

Introduction to Monte-Carlo Methods

Francesco Cerutti



RADMEP Workshop 2023

Outline

- The problem and the solution
- Probability and statistics primer
- Uncertainty
- Continuous and discrete processes
- Scoring
- Biasing

*With precious input from many colleagues
Special thanks to F. Salvat Pujol!*

The radiation transport problem



Photons
Leptons (e^\pm , μ^\pm , τ^\pm , ν)
Hadrons (n , p , π , Σ , ...)
Ions (Z, A)
Radioactive sources
Cosmic rays

Colliding particle beams
Synchrotron radiation
...

“Monoenergetic”/Spectral

Energies:
- keV-PeV
- down to thermal energies for neutrons

Arbitrary geometry

Various shapes
materials, compounds

Radiation-matter interaction

Secondary particles
Particle shower


Material activation
Magnetic and electric fields...

Measure/estimate/score:

- Energy-angle particle spectra
- Deposited energy
- Material damage
- Biological effects
- Radioactive inventories...

Terminology

- **Radiation field**: an ensemble of particles, possibly of different species ($\gamma, e^\pm, p, n, \dots$), each at a position \mathbf{r} moving with energy E along a direction $\Omega=(\theta, \phi)$ with polar and azimuthal angles θ and ϕ
- Every particle species can undergo a series of interaction mechanisms, each characterised by a **differential cross section**:
$$\frac{d^2\sigma}{d\Omega dW} \equiv \frac{\dot{N}_{\text{count}}}{|\mathbf{J}_{\text{inc}}| d\Omega dW}$$
- The **integrated cross section** σ (area) measures the likelihood of the interaction
- Consider a medium with N_0 scattering centers per unit volume
- $N_0\sigma$ gives the probability of interaction per unit path length, AKA **macroscopic cross section**
- $1/(N_0\sigma)$ gives the **mean free path** or **scattering length** between consecutive interactions



The transport equation

- Let $n_0(\mathbf{r}, E, \Omega, t=0)$ be the particle density at the radiation source with energy E , moving in a direction Ω .
- The **transport equation** determines the radiation field (consisting of several particle species i , with different energies E , and different directions Ω) at another position \mathbf{r} at a later time t by looking at the particle balance in a small volume V (with surface S)


$$\begin{aligned}
 \int_V d\mathbf{r} \frac{\partial n_i(\mathbf{r}, E, \Omega, t)}{\partial t} &= - \oint_S dA \mathbf{j}(\mathbf{r}, E, \Omega, t) \cdot \hat{\mathbf{a}} && \text{(unscattered particles)} \\
 &- N \int_V d\mathbf{r} n_i(\mathbf{r}, E, \Omega, t) v(E) \sigma(E) && \text{(particles scattered out)} \\
 &+ N \int_V d\mathbf{r} \int dE' \int d\Omega' n_i(\mathbf{r}, E', \Omega', t) v(E') \frac{d\sigma}{d\Omega'' dW''} && \text{(particles scattered in)} \\
 &+ N \int_V d\mathbf{r} \int dE' \int d\Omega' \sum_j n_j(\mathbf{r}, E', \Omega', t) v(E') \frac{d\sigma_{\text{sec},i}}{d\Omega'' dW''} && \text{(production of secondaries)} \\
 &+ \int_V d\mathbf{r} Q_{\text{source}}(\mathbf{r}, E, \Omega, t) && \text{(source)}
 \end{aligned}$$

Notation: Ω'' is a direction such that scattering angles Ω' bring it to Ω .

Solution strategies

- Transport equation to be solved for an arbitrary source density $n_0(\mathbf{r}, E, \Omega, t)$, an arbitrary geometry, and realistic interaction cross sections.
- Solution strategies:
 - Analytical: only for restricted geometries and restricted interaction models
 - Spectral: exploit symmetries and expand in appropriate basis functions. Only for restricted cases
 - Numerical quadrature integration: general, but inefficient for high-dimensional integrals
 - **Monte Carlo method**: general, efficient, can treat arbitrary radiation fields and geometries
- Monte Carlo is a stochastic method, exploiting random numbers to:
 - **Simulate** an ensemble of particle histories governed by known interaction cross sections
 - Track them in **arbitrary geometries**
 - Accumulate contribution of each track to **statistical estimator** of the desired physical observables

The origins



Stan Ulam

Richard Feynman

John von Neumann


Probability and statistics primer

Random variables

- A **random variable** X describes the outcome of a process whose value we cannot predict with certainty, but nevertheless we know:
 - Its possible values.
 - How likely each value is, governed by the **probability density function (PDF)**, $p(x)$
- Properties of $p(x)$:
 - Positive defined: $p(x) \geq 0$ for all x
 - Unit-normalized: $\int dx p(x) = 1$
 - Integral gives probability: $\int_a^b dx p(x) = P(a < x < b)$
- The **expectation value** $\langle X \rangle = \int_{-\infty}^{\infty} dx xp(x)$ measures the average value of X
- The **variance** σ^2 measures the average square deviation from $\langle X \rangle$
$$\sigma^2 = \langle (X - \langle X \rangle)^2 \rangle = \dots = \langle X^2 \rangle - \langle X \rangle^2$$
- The **standard deviation** σ is the square root of the variance and is widely used as measurement of data spread


