

The Geant4 Bertini Cascade



D.H. Wright*, M.H. Kelsey, For the GEANT4 Hadronics Working Group

SLAC National Accelerator Laboratory, Stanford, CA, USA

ARTICLE INFO

Article history:

Received 30 June 2015

Received in revised form

16 September 2015

Accepted 18 September 2015

Available online 30 September 2015

Keywords:

Hadronic physics models

Simulation

Geant4

ABSTRACT

One of the medium energy hadron–nucleus interaction models in the GEANT4 simulation toolkit is based partly on the Bertini intranuclear cascade model. Since its initial appearance in the toolkit, this model has been largely re-written in order to extend its physics capabilities and to reduce its memory footprint. Physics improvements include extensions in applicable energy range and incident particle types, and improved hadron–nucleon cross-sections and angular distributions. Interfaces have also been developed which allow the model to be coupled with other GEANT4 models at lower and higher energies. The inevitable speed reductions due to enhanced physics have been mitigated by memory and CPU efficiency improvements. Details of these improvements, along with selected comparisons of the model to data, are discussed.

© 2015 Elsevier B.V. All rights reserved.

1. Introduction and overview

The GEANT4 Bertini Cascade, as it has come to be known, has an old lineage, dating back to some of the earliest nuclear Monte Carlo codes. The code first appeared in version 5.0 of the GEANT4 toolkit [1,2] as a C++ translation and re-engineering [3] of the Fortran version of the intranuclear cascade INUCL [4]. A short description of this model, its implementation and some comparison to data is given in Ref. [5]. INUCL in turn was based on the methods of Bertini [6].

Although some of its key elements have remained, translations, improvements and extensions over the years have changed the code to such an extent that it bears little resemblance to the earlier models. The original INUCL and GEANT4 Bertini codes treated nucleon and pion interactions with nuclei up to incident energies of 3 GeV. The Bertini code has since been extended to handle all long-lived hadrons (p, n, π , K, Λ , Σ , Ξ , Ω) as projectiles, and in connection with other models, to implement lepton– and gamma–nuclear interactions. Its applicable incident energy range has been extended to ~ 20 GeV, although little validation of the model has been done above 15 GeV and most use cases do not exceed 10 GeV.

Interfaces to other GEANT4 models have also been developed. One such interface allows the Bertini-style cascade access to a more detailed treatment of pre-equilibrium and evaporation physics. Another allows it to handle the secondary particles and residual nuclei produced in high energy models. Still others enable

the Bertini cascade to provide improved hadronic final states following certain kinds of particle capture.

Throughout its recent development, the Bertini-style cascade has been used successfully for simulation in the BaBar, Alice, ATLAS, CMS, and LHCb experiments, to name a few. Their demands for more efficient simulation have driven many improvements in the code, including speed optimizations and minimization of repeated memory access and footprint.

A diagrammatic summary of the GEANT4 Bertini algorithm is shown in Fig. 1. The various stages of this algorithm, and the internal data required to support it, are described in the next section.

2. Physics modeling

2.1. Intranuclear cascade

At the core of the GEANT4 Bertini model is the concept of the intranuclear cascade of hadrons and nucleons produced through a series of interactions within the nucleus. It is essentially a classical model solving on average the Boltzmann equation for the transport of a particle through a “gas” of nucleons. The nuclear medium may be considered as a gas if the effective nucleon size is small and there are few collisions.

The traditional range of applicability for such a model is approximately

$$200 \text{ MeV} < E < 3 \text{ GeV} \quad (1)$$

where E is the kinetic energy of the particle initiating the cascade.

* Corresponding author. Tel.: +1 650 926 4743.

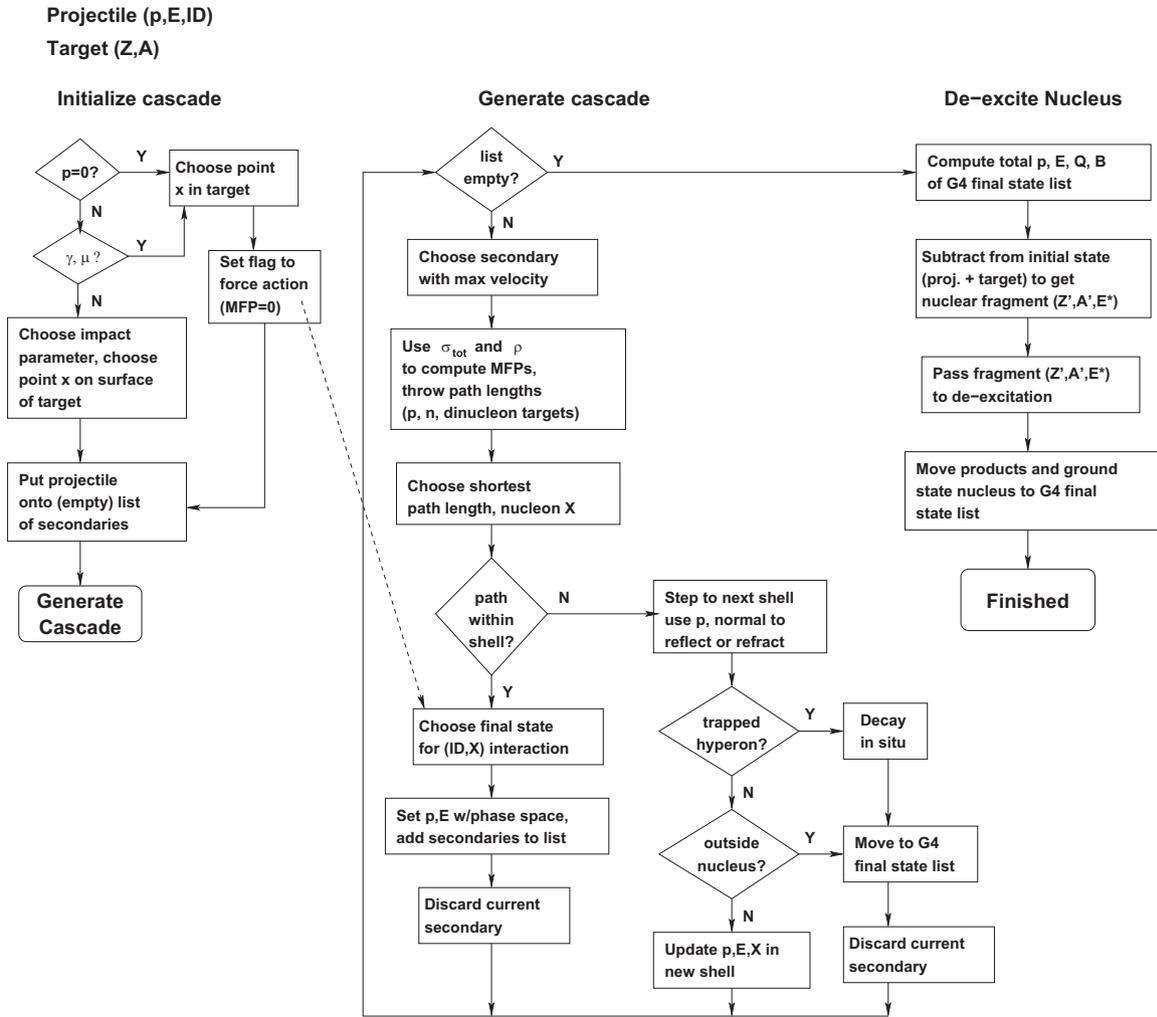


Fig. 1. General algorithm for GEANT4 Bertini-style cascade. At left, the projectile hadron (ID) with momentum p and energy E in the lab frame, and the target nucleus are configured for the initial collision. Captured hadrons and lepto- or photonuclear interactions are treated specially, with the projectile interacting at a random location X in the nucleus (see Sections 3.3, 3.4, and 4.3). The cascade development is shown in the middle section. Here MFP refers to the mean free path from which path lengths are sampled. A cumulative list of particles is maintained, and for each iteration the fastest particle is chosen to interact, which may either produce secondaries, or be propagated one “step” through the nucleus (Section 2.4). Particles which interact or escape the nucleus are removed from the list; this section is repeated until no cascade particles remain. Finally, at right is shown the build-up (using four-momentum, charge Q and baryon conservation B) and de-excitation of the compound nuclear fragment remaining at the conclusion of the cascade (Section 2.2).

At 200 MeV the deBroglie wavelength of a projectile nucleon is roughly the same as the distance between nucleons in the target nucleus. Above that projectile energy a cascade is appropriate since the nucleus may be treated as a gas of individual nucleons. Below that energy, the properties of the nucleus as a whole begin to dominate, and precompound and evaporation models should be used instead.

Above 3 GeV, a host of effects lead to a breakdown of the classical cascade view. One of these is the Lorentz contraction of the target nucleus in which the nucleons come too close together to be considered a gas. Two more are the increasingly forward-peaked final state angular distributions of hadron–nucleon interactions, and the increased time taken for secondaries to become quantum mechanically distinct from one another. Both of these reduce the probability of subsequent intranuclear interactions. Nevertheless, it is possible to modify a cascade model so that it can be applied at energies well beyond 3 GeV.

The cascade is generated when an incident particle interacts elastically or inelastically with individual protons and neutrons in the nucleus, producing a number of secondaries. The secondaries interact with other protons and neutrons, producing further secondary particles with energies consistent with the individual

interactions. All of these secondaries propagate through the nucleus, interacting in turn until the last secondaries either escape the nucleus or stop and are absorbed. At this point the residual nucleus is de-excited using the pre-equilibrium and evaporation models. Clusters of nucleons within the initial nucleus are not generally considered, except in the case of pion absorption which requires dinucleons.

The above approach is similar to that of many cascade codes now available [7–10]. Another common approach followed here is to treat the elementary collisions during the cascade, to first order, as if they occurred in vacuum. It is this feature which provides the model with its essential extensibility. As long as the elementary interactions of a given incident particle type with protons and neutrons are known, the cascade can be extended to handle it. Nuclear medium effects can then be applied as corrections.

2.2. Pre-equilibrium and evaporation models

When all cascade particles have either left or been trapped in the nucleus the de-excitation codes included with the Bertini model take over. There is no specific energy at which this occurs,

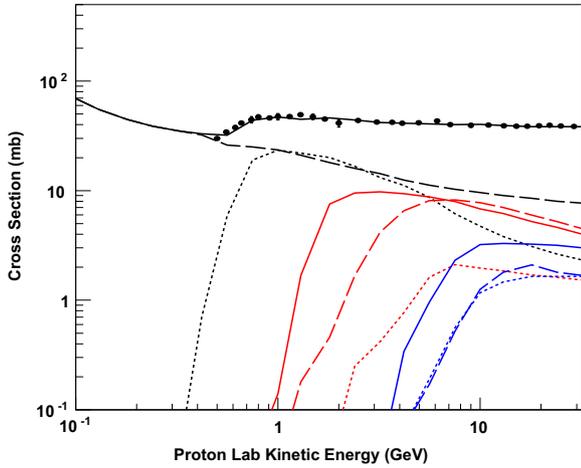


Fig. 2. Total and final state multiplicity cross-sections for n -body final states of the $p p$ reaction. Black curves represent the total (solid), two-body (dashed) and three-body (dotted) cross-sections. Red curves represent the four-body (solid), five-body (dashed) and six-body (dotted) cross-sections. Blue curves represent the seven-body (solid), eight-body (dashed) and nine-body (dotted) cross-sections. Filled circles represent a subset of the $p p$ total cross-section data [17]. (For interpretation of the references to color in this figure caption, the reader is referred to the web version of this paper.)

but it corresponds roughly to a projectile (not a secondary) energy of 200 MeV.

