her Monte Carlo simulation code at NRC. This work is partially supported by NCI Grant No. R01CA52692. One of us (G.X.D.) is partially funded by a post-graduate fellowship from the Medical Research Council of Canada.

[a]E-mail: dave@irs.phy.nrc.ca, Fax: 613-952-9865, and Tel: 613-993-2715.
[b]Currently at Toronto Sunnybrook Regional Cancer Centre, Toronto, Ontario.
[c]Also, Physics Department, Carleton University, Ottawa, Ontario.

[1]D. W. O. Rogers, A. F. Bielajew, T. R. Mackie, and S. S. Kubsad, "The OMEGA Project: Treatment planning for electron-beam radiotherapy using Monte Carlo techniques," Phys. Med. Biol. **35**, (abstract) 285 (1990).
[2]T. R. Mackie, S. S. Kubsad, D. W. O. Rogers, and A. F. Bielajew, "The OMEGA project: Electron dose planning using Monte Carlo simulation," Med. Phys. **17**, (abstract) 730 (1990).
[3]E. Mah, J. Antolak, J. W. Scrimger, and J. J. Battista, "Experimental evaluation of a 2D and a 3D electron pencil beam algorithm," Phys. Med. Biol. **34**, 1179–1194 (1989).
[4]J. Cygler, J. J. Battista, J. W. Scrimger, E. Mah, and J. Antolak, "Electron dose distributions in experimental phantoms: a comparison with 2D pencil beam calculations," Phys. Med. Biol. **32**, 1073–1086 (1987).
[5]A. S. Shiu and K. R. Hogstrom, "Pencil-beam redefinition algorithm for electron dose distributions," Med. Phys. **18**, 7–18 (1991).
[6]M. Morawska-Kaczynska and H. Huizenga, "Numerical calculation of energy deposition by broad high-energy electron beams: II. Multi-layered geometry," Phys. Med. Biol. **37**, 2103–2116 (1992).
[7]D. Jette and S. Walker, "Electron dose calculation using multiple-scattering theory: Evaluation of a new model for inhomogeneities," Med. Phys. **19**, 1241–1254 (1992).
[8]K. R. Shortt, C. K. Ross, A. F. Bielajew, and D. W. O. Rogers, "Electron Beam Dose Distributions Near Standard Inhomogeneities," Phys. Med. Bio. **31**, 235–249 (1986).
[9]D. W. O. Rogers and A. F. Bielajew, "Monte Carlo techniques of electron and photon transport for radiation dosimetry," in *The Dosimetry of Ionizing Radiation*, Vol. III, edited by K. R. Kase, B. E. Bjarngard, and F. H. Attix (Academic, New York, 1990), pp. 427–539.
[10]T. R. Mackie, "Applications of the Monte Carlo Method in Radiotherapy," in Vol. III of *Dosimetry of Ionizing Radiation*, edited by K. Kase, B. Bjarngard, and F. H. Attix (Academic, New York, 1990), pp. 541–620.
[11]A. F. Bielajew and D. W. O. Rogers, "A standard timing benchmark for EGS4 Monte Carlo Calculations," Med. Phys. **19**, 303–304 (1992).
[12]M. A. Holmes, T. R. Mackie, W. Sohn, P. J. Reckwerdt, T. J. Kinsella, A. F. Bielajew, and D. W. O. Rogers, "The application of correlated sampling to the computation of electron beam dose distributions in heterogeneous phantoms using the Monte Carlo method," Phys. Med. Biol. **38**, 675–688 (1993).
[13]H. Neuenschwander, T. R. Mackie, P. J. Reckwerdt, and J. Yang, "Improved Macro Monte Carlo method for electron beam dose calculations," Phys. Med. Biol. (in press) (1995).

[14]C. M. Ma, D. W. O. Rogers, B. A. Faddegon, G. X. Ding, J. S. Wei, A. F. Bielajew, and T. R. Mackie, "Simplified models of electron beams from a Clinac 2100C accelerator," Med. Phys. **20**, 1295 (abstract) (1993).
[15]M. Udale-Smith, "A Monte Carlo Investigation of High Energy Electron Beams Used in Radiotherapy," PhD. thesis, Leeds University (1990).
[16]M. Udale, "Monte Carlo calculations of electron beam parameters for three Philips linear accelerators," Phys. Med. Biol. **37**, 85–105 (1992).
[17]M. Udale, "A Monte Carlo investigation of surface doses for broad electron beams," Phys. Med. Biol. **33**, 939–954 (1988).
[18]C. Manfredotti, U. Nastasi, R. Ragona, and S. Anglesio, "Comparison of three dimensional Monte Carlo simulation and the pencil beam algorithm for an electron beam from a linear accelerator," Nucl. Instr. Meth. A **255**, 355–359 (1987).