Relevant examples

Uniform distribution




Basic PDF for sampling

Exponential



Probability of survival

Gaussian




Natural distribution of large samples

Generation of homogeneously distributed random numbers

- For reasons of reproducibility, we use **pseudo-random numbers**: *uniformly* distributed numbers between 0 and 1 obtained from a deterministic algorithm (*not random!*) which pass all tests of randomness.
- Needs one/several **seed values**, X_1 , from which the sequence starts: X_2, X_3, X_4, \dots
- Different seed values yield different random number sequences

Sampling from arbitrary distributions

- **Sampling**: generation of random values according to a given distribution
- Various quantities are sampled
 - step length
 - event type
 - energy loss
 - deflection...
- Fundamental problem: we know how to sample uniformly distributed values, but how do we sample from arbitrary distributions?
- There's a whole array of **sampling techniques**:
 - Inverse sampling
 - Rejection sampling
 - ...



(Analytical) Inversion

- $f(x) = x, x \in [A, B]$

$$\int_A^x f(t) dt = \int_A^x t dt = \frac{t^2}{2} \Big|_A^x = \frac{x^2 - A^2}{2}$$

$$\int_A^B f(t) dt = \frac{B^2 - A^2}{2}$$

$$F(x) = \int_A^x f(t) dt / \int_A^B f(t) dt = \frac{x^2 - A^2}{B^2 - A^2} = \xi$$

$$x = \sqrt{A^2 + (B^2 - A^2)\xi}$$

ξ random number uniformly sampled over (0,1]

- $F(x)$ can also be computed numerically and the $x = F^{-1}(\xi)$ inversion performed by a lookup table

Rejection


$$x = \xi_1$$

if $\xi_2 > f(\xi_1)/f(x_{max})$, then resample ξ_1 and ξ_2 , otherwise keep $x = \xi_1$

Simplified Monte Carlo simulation algorithm


Loop over n_p primary events:

1. Initialize source particle position and momentum
2. If particle is in vacuum, advance it to next material boundary
(or sample step length to **decay** if unstable)
3. Determine total interaction cross section at present energy and material: σ
4. Evaluate the mean free path to the next interaction: $\lambda = 1/(N\sigma)$
5. Sample step length to next interaction from $p(s) = (1/\lambda) e^{-s/\lambda}$
6. Decide nature of interaction: $P_i = \sigma_i / \sigma, i=1,2,\dots,n$
7. Sample the final state of the selected interaction mechanism i . Add generated secondary particles to the stack if any
8. Score contribution of the track/event to the desired physical observables
9. Go to 2 unless particle energy drops below user preset threshold or particle exits the geometry



Example: 100-MeV proton beam on water



10 simulated proton trajectories in water ($E_p=100$ MeV):





Statistical uncertainties

- Results from Monte Carlo simulations are affected by statistical uncertainty
- How does it depend on the number of simulated primary particles?

Example: 100-MeV proton beam on water




Example: 100-MeV proton beam on water





Example: 100-MeV proton beam on water

- Results from MC simulations are affected by statistical uncertainty
- The larger the number of primaries, the smaller the error bars



A numerical experiment



- Imagine a source emitting particles with a flat energy distribution which deposit all their energy in a detector
- Let the detector/estimator measure the average deposited energy:



$$\langle E \rangle = \frac{1}{N} \sum_{i=1}^N E_i$$


- What can one say about the estimated $\langle E \rangle$?
 - It is a random variable
 - As such, it follows a certain distribution
 - Which one? It depends on the number of events

Distribution of $\langle E \rangle$ if source emits $N=1$ particle





$$\langle E \rangle = \frac{1}{N} \sum_{i=1}^N E_i$$

Distribution $\langle E \rangle$ for 1 event




Distribution of $\langle E \rangle$ if source emits $N=2$ particles





$$\langle E \rangle = \frac{1}{N} \sum_{i=1}^N E_i$$

Distribution $\langle E \rangle$ for 2 events




Distribution of $\langle E \rangle$ if source emits $N=3$ particles





$$\langle E \rangle = \frac{1}{N} \sum_{i=1}^N E_i$$

Distribution $\langle E \rangle$ for 3 events




Distribution of $\langle E \rangle$ if source emits $N=10$ particles



$$\langle E \rangle = \frac{1}{N} \sum_{i=1}^N E_i$$


Distribution $\langle E \rangle$ for 10 events






Distribution of $\langle E \rangle$

- For sufficiently large number of contributions, the estimate mean $\langle E \rangle$ follows a Gaussian!
- The standard deviation (\sim width) of this Gaussian is a measure of the statistical uncertainty when estimating $\langle E \rangle$.
- The standard deviation (statistical uncertainty) decreases with the number of contributions N
- We now check how the statistical uncertainty drops with N


Distribution $\langle E \rangle$ for 10 events



Standard deviation



$$\langle E \rangle = \frac{1}{N} \sum_{i=1}^N E_i$$



- For a very large number of samples the detector would yield the estimated mean $\langle E \rangle = 1/2$ with $\text{sigma} = 0$.
- Statistical uncertainty decreases with the number of contributions N as $1/\text{sqrt}(N)$.