Secondaries from the cascade with energies below the nuclear potential are trapped or stopped in the nucleus. These particles are not placed in the final state particle list but are eventually decayed or recombined with the nucleus in the pre-equilibrium phase. The particles destined to recombine from a collection of particle-hole states described by the Griffin exciton model [11]. These states subsequently de-excite using parameterized level densities which are functions of Z and A . Particle emission from this process is assumed to be isotropic in the rest frame of the exciton system.

Following this stage the nucleus is further de-excited by Fermi breakup for nuclei with $A < 12$. All other nuclei undergo equilibrium evaporation which is performed by the Weisskopf statistical model [13] as implemented by Dostrovsky et al. [14]. For some heavy nuclei ($A > 100$) fission may be performed [12]. Relative probabilities for the evaporation of nucleons, light ions (d , ${}^3\text{He}$, t , α), and fission are used to implement competition between these processes.

The pre-equilibrium and evaporation models used by the current Bertini code are very similar to those used by many other cascade codes [7,8], and within GEANT4 have changed little since their translation into C++.

2.3. Extraction and implementation of the hadron–nucleon cross-sections

The total hadron–nucleus inelastic cross-sections are calculated in classes outside of the Bertini code and are the result either of parameterizations of data or calculations where there are no data. These cross-sections in part determine the location of the interaction in the detector.

Once the cascade is invoked, the total hadron–nucleon cross-section is used to determine where in the nucleus the incident hadron interacts. The partial hadron–nucleon cross-sections determine the type and multiplicity of the secondaries of this interaction. In the GEANT4 implementation of the Bertini model, these elementary cross-sections are taken from the CERN-HERA compilations [15–18] of hadron–nucleon interaction data. These cover the total and partial cross-sections for gammas, protons, antiprotons, charged pions, kaons, and hyperons incident upon

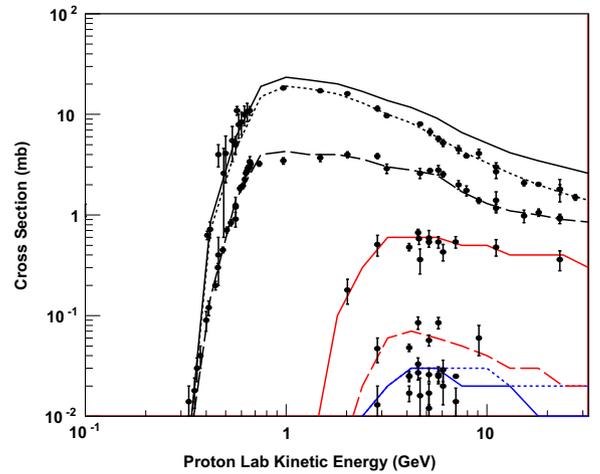


Fig. 3. Exclusive and summed three-body final state cross-sections for $p p$ scattering. Black curves represent the total three body final state (solid), the $p p \rightarrow p n \pi^+$ (dotted), and the $p p \rightarrow p p \pi^0$ (dashed) cross-sections. The solid red curve represents $p p \rightarrow p \Lambda K^+$ which has been multiplied by a factor of 10 for clarity. The dashed red curve represents $p p \rightarrow n \Sigma^+ K^+$. The blue curves represent the $p p \rightarrow p \Sigma^0 K^+$ (solid) and $p p \rightarrow p \Sigma^+ K^0$ (dashed) cross-sections. Filled circles represent experimental data. (For interpretation of the references to color in this figure caption, the reader is referred to the web version of this paper.)

protons in most cases and upon neutrons in a few cases. The data cover all energies up to 30 GeV.

The partial cross-sections used in the model are extracted from the tables and plots in the data compilations. For each exclusive reaction data set, a smooth curve is drawn through the data points and values of the curve are taken at 30 intervals roughly logarithmic in lab energy. Final state multiplicity cross-sections are then accumulated by summing the exclusive cross-section arrays of a given multiplicity. The sum of all partial cross-sections thus extracted is constrained by the measured total cross-section for the incident particle type. The error introduced by this procedure is estimated to be about the size of the error bars on the data used, or about 10%, increasing to 20–30% in energy regions where data is sparse. Above 20 GeV, many partial cross-sections could be off by factors of two.

A result of this method is shown in Fig. 2, which displays the total and final state multiplicity cross-sections for the $p p$ reaction along with a subset of its corresponding data.

The multiplicity of the hadron–nucleon reaction can now be sampled: the probability of a given final state multiplicity is the multiplicity cross-section linearly interpolated in energy and divided by the total cross-section.

With the multiplicity in hand, the exclusive final state must now be sampled. The same procedure used to find the multiplicity is used to sample the final state, except that the probability is given by the ratio of the exclusive cross-section to the multiplicity cross-section. Fig. 3 shows the total three-body cross-section for the $p p$ reaction along with the six exclusive final state cross-sections that comprise it.

A large number of partial cross-section data sets are required to implement the full range of reactions, especially when strangeness exchange and strangeness production channels are included. The data sets are organized according to incident particle type, target nucleon and final state multiplicity.

The number of final state particles that need to be considered depends on the desired validity range of the model. In the typical cascade range ($E < 5$ GeV) it is sufficient to include up to six-body final state partial cross-sections. In this region the cross-sections of higher multiplicity states are a negligible part of the total, but at 10 GeV or more they become significant and up to nine-body final

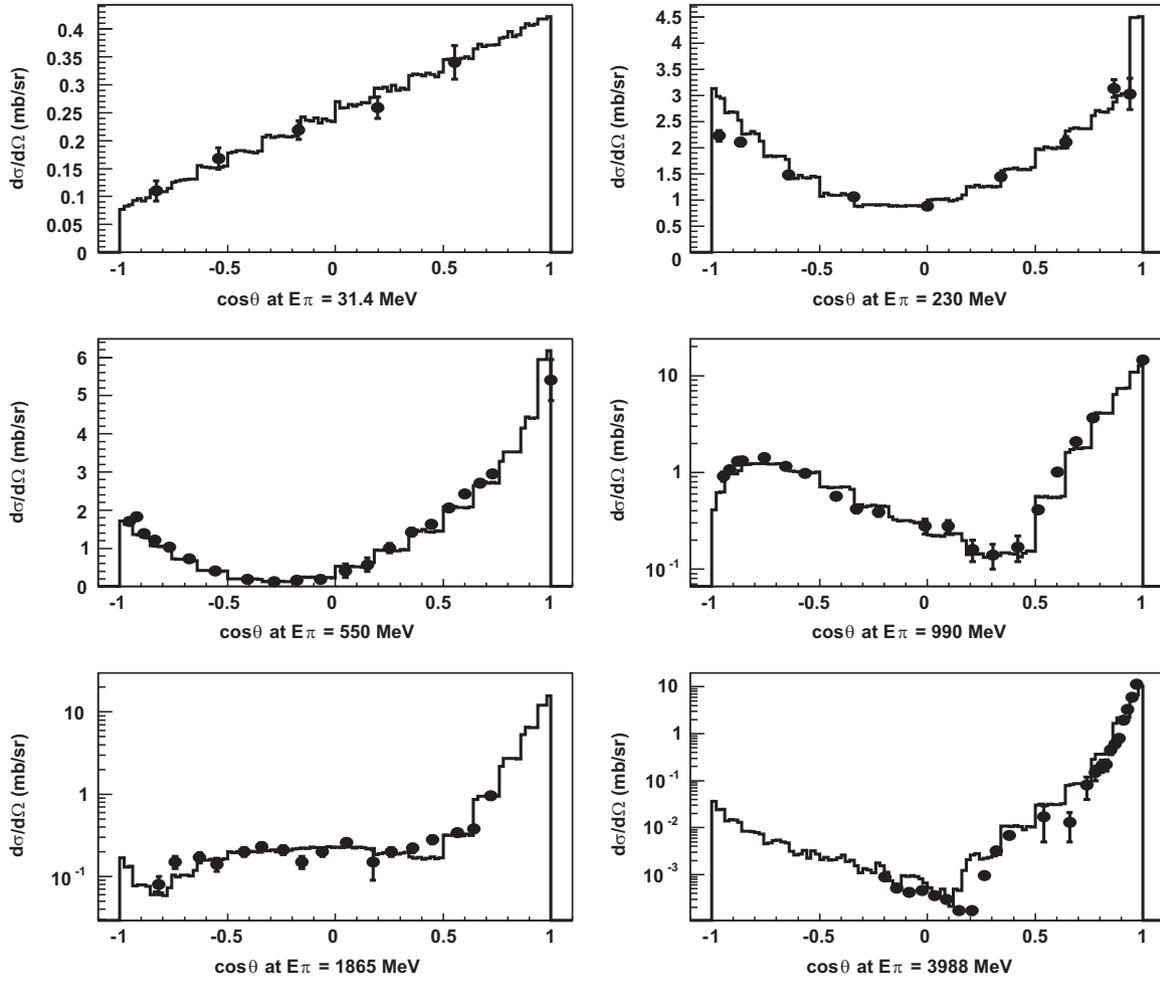


Fig. 4. Angular distributions for π^- p elastic scattering. Points are experimental data [21–27]. Below 2.6 GeV incident energy, SAID phase shift parameterizations are used. Above 2.6 GeV, experimental data were averaged over 10° bins and tabulated for interpolation.

states must be included in order to account for the total cross-section.

In many cases there are no published data for a given exclusive channel. Relatively few cross-sections are available with the neutron as either target or projectile, and even for proton targets the number of partial cross-sections available decreases with increasing incident energy and final state multiplicity. Isospin invariance allows some neutron cross-sections to be constructed from known proton cross-sections. At higher energies, where final state multiplicities increase, typically only one or two of the exclusive cross-sections at a given multiplicity have been published. In this case the missing cross-sections were estimated by taking the approach of Iljinov et al. [19], simplified to take into account only the combinatorics of the final state. For example, the partial cross-section

$$p p \rightarrow p n \pi^+ \pi^0 \pi^0 \pi^0 \quad (2)$$

can be estimated by scaling the known partial cross-section

$$p p \rightarrow p n \pi^+ \pi^+ \pi^- \pi^0 \quad (3)$$

by the factor

$$\frac{n_{\pi^+}! n_{\pi^-}! n_{\pi^0}!}{n_{\pi^+}! n_{\pi^0}!} = \frac{2}{6}. \quad (4)$$

The energy dependences of the unmeasured cross-sections are assumed to follow the generic behavior observed for most of the measured cross-sections: a steep rise from threshold followed by a slow exponential decay as energy increases. Here again, the total

measured cross-section is useful to constrain the unknown cross-sections.

Having determined the partial cross-sections as a function of energy it remains to parameterize the angular distributions of the produced particles. For 2-body final states, the SAID [20] phase shift parameterizations are used wherever possible. These cover nucleon–nucleon and pion–nucleon elastic and charge exchange reactions up to about 2.6 GeV lab energy. For particles above the SAID energy range, existing data are averaged over 10° bins and tabulated for later interpolation. The parameterization describing π^- p elastic scattering using both SAID phase shifts and tabulated data is shown in Fig. 4. This parameterization replaced the old Barashenkov polynomial representation used in GEANT4 releases older than 10.1, and which failed to reproduce data at energies above 2–3 GeV.