[19]C. Manfredotti, U. Nastasi, R. Marchisio, C. Ongaro, G. Gervino, R. Ragona, S. Anglesio, and G. Sannazzari, "Monte Carlo simulation of dose distribution in electron beam radiotherapy treatment planning," Nucl. Instr. Meth. A **291**, 646–654 (1990).
[20]A. A. Al-Beteri and D. E. Raeside, "Optimal electron-beam treatment planning for retinoblastoma using a new three-dimensional Monte Carlo-based treatment planning system," Med. Phys. **19**, 125–135 (1992).
[21]T. R. Mackie, J. W. Scrimger, and J. J. Battista, "A convolution method of calculating dose for 15 MV x-rays," Med. Phys. **12**, 188–196 (1985).
[22]R. Mohan, C. Chui, and L. Lidofsky, "Differential pencil beam dose computation model for photons," Med. Phys. **13**, 64–73 (1986).
[23]P. Andreo, A. Brahme, A. E. Nahum, and O. Mattsson, "Influence of energy and angular spread on stopping-power ratios for electron beams," Phys. Med. Biol. **34**, 751–768 (1989).
[24]G. X. Ding, D. W. O. Rogers, and T. R. Mackie, "Calculation of stopping-power ratios using realistic clinical electron beams," Med. Phys. **22**, 489–501 (1995).
[25]P. L. Petti, M. S. Goodman, T. A. Gabriel, and R. Mohan, "Investigation of buildup dose from electron contamination of clinical photon beams," Med. Phys. **10**, 18–24 (1983).
[26]P. L. Petti, M. S. Goodman, J. M. Sisterson, P. J. Biggs, T. A. Gabriel and R. Mohan, "Sources of electron contamination for the Clinac–35 25–MV photon beam," Med. Phys. **10**, 856–861 (1983).
[27]ICRU, "Specification of high-activity gamma-ray sources," International Commisssion on Radiation Units and Measurements Report 18 (ICRU Bethesda U.S.A.) (1971).
[28]D. W. O. Rogers, G. M Ewart, A. F. Bielajew, and G. van Dyk, "Calculation of Electron Contamination in a ^{60}Co Therapy Beam," in *Proceedings of the IAEA International Symposium on Dosimetry in Radiotherapy* (IAEA, Vienna, 1988), Vol. 1, pp. 303–312.
[29]D. W. O. Rogers, G. M Ewart, A. F. Bielajew, and G. van Dyk, "Calculation of Contamination of the ^{60}Co Beam from an AECL Therapy Source," NRC Report No. PXNR-2710 (1985).
[30]K. Han, D. Ballon, C. Chui, and R. Mohan, "Monte Carlo simulation of a cobalt-60 beam," Med. Phys. **14**, 414–419 (1987).
[31]R. Mohan, C. Chui, and L Lidofsky, "Energy and angular distributions of photons from medical linear accelerators," Med. Phys. **12**, 592–597 (1985).
[32]S. S. Kubsad, T. R. Mackie, M. A. Gehring, D. J. Misisco, B. R. Paliwal, M. P. Mehta, and T. J. Kinsella, "Monte Carlo and convolution dosimetry for stereotactic radiosurgery," Int. J. Radiation Oncology Biol. Phys. **19**, 1027–1035 (1990).
[33]A. Kosunen and D. W. O. Rogers, "Beam Quality Specification for Photon Beam Dosimetry," Med. Phys. **20**, 1181–1188 (1993).
[34]E. L. Chaney, T. J. Cullip, and T. A. Gabriel, "A Monte Carlo study of accelerator head scatter," Med. Phys. **21**, 1383–1390 (1994).
[35]A. F. Bielajew, R. Mohan, and C. S. Chui, "Improved bremsstrahlung photon angular sampling in the EGS4 code system," National Research Council of Canada Report No. PIRS-0203 (1989).
[36]B. A. Faddegon, C. K. Ross, and D. W. O. Rogers, "Angular distribution of bremsstrahlung from 15 MeV electrons incident on thick targets of Be, Al and Pb," Med. Phys. **18**, 727–739 (1991).
[37]B. A. Faddegon and D. W. O. Rogers, "Comparisons of thick-target bremsstrahlung calculations by EGS4 and ITS," Nuc. Inst. Meth. A **327**, 556–565 (1993).
[38]A. F. Bielajew and D. W. O. Rogers, "PRESTA: The Parameter Reduced Electron-Step Transport Algorithm for Electron Monte Carlo Transport," Nuclear Instrum. Methods B **18**, 165–181 (1987).
[39]W. R. Nelson, H. Hirayama, and D. W. O. Rogers, "The EGS4 Code System," Stanford Linear Accelerator Center Report No. SLAC-265 (Stanford Calif) (1985).
[40]W. R. Nelson and D. W. O. Rogers, "Structure and Operation of the EGS4 code system," in *Monte Carlo Transport of Electrons and Photons Below 50 MeV*, edited by T. M. Jenkins, W. R. Nelson, A. Rindi, A. E. Nahum, and D. W. O. Rogers (Plenum, New York, 1989), pp. 287–306.