Central Limit Theorem

- In the numerical experiment above, the “detector” plays the role of the expectation value of any physical observable estimated à la Monte Carlo.
- If the estimator receives a sufficiently large number of contributions N , the distribution of its expectation value tends to a Gaussian centered around the true expectation value, with standard deviation that goes like $1/\sqrt{N}$. That is, **the statistical uncertainty of a MC estimate reduces as $1/\sqrt{N}$ with the number of primary events.**
- This is essentially the [Central Limit Theorem](#).
- Note that:
 - When doing a Monte Carlo simulation, quoting a result without a measure of the statistical uncertainty is meaningless.
 - Quoting a result obtained with a low number of contributions is dangerous: the distribution of the mean may still be far from the Gaussian centered around the actual expectation value!

Recommendations

- Remember that the variance itself is a random variable. E.g., runs with the same number of primaries but different random seeds will give different values of the variance. The larger the number of primaries, the smaller the difference (MC results for different seeds converge to the same value).
- It is wise to examine how convergence is attained: verify that error bars drop with $1/\sqrt{N}$. Sudden/isolated spikes indicate poor sampling in some corner of phase space.
- It is often a good idea to plot 2D and 3D distributions. The human eye is a good tool for judging statistical convergence of 2D/3D estimators!

Statistical uncertainty

Statistical errors, due to sampling (in)efficiency

<u>Relative error</u>	<u>Quality of Tally</u>	<i>(from an old version of the MCNP Manual)</i>
50 to 100%	Garbage	
20 to 50%	Factor of a few	
10 to 20%	Questionable	
< 10%	Generally reliable	

- The MCNP guideline is based on experience, not on a mathematical proof. But it has been generally confirmed also working with other codes.
- Small penetrations and cracks in a geometry are very difficult to handle by MC, because the “detector” is too small and too few non-zero contributions can be sampled, even by biasing.

Assumptions, limitations, and sources of uncertainty

Assumptions and limitations

- Materials are static, homogeneous, and isotropic.
- Radiation transport is treated as a Markovian process: the fate of a particle depends only on its actual state, and not on its history.
- Material properties are not affected by previous histories.
- Particles follow trajectories and interact with individual atoms/electrons/nuclei.
 - A general order-of-magnitude measure: the particle's de Broglie wavelength must be small compared to typical interatomic distances (Angstroem).


Validity of the trajectory picture

- De Broglie wavelength:

$$\lambda_{\text{dB}} = \frac{hc}{\sqrt{E_K(E_K + 2m_0c^2)}}$$

where E_K is the particle's kinetic energy, m_0 is its rest mass, h is the Planck constant, and c is the speed of light.

- Typical interatomic distances are in the order of \sim Angstroem.
- E.g.: MC simulation of electron transport at energies much below 100 eV is questionable.
- The assumption of scattering on single target puts a lower energy bound on applicability of MC



Systematic uncertainties [1/2]

- We have discussed **statistical uncertainties** above.
- That's only part of the uncertainty in the results of any MC simulation. The rest are **systematic uncertainties**, due to:
 - Adopted physics models: different codes are based on different physics models. Some models are better than others. Some models are better in a certain energy range. Model quality is best shown by benchmarks at the microscopic level (e.g. thin targets)
 - Transport algorithm: due to imperfect algorithms, e.g., energy deposited in the middle of a step, inaccurate path length correction for multiple scattering, missing correction for cross section and dE/dx change over a step, etc. Algorithm quality is best shown by benchmarks at the macroscopic level (thick targets, complex geometries)
 - Cross-section data uncertainty: an error of 10% in the absorption cross section can lead to an error of a factor 2.8 in the effectiveness of a thick shielding wall (10 attenuation lengths). Results can never be better than allowed by available experimental data

Systematic uncertainties [2/2]

- Systematic errors due to incomplete knowledge:
 - material composition not always well known. E.g. concrete/soil composition (how much water content? Can be critical)
 - beam losses: most of the time these can only be guessed. Close interaction with engineers and designers is needed.
 - presence of additional material, not well defined (cables, supports...)
 - Is it worth to do a very detailed simulation when some parameters are unknown or badly known?
- Systematic errors due to simplification:
 - Geometries that cannot be reproduced exactly (or would require too much effort)
 - Air contains humidity and pollutants, has a density variable with pressure

Errors, bugs, mistakes


- Monte Carlo codes can contain bugs:
 - Physics bugs
 - Programming bugs (as in any other software, of course)
- User mistakes:
 - mistyping the input (precious help from the user interface)
 - error in user code
 - wrong units
 - wrong normalization: quite common
 - unfair biasing: energy/space cuts cannot be avoided, but must be done with much care

Continuous processes

Too frequent


- **EMFP**: mean free path between consecutive Coulomb scattering
- **IMFP**: mean free path between consecutive ionization losses
- **RANGE**: estimated distance traveled to rest

- Estimate number of ionization losses:
 - $N = \text{RANGE} / \text{IMFP}$
 - e.g. for a 1-MeV electron, $N \sim 10^4$
- Estimate number of Coulomb scatterings:
 - $N = \text{RANGE} / \text{EMFP}$
 - e.g. for a 1-MeV electron, $N \sim 10^4$
- **Too many to simulate explicitly!**
- A more practical approach is necessary to keep CPU time within acceptable bounds.