For other projectiles, data [16,18] indicate that the full angular range can be approximately described by the function

$$Ae^{-Bp_{cm}^2(1-\cos\theta)}, \quad \theta \leq \theta_0 \\ Ce^{Dp_{cm}^2(1-\cos\theta)}, \quad \theta > \theta_0 \quad (5)$$

which is required to be continuous at θ_0 . p_{cm} is the center of mass momentum and θ is the scattering angle. The parameters B , D and θ_0 depend on the projectile energy. This parameterization is able to describe reasonably the backward and forward scattering present at lower energies, as well as the very forward scattering at higher energies. For the three-body and n -body final states of all

incident particle types, the scattering angle is sampled from uniform phase space.

In the partial cross-section approach resonances, such as the Δ , N^* and others, are not treated explicitly. Rather the energy and angle dependence of decay products in the resonance region is taken directly from measured partial cross-sections.

2.4. Nuclear model and medium corrections

The extended Bertini model belongs to the class of so-called “space-like” intranuclear cascades, and implements the target nucleus as a continuous medium. The nuclear potential is represented by a set of concentric, constant-density shells whose depths approximate the Woods–Saxon shape. The number of shells varies from 1 for light nuclei, to 3 for medium nuclei, to 6 for heavy nuclei. The functional form for the outer shell edges of a three-shell nucleus is

$$r_i = c \ln \left(\frac{1 + e^{-R_N/c}}{\alpha_i} - 1 \right) + R_N \quad (6)$$

where $R_N = aA^{1/3} + bA^{-1/3}$, c is proportional to the nuclear skin depth, and the $\alpha_i = [0.7, 0.3, 0.01]$. The depth of each shell is found by integrating the Woods–Saxon shape between the shell radii, and the density by dividing the number of protons or neutrons by the resulting volume.

The Fermi momentum $p_{F,i}$ is proportional to the cube root of the density of the i th shell calculated above and is used to calculate the potential seen by nucleons traveling in that shell,

$$V_{N,i} = \frac{p_{F,i}^2}{2M_N} + B_N(A, Z) \quad (7)$$

where B_N is the nucleon binding energy. This potential follows from the Pauli exclusion of nucleons and does not apply to other hadrons. In this model, non-nucleons see a potential which has constant depth over all shells with the value

$$V_{hadron} = \begin{cases} 7 \text{ MeV} & \text{for pions} \\ 15 \text{ MeV} & \text{for kaons} \\ 30 \text{ MeV} & \text{for hyperons.} \end{cases} \quad (8)$$

Hadrons travel in straight line segments through the potential. When a shell boundary is encountered the hadron is either reflected or transmitted (with refraction) according to the differential potential of the adjacent shells.

A hadron may interact when it reaches its proposed interaction point which is determined by sampling the path length distribution. This distribution is governed by the mean free path calculated from the shell density and total cross-sections discussed earlier.

Because the mean free path calculation relies on the parameterized or tabulated free-space cross-sections, the sampled interaction point is not influenced by nuclear structure effects. These are taken into account after the fact by the inclusion of Pauli blocking, Fermi motion, hard-core nucleon repulsion and the trailing effect. These effects combine to reduce the effective cross-section, and increase the path length, of the hadrons passing through the nucleus.

Pauli blocking and Fermi motion were included in the original Bertini model and depend as usual on the Fermi momentum, which in turn depends on the local nuclear density.

Hard-core nucleon repulsion prevents other nucleons from approaching closely. To reflect this in the model, secondaries produced in a collision are not allowed to interact within a distance D_{HC} which is roughly the size of the nucleon radius.

The “trailing effect” [28] was recently added: it takes into account the removal of nucleons during the cascade, thereby

reducing the local nuclear density seen by subsequent particles passing through. At an interaction point determined for a given track, the locations of all previous collisions are queried. If the proposed interaction point is within a distance D_T of any past collision point, the interaction is canceled and the track passes through unchanged. D_T is a distance roughly the same size as the nucleon radius.

One feature implemented in some cascade codes is the formation time or formation zone. At high enough energies the products of a reaction do not become quantum mechanically distinct from one another until they have traveled a significant distance through the nucleus. Over this distance they cannot interact and the resulting interaction probability is reduced. The formation zone was implemented in the extended Bertini cascade, but no significant improvement in the comparison to data was observed, so it was removed.

The parameters D_{HC} and D_T are not well-determined by experiment and are usually treated as tunable by most cascade codes. The nuclear radius, R_{nuc} , is well measured by electron scattering, but variations of this parameter are sometimes made to account for the longer range of the electromagnetic interaction or to include more of the nuclear density tail. The current Bertini code maintains the parameter values found in the original INUCL code, but in the near future a global tuning of these parameters will be attempted in order to improve agreement with data.

3. Extensions

3.1. Higher energy

The first extension of the GEANT4 Bertini cascade enabled its use for particles of higher incident energy and appeared in version 6.0. Until then no satisfactory hadronic model existed to cover the energy gap between the traditional cascade codes and the parton string models. The extension for nucleons and pions up to 30 GeV more than filled that gap. As mentioned above, this required that final states with higher multiplicities be added.

3.2. Strange particles

The second extension, in GEANT4 version 7.1, was the inclusion of strange particles, that is kaons and hyperons, as projectiles and reaction products. The approach followed here was the same as that for pions and nucleons: generate final states based on measured partial cross-sections and not on the formation and decay of resonances. The large number of possible strange resonances made the latter choice intractable.

The partial cross-sections for strange particle induced reactions were taken from the CERN data collections [16,18]. For higher multiplicity final states, the data is rather sparse, especially at higher incident energies. For this reason, the partial cross-sections for strange particles on nucleons include final state multiplicities only up to seven. As in the non-strange case, quark symmetries, combinatorics and total cross-section constraints were used to construct the unmeasured partial cross-sections.

An added complication of strange particle reactions in nuclei is that the final state may include a hypernucleus. At this time GEANT4 cannot track hypernuclei, so the Bertini cascade does not produce them. Instead of embedding a sufficiently low energy hyperon in the nucleus, the model simply decays any such particles which appear at the end of the intranuclear cascade and adds the products of the decay to the final state. Non-hadronic decay products such as neutrinos, muons, etc., are not handled by Bertini, but instead sent directly to tracking.

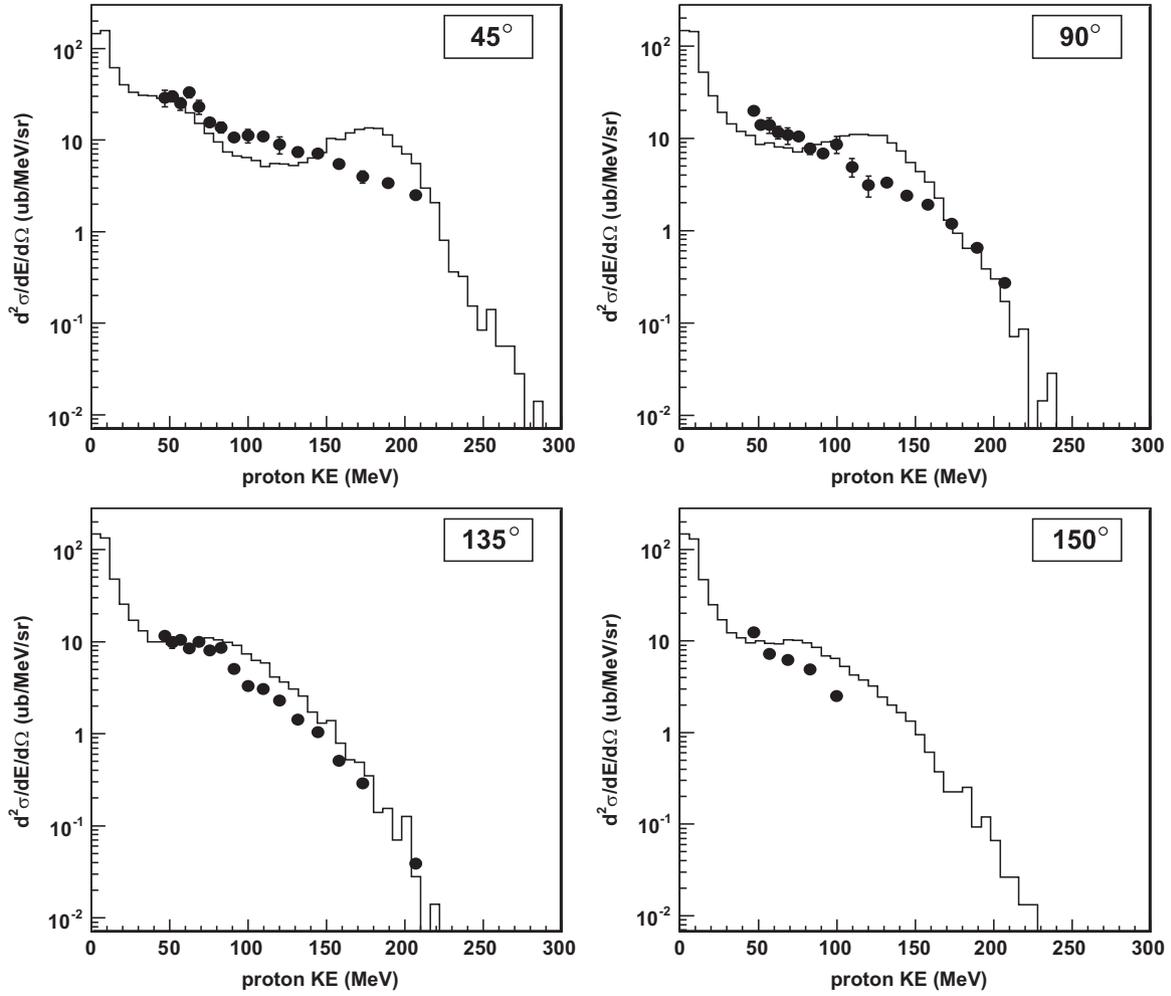


Fig. 5. Proton doubly differential cross-section in units of $\mu\text{b}/\text{MeV}/\text{sr}$ for 300 MeV gammas on copper. The Bertini model predictions are represented by histograms, and the data, from Schumacher et al. [29], are represented by dots.

3.3. Photonuclear

Enough γ - p data exist [18] that the partial cross-section approach can be used successfully in a γ -initiated cascade. Once the initial γ interacts inelastically within the nucleus, the hadronic secondaries sustain a cascade.

Because the γ - p cross-section is about two orders of magnitude smaller than a hadron-nucleon cross-section, the mean free path of the γ is longer than the diameter of the target nucleus and an interaction is unlikely. However, when the GEANT4 process determines from the total cross-section that an interaction has occurred, the model is required to produce a final state. In the case of gammas, the model forces the interaction at a point chosen by sampling from a nuclear-density-weighted path which follows the initial direction of the γ .

The γ - p data used to set up the partial cross-sections cover incident energies from 0 to about 20 GeV, but there is little detail in the region below 30 MeV where the Giant Dipole Resonance (GDR) occurs. The total γ -nuclear cross-section does include this feature and it is reproduced surprisingly well by the model considering that the GDR is a collective nuclear state while the Bertini cascade deals only with incoherent scattering.