[41]A. J. Cook, "Mortran3 user's guide," SLAC Computa. Res. Group Tech. Memo. No. CGTM 209 (1983).
[42]R. Brun, O. Couet, C. Vandoni, and P. Zanarini, "PAW User Guide," CERN Computer Center, Geneva, Switzerland (1992).
[43]A. F. Bielajew and P. E. Weibe, "EGS-Windows - A Graphical Interface to EGS," NRCC Report No. PIRS-0274 (1991).
[44]Y. Namito, W. R. Nelson, S. M. Seltzer, A. F. Bielajew, and D. W. O. Rogers, "Low-energy x-ray production studies using the EGS4 code system," Med. Phys. **17**, (abstract) 557 (1990).
[45]J. Cygler and D. W. O. Rogers, "Monte Carlo calculations of tungsten mesh electron beam modifier," Phys. Med. Biol. A **39**, (abstract) 677 (1994).
[46]D. W. O. Rogers, "How Accurately Can EGS4/PRESTA Calculate Ion Chamber Response?," Med. Phys. **20**, 319–323 (1993).
[47]G. X. Ding, D. W. O. Rogers, and T. R. Mackie, "Calculation of backscatter into beam monitor chamber for electron beams," Med. Phys. **21**, (abstract) 923 (1994).
[48]B. A. Faddegon, C. K. Ross, and D. W. O. Rogers, "Forward Directed Bremsstrahlung of 10–30 MeV Electrons Incident on Thick Targets of Al and Pb," Med. Phys. **17**, 773–785 (1990).
[49]H. Kahn, "Use of different Monte Carlo sampling techniques," in *Symposium on Monte Carlo Methods*, edited by H. A. Meyer (Wiley, New York, 1956), pp. 146–190.
[50]E. D. Cashwell and C. J. Everett, *Monte Carlo Method for random walk problems*, (Pergamon, New York, 1959).
[51]B. Nilsson and A. Brahme, "Absorbed dose from secondary electrons in high energy photon beams," Phys. Med. Biol. **24**, 901–912 (1979).
[52]D. W. O. Rogers and A. F. Bielajew, "The Use of EGS for Monte Carlo Calculations in Medical Physics," Report No. PXNR-2692 (National Research Council of Canada, Ottawa, Canada K1A 0R6, 1984).
[53]D. W. O. Rogers, "Low energy electron transport with EGS," Nucl. Inst. Meth. **227**, 535–548 (1984).
[54]C. M. Wells, T. R. Mackie, M. B. Podgorsak, M. A. Holmes, N. Papanikoloau, P. J. Reckwerdt, J. Cygler, D. W. O. Rogers, A. F. Bielajew, D. G. Schmidt, and J. K. Muehlenkamp, "Measurements of electron dose distribution near inhomogeneities using a plastic scintillator detector," Int. J. Radiat. Oncol. Biol. Phys. **29**, 1157–1165 (1994).
[55]J. O. Deasy, "Electron energy and angular distributions in radiotherapy," PhD. thesis University of Kentucky (1992).
[56]P. Andreo and A. Fransson, "Stopping-power ratios and their uncertainties for clinical electron beam dosimetry," Phys. Med. Biol. **34**, 1847–1861 (1989).
[57]C. K. Ross and K. R. Shortt, "Energy Calibration of the 90-B Beam Line of the IRS Linac," NRC Report No. PIRS-0021 (1985).
[58]D. W. O. Rogers, "The role of Monte Carlo simulation of electron transport in radiation dosimetry," Int. J. Appl. Radiation and Isotopes **42**, 965–974 (1991).
[59]G. X. Ding, D. W. O. Rogers, and T. R. Mackie, "Mean energy, energy-range relationships and depth-scaling factors for clinical electron beams," submitted to Med. Phys., Apr. 1994.
[60]A. Kassaee, X. Chen, S. Ayyalasomayajula, and P. Bloch, "Electron beam characteristic with modified electron cones and scattering foil on a Varian 2100C accelerator," Med. Phys. **20**, (abstract) 887 (1993).
[61]A. Kassaee, M. D. Altschuler, S. Ayyalasomayajula, and P. Bloch, "Influence of cone design on electron beam characteristics on clinical accelerators," Med. Phys. **21**, 1671–1676 (1994).
[62]D. Sheikh-Bagheri, D. W. O. Rogers, and T. R. Mackie, "Monte Carlo model of the 6 MV photon beam from a SL25 accelerator," Med. Phys. **21**, 1368 (abstract) (1994).
[63]G. X. Ding and D. W. O. Rogers, "Energy spectra, angular spread, and dose distributions of electron beams from various accelerations used in radiotherapy," NRC Report, No. PIRS-439 (NRC, Ottawa, March 1995). (See also http://www.irs.inms.nrc.ca/inms/irs/papers/PIRS439/pirs439.html).