Ionization energy losses

Two different treatments: **small** vs **large** energy losses



- Ionization loss cross sections $\propto 1/T^2$
→ Small losses are dominant
Expensive CPU-wise to simulate them all
- Need for aggregation along particle step:
 - Determine average energy loss per unit path length up to T_δ (restricted stopping power)
 - Random fluctuations applied on top

Energy is deposited **locally along the step**

- Transferred energy sets target electron in motion
→ δ ray
- δ rays
 - are energetic
 - can transport energy away from origin
→ are explicitly produced/transported
- Cross sections depends on projectile
 - Moller (e-)
 - Bhabha (e+)
 - generic spin-0, spin-1/2...

Coulomb scattering

- Use of an algorithm based on the Moliere multiple-scattering theory
 - **screened Rutherford** differential cross section for an individual collision

$$\frac{d\sigma_{\text{mol}}}{d\Omega} = \left[\frac{z^2 Z^2 e^4}{4c^4 \beta^2 E^2 \sin^4 \frac{1}{2}\theta} \right] \left[\frac{(1 - \cos \theta)^2}{(1 - \cos \theta + \frac{1}{2}\chi_a^2)^2} \right]$$

- small-angle approximation
- via analytical manipulations → **minimum applicable step length (energy-dependent)**
- at every step t, aggregate deflection is sampled from F_{Mol}

$$F_{\text{Mol}}(\theta, t) d\Omega = 2\pi\chi d\chi \left[2e^{-\chi^2} + \frac{1}{B} f_1(\chi) + \frac{1}{B^2} f_2(\chi) + \dots \right] \left[\frac{\sin \theta}{\theta} \right]^{\frac{1}{2}}$$

$$f_n(\chi) = \frac{1}{n!} \int_0^\infty u du J_0(\chi u) e^{-u^2/4} \left(\frac{u^2}{4} \ln \frac{u^2}{4} \right)^n$$


- There are situations where the Moliere theory is not applicable:
 - Transport in residual gas
 - Interactions in thin geometries like wires or slabs (few elastic collisions)
 - Electron spectroscopies at sub-10-keV energies
 - Micro-dosimetry