At a more traditional cascade energy, 300 MeV, a comparison of the extended Bertini prediction and data for gammas on a copper target is shown in Fig. 5.

3.4. Muon capture

Although the Bertini cascade was designed to deal with hadronic interactions, it is clear from the previous section that as long as a prescription exists for exclusively generating hadrons in the final state, the initial particle does not have to be a hadron. This is true for the case of captured muons (μ^-) in which the muon is absorbed on a single proton or on a proton-proton or proton-neutron pair.

The Bertini model does not deal with the atomic cascade phase of muon capture. It assumes that the μ^- has no kinetic energy and that its wave function is spread uniformly over the entire target nucleus. The interaction point can thus be chosen at random within the nuclear volume. A more sophisticated sampling involving the nuclear wave function or nuclear density has not yet been implemented.

At this point absorption on a proton may occur according to the $\mu^- p \rightarrow n \nu_\mu$ cross-section in free space, or absorption on a nucleon pair may occur according to the same parameterized cross-section used for pion absorption on a nucleon pair. In the latter case the final state consists of two nucleons and a muon neutrino whose momenta are given by phase space. Muon capture on a nucleon pair dominates, occurring roughly three times as often as capture on a proton. After capture, the intranuclear cascade proceeds with the propagation of the ejected nucleons.

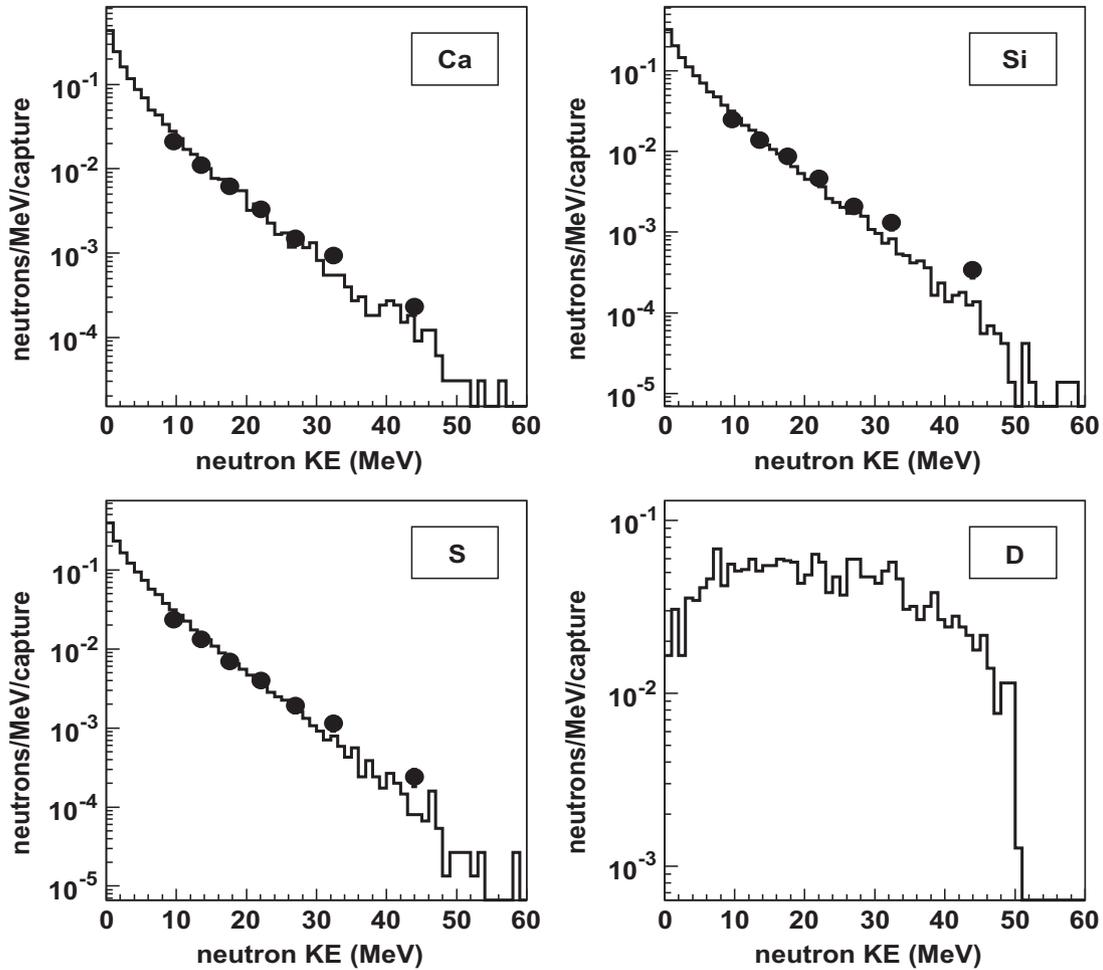


Fig. 6. Neutron kinetic energy spectrum from μ^- capture on Ca, Si, S and deuterium. The histograms represent the Bertini model predictions and the data (dots) are from Sundelin and Edelstein [30].

Experimental data for muon capture on several nuclei are shown in Fig. 6 where they are compared to the Bertini prediction. Although there are no data shown for deuterium, this target was included to demonstrate the phase space nature of the neutron energy spectrum.

3.5. Coalescence

Typical intranuclear cascade models do not produce energetic deuterons, tritons, ^3He particles or α particles, except during the pre-equilibrium and de-excitation phases at the end of the cascade. As a result there is usually a deficit of these particles at higher energies. One solution to this problem [31] is to take nucleons in the cascade stage and group them into light nuclei if their momenta are close to one another.

This approach was taken in the extension of the GEANT4 Bertini cascade. All nucleons present at the end of the cascade are sorted according to their relative momenta. If the momenta of a proton and neutron are found to be within 90 MeV/c of one another, they are removed from the final state list and replaced with a deuteron. The remaining nucleons are considered as triplets and quadruplets; tritons and alphas are formed if their constituent nucleons differ in momenta by less than 108 and 115 MeV/c, respectively. The relative momentum values used come from Toneev and Gudima [31]. No clustering beyond four nucleons is done, except for the nuclear fragments produced in Fermi break-up.

4. Interfaces

4.1. Precompound model

As mentioned in Section 2.2, the pre-equilibrium and nuclear evaporation models native to the Bertini code are relatively simple and highly parameterized. A more detailed option, `G4PreCompoundModel` [32] offers better precision at energies below 75 MeV, although at the cost of greater CPU time. In order to use the extended Bertini cascade for intermediate energies and a more precise model at low energies, an interface connecting Bertini (the main `G4CascadeInterface` class) and `G4PreCompoundModel` was developed. It can be activated in user code by invoking the `usePreCompoundDeexcitation()` method of `G4CascadeInterface`.

4.2. High energy

For high energy reactions ($E \gtrsim 10$ GeV) GEANT4 offers the quark gluon string (QGS) [33] and Fritiof (FTF) [34,35] models. Each of these models simulates an initial hadron–nucleus interaction which leaves the target nucleus in a highly disorganized state. The energy of the residual nucleus is often too high to be de-excited by a pre-equilibrium model and one alternative is a cascade model. An interface to use the cascade models as a “back-end” to the high energy models has been available for many years in GEANT4 but an implementation using the Bertini cascade is relatively recent.

The method `G4CascadeInterface::Propagate()` takes the secondaries of the initial high energy interaction, with their points

of origin and formation times, and inserts them into the Bertini nuclear model, thus initiating a cascade. The cascade then proceeds as before through the pre-equilibrium and nuclear de-excitation stages. This method is currently in testing and has not yet been released for use.

4.3. Interaction of stopped hadrons

Negatively charged hadrons coming to rest can be captured by a nucleus and leave it in an excited state. The excitation energy is generally too low to initiate an intranuclear cascade, but the pre-equilibrium stage of the model can handle the interaction by treating the captured particle as part of an exciton particle-hole state.

The pre-equilibrium module of the Bertini code is also used by other models for stopped π^- , K^- and Σ^- . For the π^- case, simple absorption occurs followed by the decay of the particle-hole state and nuclear evaporation. Kaons and sigmas are assumed to interact with a single nucleon, with available two-body final states sampled randomly from the zero-energy bins of the tabulated $K^- p$ and $K^- n$ partial cross-sections. The zero-energy values are shown in Tables 1 and 2. The resulting secondaries are propagated in the nuclear potential. When strange secondaries, most often hyperons, become trapped in the nuclear potential (Eq. (8)), they decay weakly, and their decay products are added to the cascade.

4.4. Electro- and muon-nuclear reactions

The extended Bertini model treats lepto-nuclear reactions as a virtual photon exchanged between the lepton and a nucleon in the target. At typical cascade energies it is safe to ignore any weak component in the exchange. If the virtual photon mass modulus is small compared to the energy transfer, the virtual photon can be treated as real and then interacted within the nucleus as mentioned in Section 3.3.

For incident electrons and positrons `G4ElectroVDNuclearModel` generates an electromagnetic vertex at which a virtual photon is created. This is done using the `G4ElectroNuclearCrossSection` [36] class to sample the nuclear scattering in the two-dimensional momentum and energy transfer space to get the virtual photon four-momentum. The Weiszäcker-Williams [37] method is then used to obtain from this an equivalent flux of real photons.

If the real photon energy is below 10 GeV, the extended Bertini cascade is used to initiate the intranuclear cascade. Higher energy photons are converted to a π^0 and then given to the FTF high energy model for interaction. This step exploits the fact that high energy photons act like hadrons. The 10 GeV boundary between

the two models represents the highest energy gammas for which the Bertini cascade has been tested for incident photons.

The procedure for muon-nuclear scattering is similar. The `G4MuonVDNuclearModel` uses the Borog and Petrukhin [38] muon-nuclear cross-sections to generate the electromagnetic vertex and virtual photon. The hadronic vertex is treated in exactly the same way as the electro-nuclear case.

5. Speed and memory optimization

One of the goals of the extended Bertini model was to improve its physics content without unduly increasing its computation time and memory footprint. This was accomplished through efficient coding and reduction of time-consuming physics detail where possible.

5.1. Physics choices

A list of the time-saving physics choices for the extended Bertini model is given here. For each item it was verified that the simplifying assumption made did not negatively impact the agreement with data.

- A relatively simple nuclear model was chosen. The individual nucleons making up the nucleus are not explicitly represented unless a scattering occurs on one of them. It was also deemed unnecessary to implement a smooth Woods-Saxon nuclear density, to avoid time-consuming (multistep) propagation through a varying potential.
- The most precise way to propagate hadrons through a nuclear potential would be to solve the equations of motion and step the particles along the resulting curved trajectories. The `GEANT4 Binary Cascade` code does this [10]. Propagation in straight line segments, however, is the faster method and was chosen.

The above two items were already features of the INUCL model. Two new additions assume that the energies and angles of secondaries produced in the nucleus need not be determined to high precision. Any graininess due to this approximation will be washed out over the course of multiple collisions within the nucleus.

- Linear interpolation, instead of quadratic, cubic or logarithmic, is used to retrieve the total and partial cross-sections from their look-up tables. This shows up, in Figs. 2 and 3 for example, as cross-section curves which are not perfectly smooth.
- Large angle bins are used when preparing the angular distribution look-up tables from data. Ten-degree bins were found to be sufficiently small. An artifact of this can be seen in the stair-step behavior of the histograms of Fig. 4.