→ **single scattering**

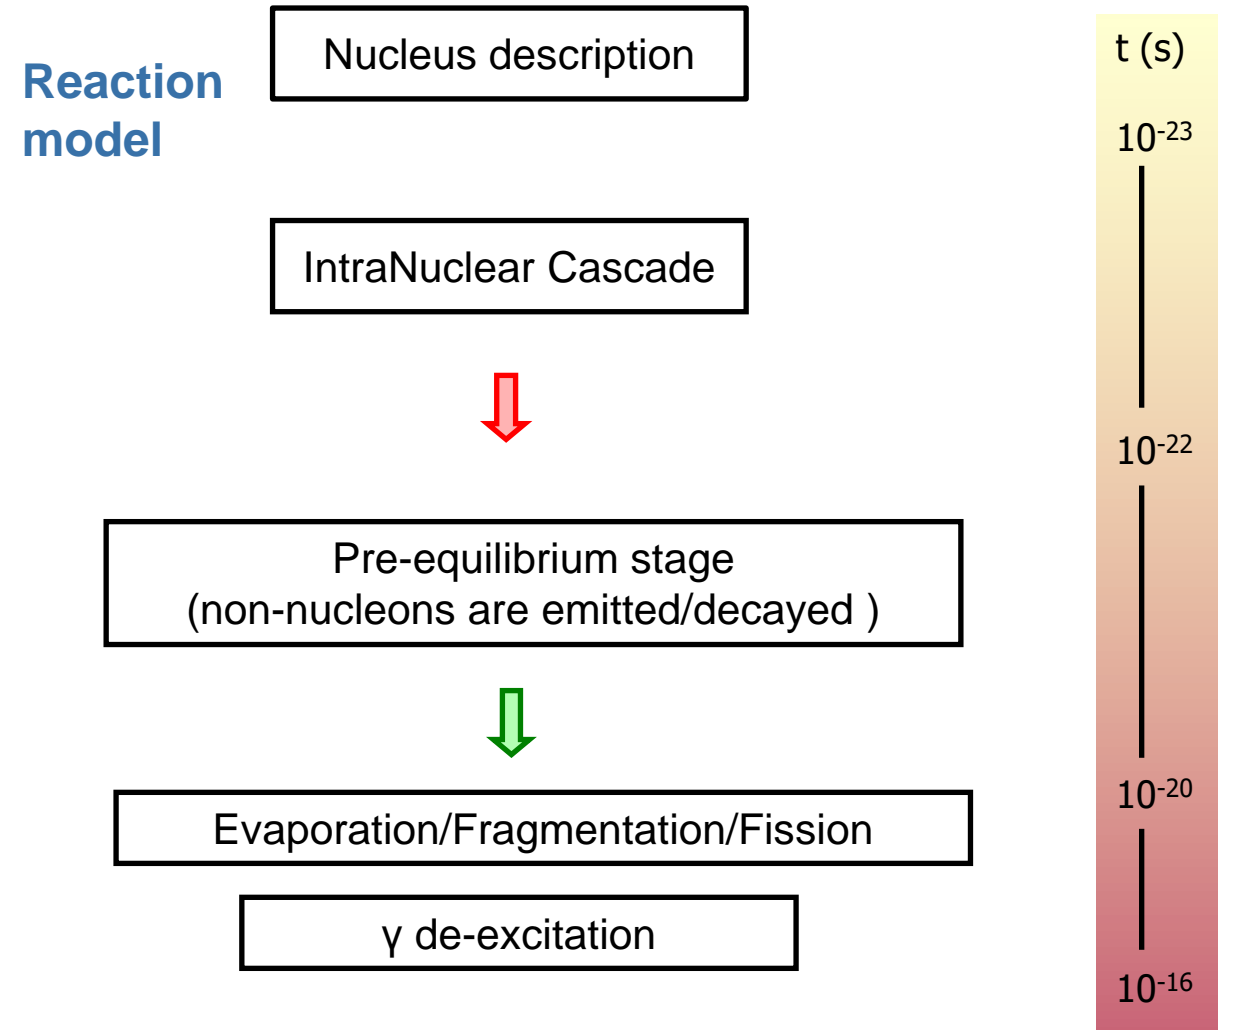
Discrete processes

Nuclear reactions

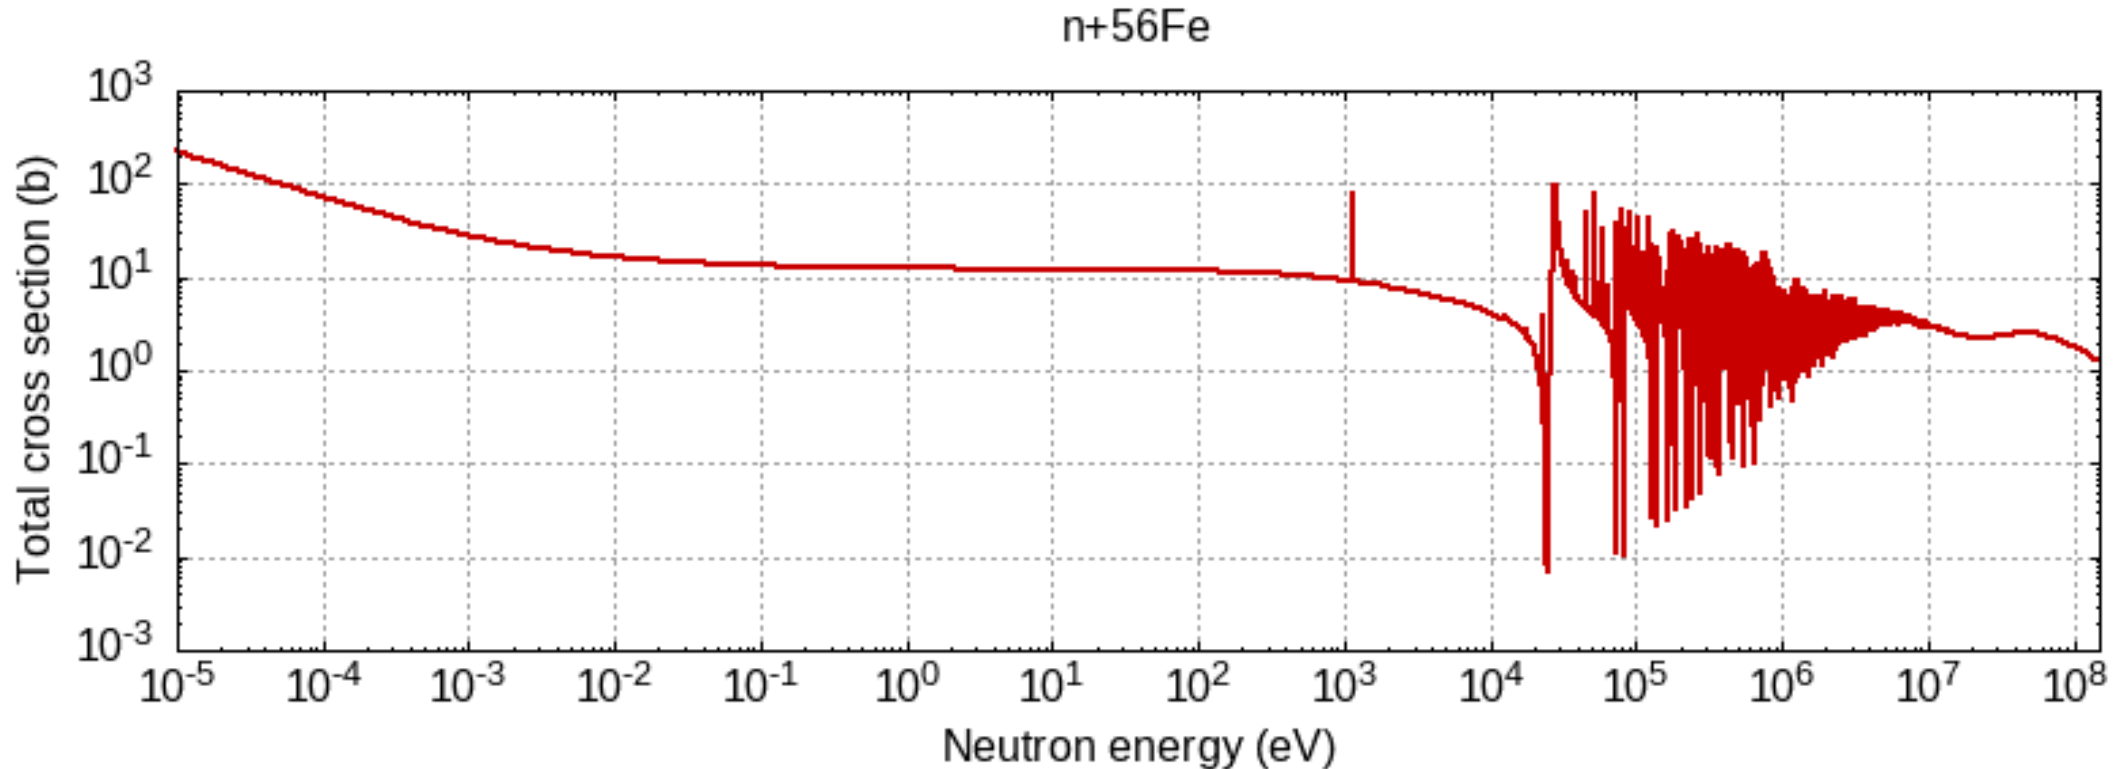
i. to decide the process occurrence