Table 1

Two-body final states sampled for capture of K^- on protons (left) and neutrons (right), with indicated probabilities.

$K^- p \rightarrow$	P
$K^- p$	0.13361
$\bar{K}^0 n$	0.07355
$\Lambda \pi^0$	0.05812
$\Sigma^+ \pi^-$	0.15441
$\Sigma^0 \pi^0$	0.20559
$\Sigma^- \pi^+$	0.37300
$K^- n \rightarrow$	P
$\Lambda \pi^-$	0.00100
$\Sigma^0 \pi^-$	0.00040
$\Sigma^- \pi^0$	0.00032

Table 2

Two-body final states sampled for capture of Σ^- on protons (left) and neutrons (right), with indicated probabilities.

$\Sigma^- p \rightarrow$	P
Λn	0.35750
$\Sigma^0 n$	0.25888
$\Sigma^- p$	0.24563
$\Sigma^- n \rightarrow$	P
$\Sigma^- n$	0.13799

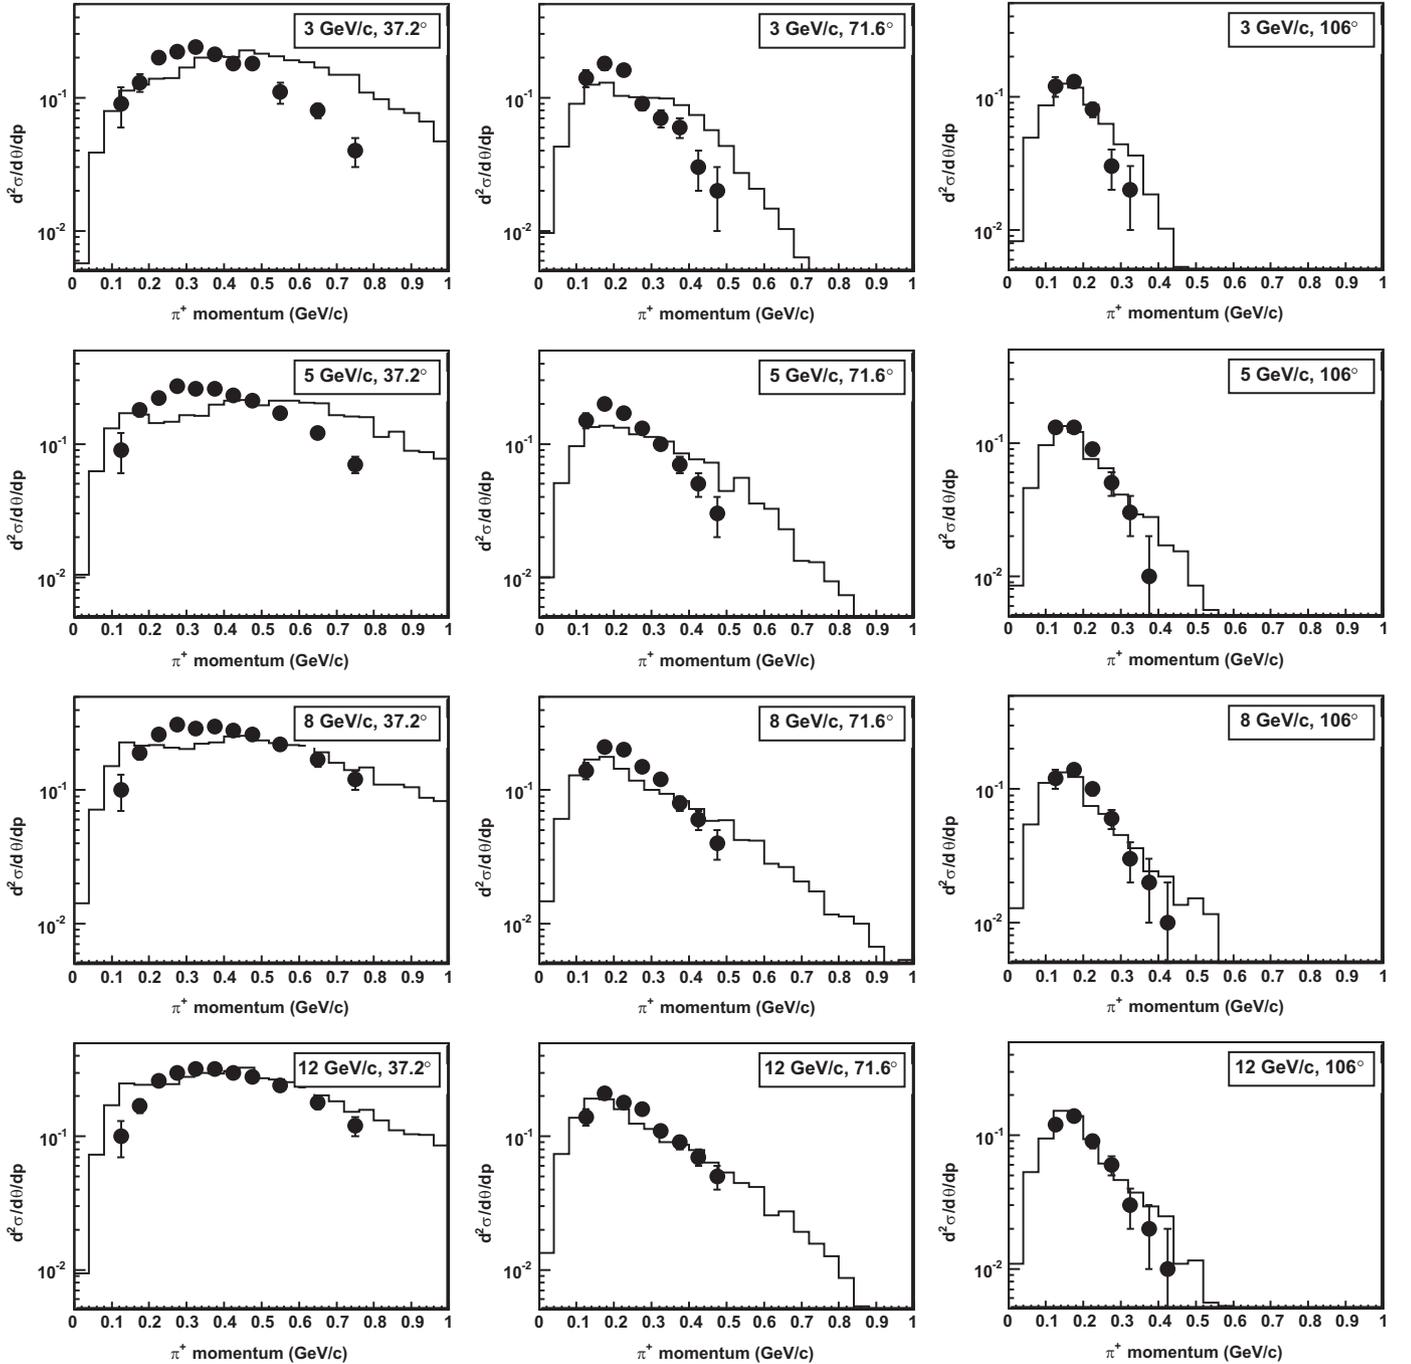


Fig. 7. Doubly differential cross-section in units of b/rad/GeV/c for $p+C \rightarrow \pi^+ X$ at 3, 5, 8 and 12 GeVp with lab angles 37.2° , 71.6° and 106° , respectively. Histograms represent the Bertini model predictions and dots represent the HARP [43] data.

5.2. Efficient coding

The initial implementation of the GEANT4 Bertini cascade was essentially a “translation” of the INUCL FORTRAN-77 code into C++ syntax and basic class structure. While functional (see Section 6), the code suffered operationally from both CPU and memory usage. In 2010, a substantial effort was made to reduce sources of unnecessary overhead throughout the Bertini code base; the results of these efforts were included in the GEANT4 9.4 release.

The CMS Collaboration [39] reported that the Bertini cascade contributed approximately 25% of the total CPU time and a majority of the memory usage in their detector simulation. They made these observations using the *IgProf* [40] profiling tool, which

measures CPU usage, memory allocation (create/delete cycles), and memory usage, all at the granularity of individual class function calls. This tool was used to identify “hot spots” in the Bertini code, as well as elsewhere in the GEANT4 software, as targets for improvement or redesign.

The primary source of both CPU and memory usage in the GEANT4 Bertini cascade code was memory “churn”, the frequent and unnecessary creation and deletion of data objects during code execution. This leads to fragmentation of memory, and reduced CPU efficiency. Churn also tends to increase overall CPU usage, due to the work required for allocation, deallocation and copying of data objects and their contents.

Significant improvements in the structure of the code essentially eliminated memory churn in the GEANT4 Bertini code. Most instances of data structures passed and returned by functions have been replaced by reusable data buffers shared across functions, and passed by reference (both read-only and modifiable) where necessary. As a result, the memory churn for a representative benchmark (1000 p -Pb interactions) was reduced from 210 MB (210 kB/interaction) with 2000 page faults to 17 MB (17 kB/interaction) with zero page faults. CPU efficiency (CPU time vs. job duration) was also increased to 99%, compared to 65–70% prior to the code interventions. For the CMS simulation, overall memory churn was reduced by 5–8%.

5.3. Multithreaded operation

With the development of a multithread-capable implementation of GEANT4 (release 10.0) [41], some modifications to the Bertini code base were required to address issues of thread safety and memory footprint.

The multithreaded design of GEANT4 creates a separate instance of each physics process and model per thread; thus, buffers used by each class, when defined as class data members, are automatically local to each thread.

In the GEANT4 Bertini code the data tables of final state differential cross-sections (Section 2.3) had been implemented as global, static objects, but with modifiable internal buffers. By moving just the modifiable buffers to be class data members (and hence thread-local), thread collisions were eliminated without sacrificing the memory efficiency of having the data tables shared in global memory.

Once the software changes required for multithreaded operation were made, the Bertini benchmark job showed essentially the same memory and CPU performance as before.