ii. to decide the reaction final state



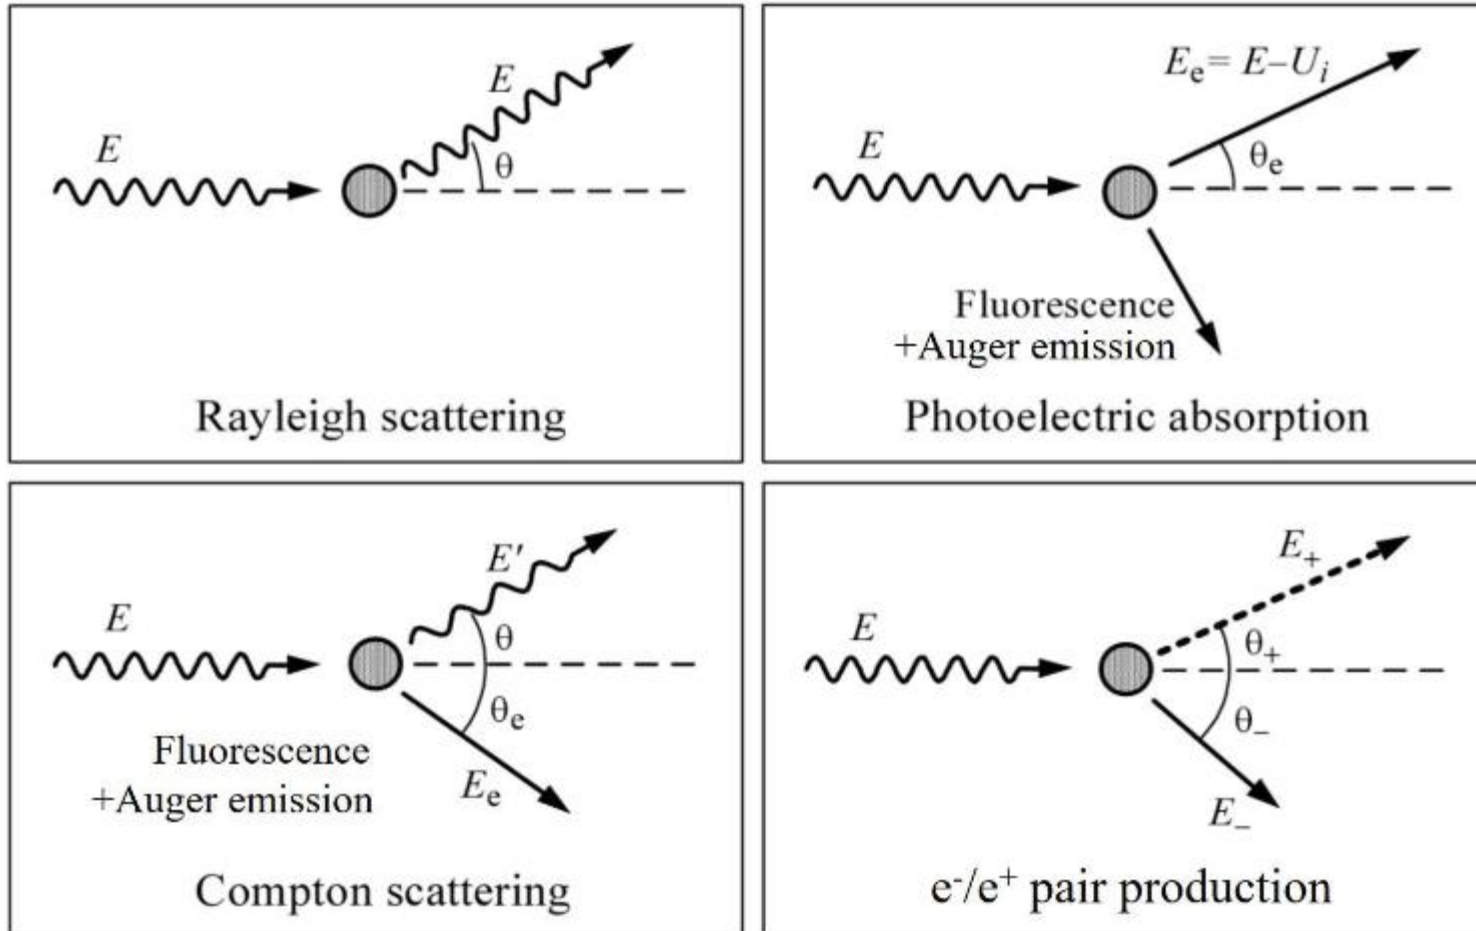
Low-energy neutrons



No general effective model is able to reproduce the complex resonance structure.