6. Validation

Thin-target data are used to validate the evolution of models. A thin enough target will ensure that only one nuclear interaction

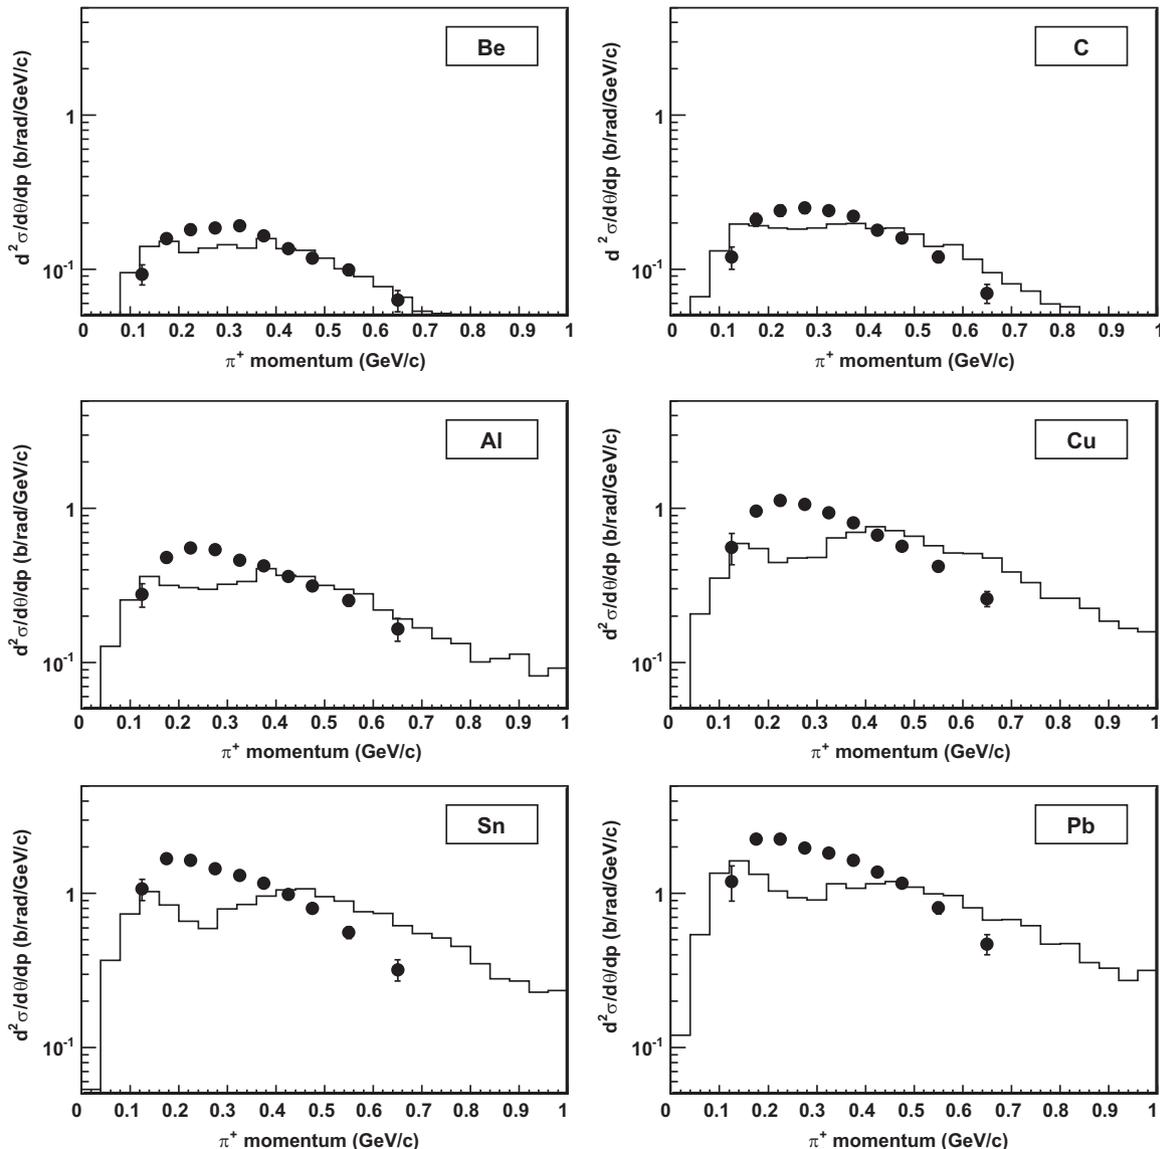


Fig. 8. Doubly differential cross-sections for π^+ produced from 8 GeV/c protons on Be, C, Al, Cu, Sn and Pb targets, for $0.35 < \theta_{\pi^+} < 2.15$ rad. Histograms represent the Bertini model predictions and dots represent the HARP [43] data.

has occurred, thereby reducing uncertainty in the angular distribution and energy dependence. Thick-target and calorimeter data are useful for testing the integrated effect of a model when it is used in conjunction with all other relevant physics models.

The GEANT4 Bertini model has been heavily validated and only a small sample of the available comparisons to data are shown here. The interested reader can examine a much larger set of comparisons at the GEANT4 hadronic validation web site [42].

6.1. Thin-target comparisons

A large amount of data was collected by the HARP collaboration [43] over most of the energy range of the extended Bertini cascade. $pA \rightarrow \pi X$ cross-sections, doubly differential in angle and

momentum, were measured from 3 to 12 GeV/c incident proton momentum and represent a consistent set of data over a range of energies, angles and target masses. Such a data set provides an excellent test of a medium energy cascade code and it is used here for comparison. Here and elsewhere in this paper, the Bertini cascade of GEANT4 version 10.1 was used.

Samples of the energy and angle dependence predicted by the Bertini cascade are shown in Fig. 7 where they are compared to the HARP data. Here the carbon data were chosen for comparison but the trends observed for this nucleus hold for the other measured targets as well.

The energy and angle dependence is reasonably well reproduced, with exceptions at lower momenta (3 GeV/c) and smaller angles; the model over-produces π^+ secondaries with momenta

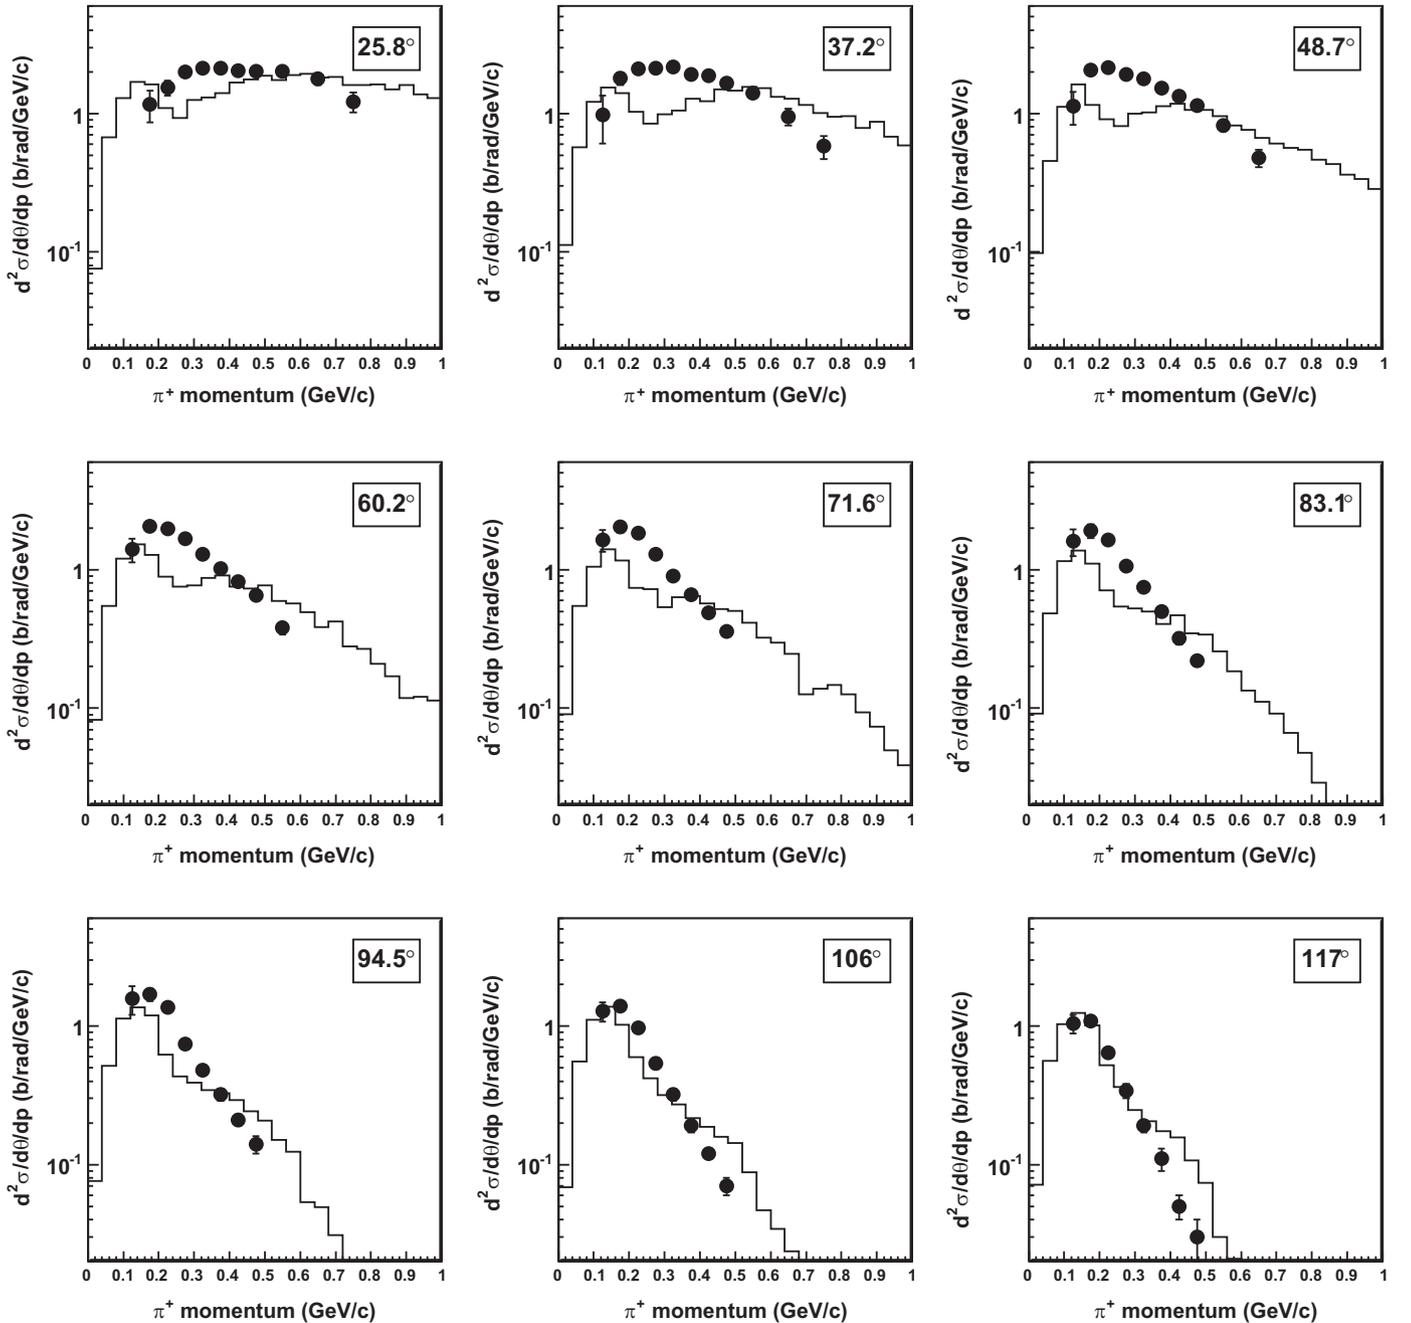


Fig. 9. Doubly differential cross-section for $p+Ta \rightarrow \pi^+ X$ at 8 GeV/c incident momentum. Histograms represent the Bertini model predictions and dots represent the HARP [43] data.