Transport codes rely on **libraries of evaluated nuclear data** to describe neutron interactions below ~20 MeV for a reasonably comprehensive list of isotopes, including *elastic scattering*, *capture*, *fission*, and *many explicit inelastic channels*.

Photon interactions





[PENELOPE manual]

+ photonuclear reactions

+ μ^+ / μ^- pair production

Electron/positron interactions



[PENELOPE manual]

+ electronuclear reactions

Thresholds

In a MC simulation, particles are tracked until either

- they leave the simulation geometry;
- their energy drops below a predefined value, the **transport threshold**.

Too high thresholds imply premature tracking end, preventing further traveling.

Too low thresholds imply waste of time with no gain in the simulation results.

Thresholds should be set such as the particle residual range is smaller than the desired resolution.

Electromagnetic particle thresholds play a major role in CPU-time consumption.

Photons travel farther than electrons of the same energy, hence require lower thresholds.

Scoring

The concept

- It is said that Monte Carlo (MC) is a “**mathematical experiment**”; the MC equivalent of the result of a real experiment (*i.e.*, of a **measurement**) is called an **estimator**
- Just as a real measurement, an estimator is obtained by sampling from a statistical distribution and has a **statistical error** (and in general also a **systematic** one)
- There are often several different techniques to measure the same physical quantity: in the same way, **the same quantity can also be calculated using different kinds of estimators**

What:

Energy deposition and associates (power, dose), fluence or current versus energy, angle or other kinematic variables, time, DPA, residual activity...

Where?

In geometry regions, across boundaries,
on independent spatial grids

When?

At the **end of a history batch** or at
each event

Definitions

- N : number of identical particles
- N_0 : number of atoms per unit volume
- λ : **mean free path**, *i.e.* average distance travelled by a particle in a material before an interaction. It depends on the material, particle type and energy
- l : **total distance travelled**
- v : **average particle velocity**

Cross section

- $\Sigma [cm^{-1}] = 1/\lambda [cm]$: **macroscopic cross-section**, *i.e.* probability of interaction per unit distance. It depends on the material, particle type and energy.
- $\sigma = \frac{\Sigma}{N_0}$ (atom effective area [barn = 10^{-24} cm²]) : **microscopic cross-section**, *i.e.*
 - the **area of an atom weighted with the probability of interaction** (hence the name “cross-section”)
 - or the **probability of interaction per unit length, with the length measured in atoms/cm²**
- The microscopic and macroscopic cross-section have a similar physical meaning of “probability of interaction per unit length”, with length measured in different units. Thus, **the number of interactions** can be obtained by multiplying them by the corresponding particle **track-length**

Reaction rate and fluence


- $R = Nl\Sigma = \Phi V\Sigma$: number of reactions in a given time interval inside the volume V (where Φ is the fluence and the product $\Sigma\Phi$ is integrated over energy or velocity)
 - $\dot{R} = N \frac{dl}{dt} \Sigma = Nv\Sigma$: **reaction rate**
 - $\frac{d\dot{R}}{dV} = \frac{dN}{dV} v \Sigma = n(\mathbf{r}, v)v\Sigma$: reaction rate inside the volume element dV
- $\Phi(\mathbf{r}, v) = n(\mathbf{r}, v)dl$ [cm^{-2}] : **fluence**, *i.e.* time integral of the flux density
 - Fluence is expressed in **particles per cm^2** but in reality represents the **density of particle tracks [cm / cm^3]** !
- $\dot{\Phi}(\mathbf{r}, v) = n(\mathbf{r}, v)v$ [$\text{cm}^{-3} \text{cm/s} = \text{cm}^{-2} \text{s}^{-1}$] : **fluence rate** or **flux density**

Fluence vs current

Surface crossing estimation

- Consider the volume generated by a surface S times an infinitesimal thickness dt .

A particle incident with *an angle θ with respect to the normal to the surface S* travels a segment $dt/\cos\theta$ inside the volume.



- The **average fluence F over the surface S** is defined as:

$$\Phi = \lim_{dt \rightarrow 0} \frac{\sum_i \frac{dt}{\cos \theta_i}}{S dt}$$

total tracklength inside the volume
volume

- While the **average current J over the surface S** is given by the number of particles crossing the surface divided by the surface area: $J = N/S$

- Fluence is **independent** of the orientation of the surface S , while current is **not** !


- On a flat surface in an isotropic particle field $J = F/2$

- Current is meaningful in case one needs to count particles (e.g. for a signal trigger)
- But to estimate dose, activation, radiation damage, instrument response... the relevant quantity to be used is fluence, since it is proportional to the interaction rate


Energy deposition profiles

From 3D spatial meshes of energy deposition density (GeV/cm^3), different 1D plots can be obtained:

- taking the average of the values for each z-bin (i.e. over all X, Y or R, Φ)