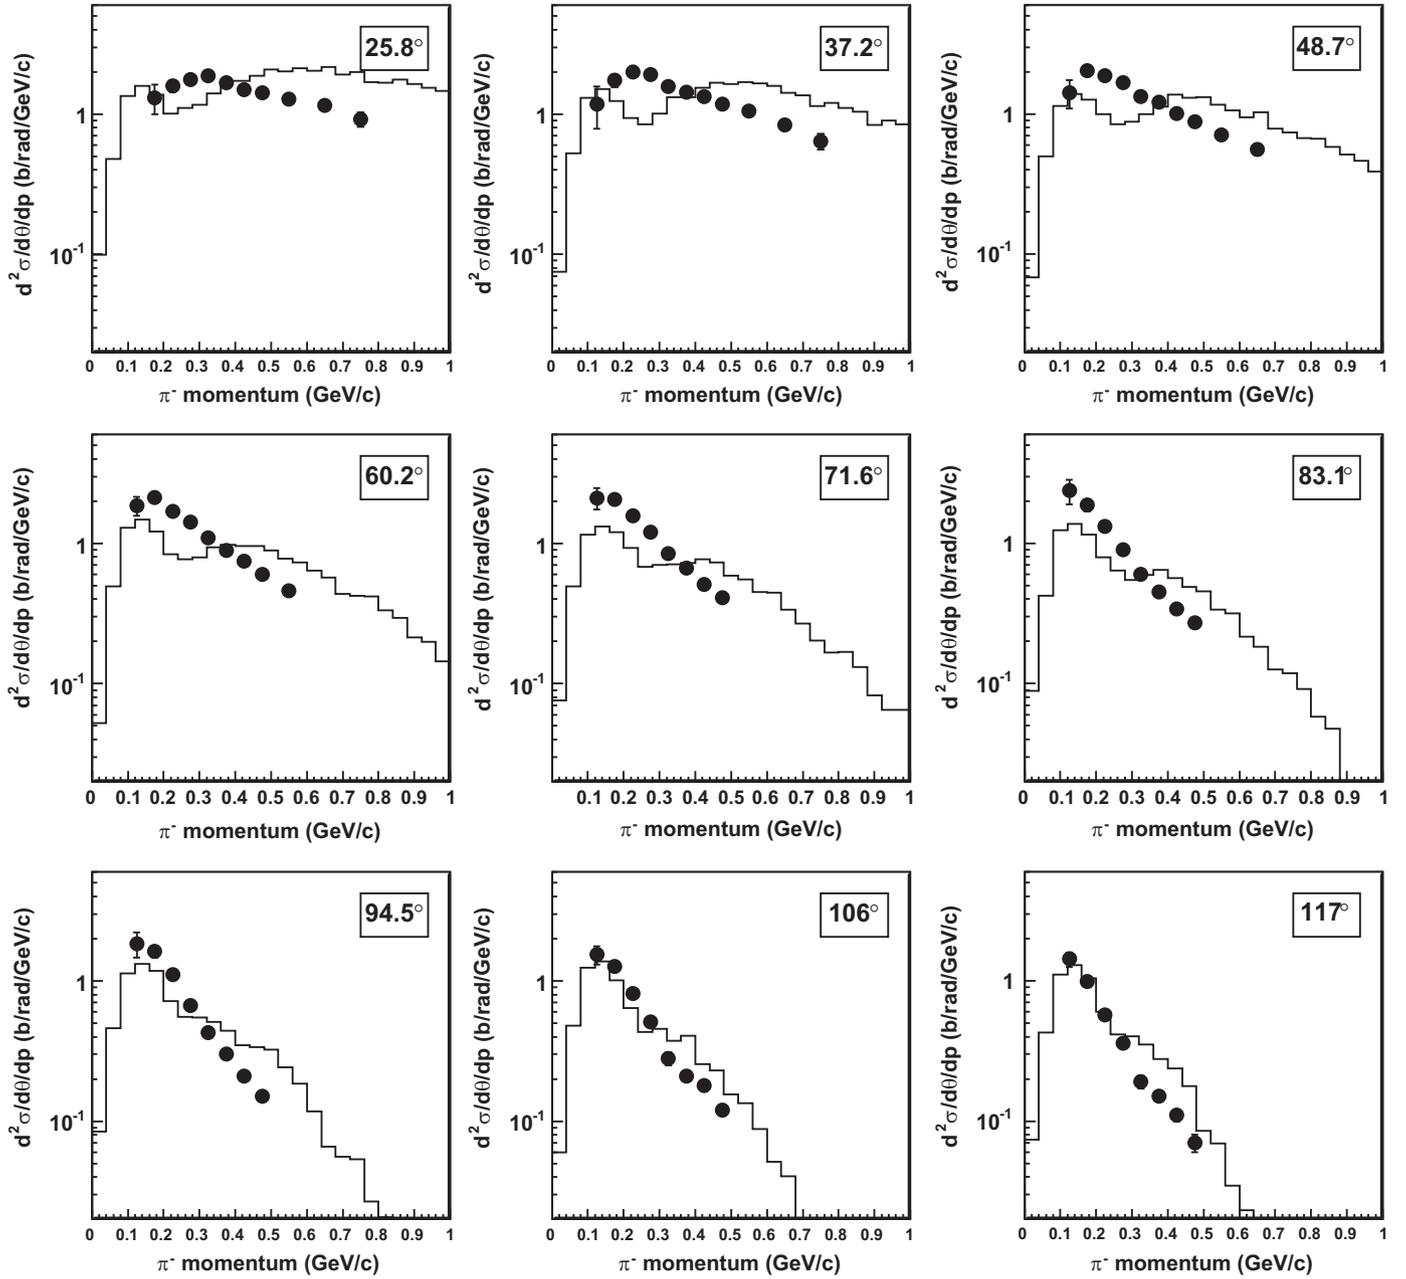


Fig. 10. Doubly differential cross-section for $p + \text{Ta} \rightarrow \pi^- X$ at 8 GeV/c incident momentum. Histograms represent the Bertini model predictions and dots represent the HARP [43] data.

above 0.4–0.5 GeV/c. The magnitude of this trend, also seen for π^- , decreases with increasing angle and incident momentum.

The A -dependence for 8 GeV/c protons on various targets is shown in Fig. 8. Here another trend is seen: for secondary momenta between 0.2 and 0.4 GeV/c the model underestimates the data for all targets. However, the normalization of the cross-section vs. A appears to be correct. A possible explanation of the deficit between 0.2 and 0.4 GeV/c is that the higher energy secondaries produced in the first interaction are not sufficiently moderated by the nuclear medium before escaping the nucleus. Increasing the amount of moderation would move higher energy particles into the deficit region and reduce the over-production seen above 0.6 GeV/c.

A comparison of π^+ vs. π^- secondaries from 8 GeV/c protons on Ta is shown in Figs. 9 and 10. No qualitative differences between the π^+ and π^- spectra are evident. The trend of the model to over-produce higher energy pions persists in both cases,

being more pronounced for π^- than for π^+ . The deficit between 0.2 and 0.3 GeV is also clear for both species at smaller angles.

6.2. Calorimeters and thick targets

LHC calorimeters have driven a large amount of development in the Bertini model. An early result from the comparison of LHC test beam data to simulation showed that showers produced by GEANT4 were too narrow. At that time all incident nucleons and pions were handled by one of the quark–gluon string models, which immediately passed medium and low energy particles to a precompound model. An intranuclear cascade stage was added when it became clear that the precompound model could not produce the hadron multiplicities expected at medium energies. The Bertini cascade was chosen because it was known from thin-target validations to produce more medium and low energy hadronic secondaries over a larger angular range than other

cascade models. As a result, the hadronic component of the showers broadened significantly. Fig. 11 plots the shower width vs. energy with and without the Bertini cascade.

A regular part of GEANT4 testing controls software evolution by monitoring the shower length and width in a simplified calorimeter. Although the Bertini cascade is only one of several models included, it occupies a critical energy range; thus the test provides a valuable constraint for any model tuning.

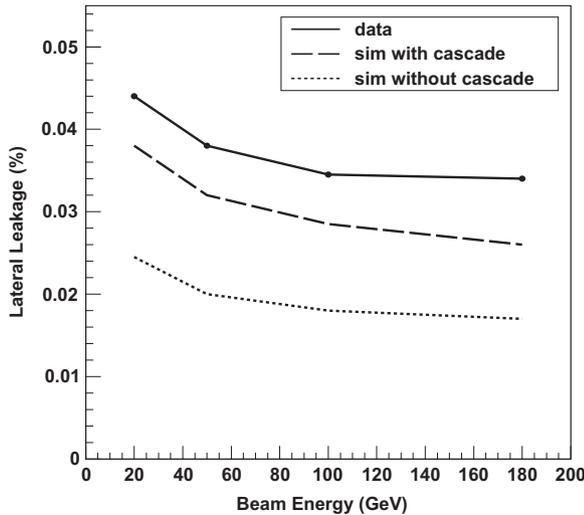


Fig. 11. Effect of adding Bertini cascade to high energy shower simulations. Lateral leakage is a measure of shower width and data are from a pion test beam on an iron/scintillator calorimeter [44]. Dashed and dotted curves are GEANT4 version 9.2 simulations with and without the Bertini cascade.

Because they effectively integrate angle and energy distributions, thick-target data have some sensitivity in the energy-angle space that thin-target data lack. A good example of this is the low energy-low angle region where some cross-sections are large and yet not well measured in thin-target experiments due to limitations in acceptance.

Some thick-target tests also allow more concentration on hadronic models, being designed to filter out electromagnetic effects as much as possible. One of these is the SATIF benchmark series. Fig. 12 shows a recent comparison made for SATIF-12 [45] which looked at the performance of a particular collection of physics processes and cross-sections, referred to in GEANT4 as a physics list. This particular physics list, named FTFFBERT, relies on the Bertini model for particle energies below 4 GeV. The SATIF comparison simulated protons incident upon a long Hg target, producing hadronic secondaries which traveled through absorbers of various thicknesses and were collected in activation counters. At the lower energy (2.83 GeV) the Bertini cascade played a central role in hadron production. At the higher energy (24 GeV) the FTFF model was more important but Bertini was still used to propagate low energy secondaries, along with the GEANT4 high precision neutron model. The data in these two regimes therefore test model performance in different energy-angle spaces.

7. Conclusion and future work

The GEANT4 Bertini cascade has proven to be a successful and very extensible model for simulating hadron, photon and lepton interactions with nuclei up to incident momenta of 12 GeV/c. Thin-target validations demonstrate good behavior over this range. Thick-target validations are scanty, but performance there is also good.

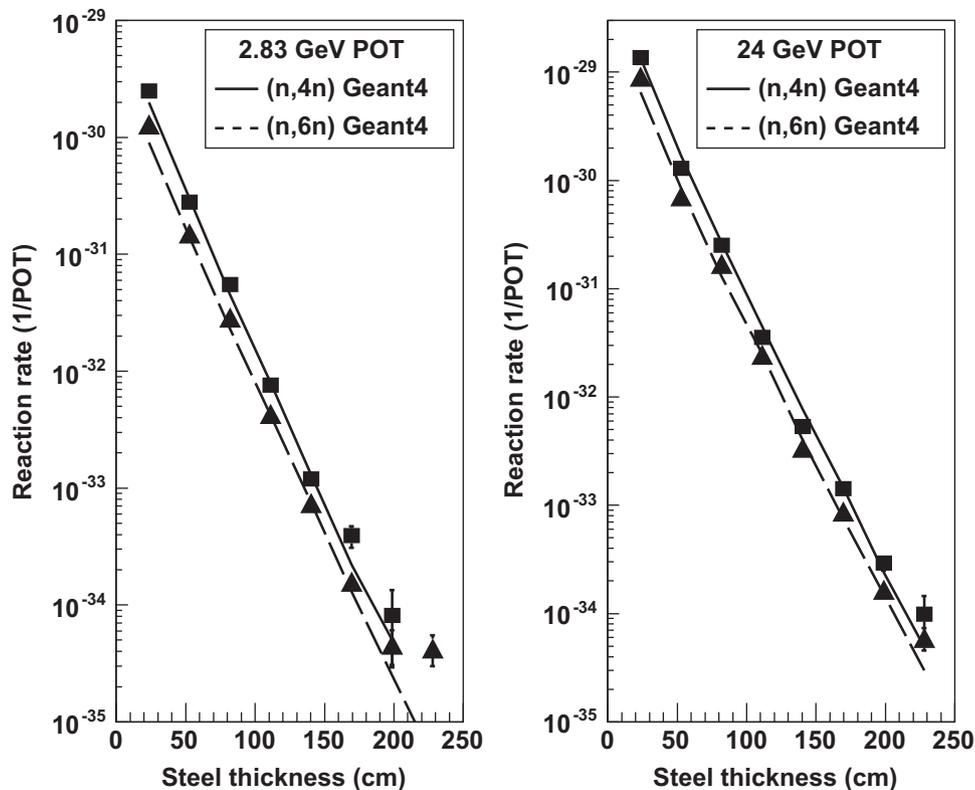


Fig. 12. Rate of ^{209}Bi ($n,4n$) ^{206}Bi (squares) and ^{209}Bi ($n,6n$) ^{204}Bi (triangles) reactions following 2.83 GeV (left) and 24 GeV (right) protons on a Hg target and subsequent penetration of hadrons through various thicknesses of steel. Data, normalized to number of protons on target (POT), are from SATIF-9 [46]. The corresponding curves were generated using the FTFF_BERT physics list from GEANT4 version 10.0.

The CPU efficiency of the code was increased from 70% to 99% as a result of profiling analysis and code re-use. The migration to multithreading also increased efficiency by taking advantage of modern multi-core computer architectures. Memory churn was significantly reduced between 8% and 92%, depending on the application.

Improvements in the model will continue to be made. In the near future a multi-parameter tuning will be performed based on a large set of thin-target data and some thick-target data. The SATIF thick-target data mentioned above will be used, as well as that by Leray [47]. In addition to the thin-target data sets mentioned above, the sets of Refs. [48–51] will be used, among others. These data are almost all produced by proton or pion beams, and look at proton, neutron and pion production.

It is also likely that the GEANT4 Bertini cascade will continue to be extended as the need arises, including improvements to the strange particle interactions, addition of anti-proton and anti-neutron induced reactions, and the addition of charmed meson induced reactions.

Acknowledgments

This paper was written on behalf of the GEANT4 hadronic working group, many of whom made valuable contributions, and the GEANT4 collaboration. We gratefully acknowledge the support and hospitality of the CERN Physics Software group. This work was supported in part by DOE contract numbers DE-AC02-76CH00016 and DE-AC03-76SF00515, and DOE Office of Nuclear Physics award number DE-FG02-099ER41110.

References

- [1] S. Agostinelli, GEANT4 Collaboration, *Nuclear Instruments and Methods in Physics Research Section A* 506 (2003) 250.
- [2] J. Allison, GEANT4 Collaboration, *IEEE Transactions on Nuclear Science NS-53* (2006) 270.
- [3] E. Yu, Titarenko, et al., Experimental and Computer Simulations Study of Radionuclide Production in Heavy Materials Irradiated by Intermediate Energy Protons, 1999. [arXiv:nucl-ex/9908012](https://arxiv.org/abs/nucl-ex/9908012).
- [4] N.V. Stepanov, ITEP preprint ITEP-55, Moscow, 1988.
- [5] A. Heikkinen, N. Stepanov, J.P. Wellisch, Bertini intranuclear cascade implementation in GEANT4, in: *Computing in High Energy Physics 2003*, La Jolla, California, USA, March 24–28, 2003.
- [6] H.W. Bertini, *Physical Review* 131 (1963) 1801; H.W. Bertini, Monte Carlo Calculations on Intranuclear Cascades, Oak Ridge National Laboratory Report ORNL-3383, 1963.
- [7] S. Mashnik, A. Sierk, CEM03.03 User Manual, Los Alamos National Laboratory, Report LA-UR-12-01364, 2012.
- [8] A. Ferrari, P.R. Sala, A. Fasso, J. Ranft, FLUKA: A Multi-Particle Transport Code, CERN-2005-10, INFN/TC_05/11, SLAC-R-773, 2005.
- [9] A. Boudard, J. Cugnon, J.-C. David, S. Leray, D. Mancusi, *Physical Review C* 87 (2013) 014606.
- [10] G. Folger, V.N. Ivanchenko, J.P. Wellisch, *European Physical Journal A* 21 (2004) 407.
- [11] J.J. Griffin, *Physical Review Letters* 17 (1966) 478; J.J. Griffin, *Physics Letters* 24B (1967) 5.
- [12] P. Fong, *Statistical Theory of Fission*, Gordon and Breach, New York, 1969.
- [13] V. Weisskopf, *Physical Review* 52 (1937) 295.
- [14] I. Dostrovsky, Z. Fraenkel, G. Friedlander, *Physical Review* 116 (1959) 683; I. Dostrovsky, Z. Fraenkel, P. Rabinowitz, *Physical Review* 118 (1960) 791.
- [15] V. Flaminio, W.G. Moorhead, D.R.O. Morrison, N. Rivoire, *Compilation of Cross Sections I: π^+ and π^- Induced Reactions*, CERN-HERA 83-01, 1983.
- [16] V. Flaminio, W.G. Moorhead, D.R.O. Morrison, N. Rivoire, *Compilation of Cross Sections II: K^+ and K^- Induced Reactions*, CERN-HERA 83-02, 1983.
- [17] V. Flaminio, W.G. Moorhead, D.R.O. Morrison, N. Rivoire, *Compilation of Cross Sections III: p and \bar{p} Induced Reactions*, CERN-HERA 84-01, 1984.
- [18] S.I. Alekhin, et al., *Compilation of Cross Sections IV: γ , ν , Λ , Σ , Ξ and K_L^0 Induced Reactions*, CERN-HERA 87-01, 1987.
- [19] A.S. Iljinov, et al., *Nuclear Physics A* 616 (1997) 575.
- [20] W.J. Briscoe, D. Scott, I.I. Strakovsky, R.L. Workman, Institute of Nuclear Studies of the George Washington University database (unpublished) (<http://gwudac.phys.gwu.edu/>).
- [21] D.E. Knapp, K.F. Kinsey, *Physical Review* 131 (1963) 1822.
- [22] L.K. Goodwin, R.W. Kenney, V. Perez-Mendez, *Physical Review* 122 (1961) 655.
- [23] P.M. Ogden, D.E. Hagge, J.A. Helland, *Physical Review* 137 (1965) B1115.
- [24] J.A. Helland, et al., *Physical Review* 134 (1964) B1079.
- [25] K.W. Lai, L.W. Jones, M.L. Perl, *Physical Review Letters* 7 (1961) 125.
- [26] M.L. Perl, L.W. Jones, C.C. Ting, *Physical Review* 132 (1963) 1252.
- [27] K.A. Jenkins, et al., *Physical Review D* 21 (1980) 2445.
- [28] A.S. Iljinov, M.V. Kazarnovsky, E.Ya. Paryev, *Intermediate-Energy Nuclear Physics*, CRC Press, Boca Raton, Florida, 1994.
- [29] R. Schumacher, et al., *Physical Review C* 25 (1982) 2269.
- [30] R.M. Sundelin, R.M. Edelstein, *Physical Review C* 7 (1973) 1037.
- [31] V.D. Toneev, K.K. Gudima, *Nuclear Physics* 173c (1983) A400.
- [32] J.M. Quesada, et al., *Progress in Nuclear Science and Technology* 2 (2011) 936.
- [33] G. Folger, J.P. Wellisch, *String Parton Models in GEANT4*, *Computing in High Energy Physics 2003*, La Jolla, California, USA, March 24–28, 2003, [arXiv:nucl-th/0306007](https://arxiv.org/abs/nucl-th/0306007).
- [34] B. Andersson, G. Gustafson, B. Nilsson-Almqvist, *Nuclear Physics B* 281 (1987) 289; B. Nilsson-Almqvist, E. Stenlund, *Computer Physics Communications* 43 (1987) 387.
- [35] V.V. Uzhinsky, The Fritiof model in GEANT4, in: *Proceedings of the International Conference on Calorimetry for the High Energy Frontier (CHEF 2013)*, Paris, April 22–25, 2013.
- [36] M.V. Kosssov, H.P. Wellisch, Section 45.3 of the GEANT4 Physics Reference Manual, 2002 (<http://geant4.web.cern.ch/geant4/UserDocumentation/UsersGuides/PhysicsReferenceManual/fo/PhysicsReferenceManual.pdf>).
- [37] K.F. von Weizsacker, *Zeitschrift für Physik* 88 (1934) 612; E.J. Williams, *Physical Review* 45 (1934) 729.
- [38] V.V. Borog, A.A. Petrukhin, in: *Proceedings of the 14th International Conference on Cosmic Rays*, Munich, vol. 6, 1975, p. 1949.
- [39] CMS Collaboration, *Journal of Instrumentation* 3 (2008) S08004.
- [40] G. Eulisse, L. Tuura, IgProf, The Ignominious Profiler, 2013 (<http://igprof.org/>) (retrieved 2.09.2014).
- [41] S. Ahn, J. Apostolakis, et al., GEANT4: bringing multi-threading into GEANT4 production, in: *Joint International Conference on Supercomputing in Nuclear Applications and Monte Carlo*, Paris, France, October 27–31, 2013 (2014). (<http://g4validationfnal.gov8080/G4WebAppNG/Display/Testxhtml>).
- [42] M.G. Catanesi, et al., *Physical Review C* 77 (2008) 055207.
- [44] P. Adragna, et al., *Nuclear Instruments and Methods in Physics Research Section A* 615 (2010) 158.
- [45] T. Koi, Shielding benchmarks for GEANT4 version 10, in: *Proceedings of the 12th Workshop on Shielding Aspects of Accelerators, Targets and Irradiation Facilities (SATIF-12)*, Fermi National Accelerator Laboratory, March 2014, in preparation.
- [46] E. Sartori, in: *Proceedings of the 9th Workshop on Shielding Aspects of Accelerators, Targets and Irradiation Facilities (SATIF-9)*, Oak Ridge National Laboratory, April 21–23, 2008, (OECD Documents, Paris, 2009).
- [47] S. Leray, *Physical Review C* 65 (2002) 044621.
- [48] T. Shibata, *Nuclear Physics A* 408 (1983) 525.
- [49] Yu.D. Bayukov, et al., *Soviet Journal of Nuclear Physics* 42 (1985) 116.
- [50] D.R.F. Cochran, et al., *Physical Review D* 6 (1972) 2085.
- [51] H. En'yo, *Physics Letters* 159B (1985) 1.