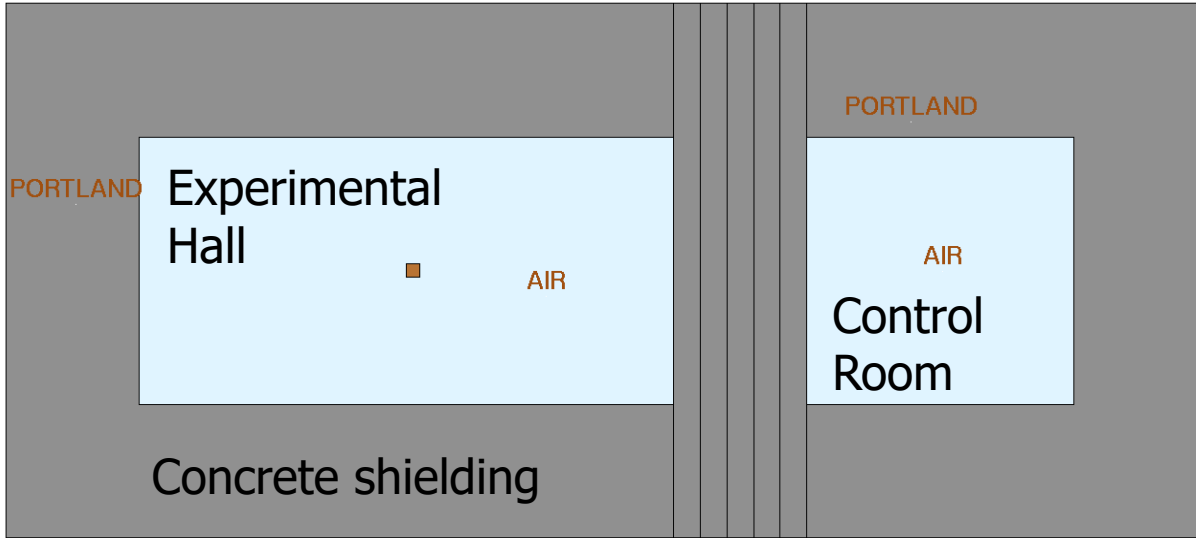


- taking the highest value for each z-bin


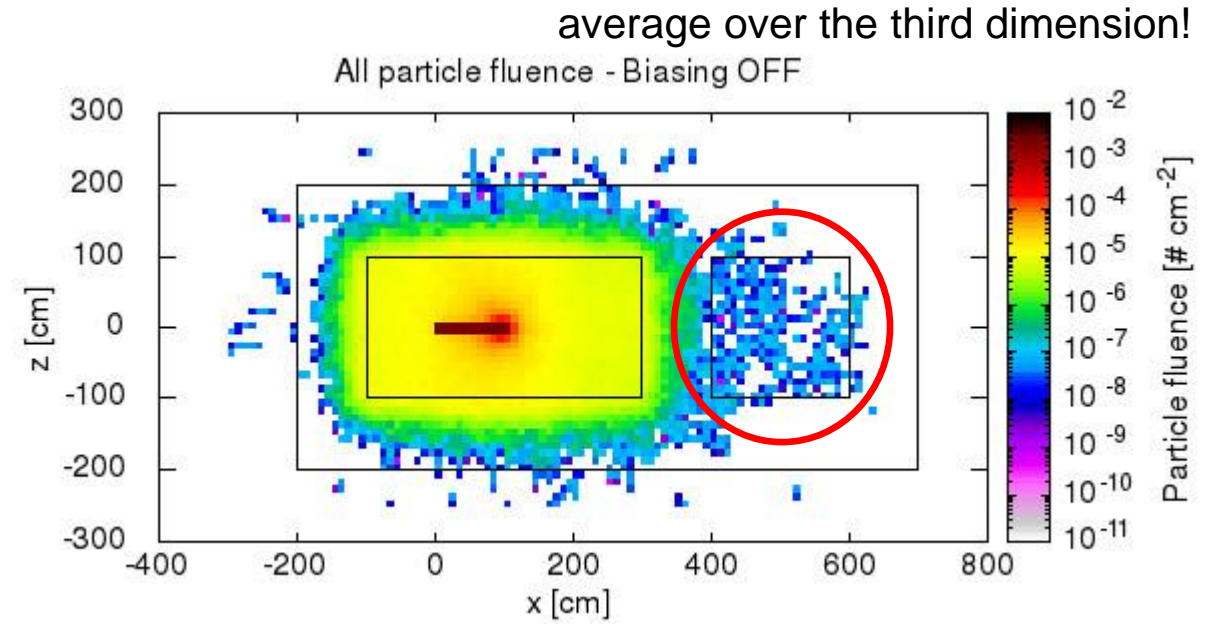


Biassing

The problem



Quantification of radiation levels in the CR badly converges.



The concept

Deliberate alteration of simulation parameters to reduce the product between **variance and CPU time** (figure of merit).

To preserve a sound mathematical treatment, this bias is compensated for by changing **the statistical weight of simulated particles**. Nonetheless, **physical correlations are no longer preserved**.

Several techniques are employed (not exhaustive list):

- *Region Importance Biasing*
- *Mean Free Path Biasing*
- *Leading Particle Biasing*
- *Multiplicity Tuning*
- *Decay-length Biasing*
- *Weight Windows*
- *Neutron non-analogue absorption*
- *Direction biasing*


They require **user's** time to be implemented as well as active reasoning and experience!

Splitting

- Moving toward a higher importance (I) region
- $n = I_2/I_1$ particle *replicas* are created
- *Weight* of replicas is $w = 1/n = I_1/I_2 < 1$
- Total weight of all replicas is equal to the weight of the original particles

Reduces variance


WARNING: The variance reduction achieved through particle replicas is **authentic** only if they experience independent histories, which may **not** be the case when they travel in **low-density regions** (air, vacuum!).



Russian Roulette

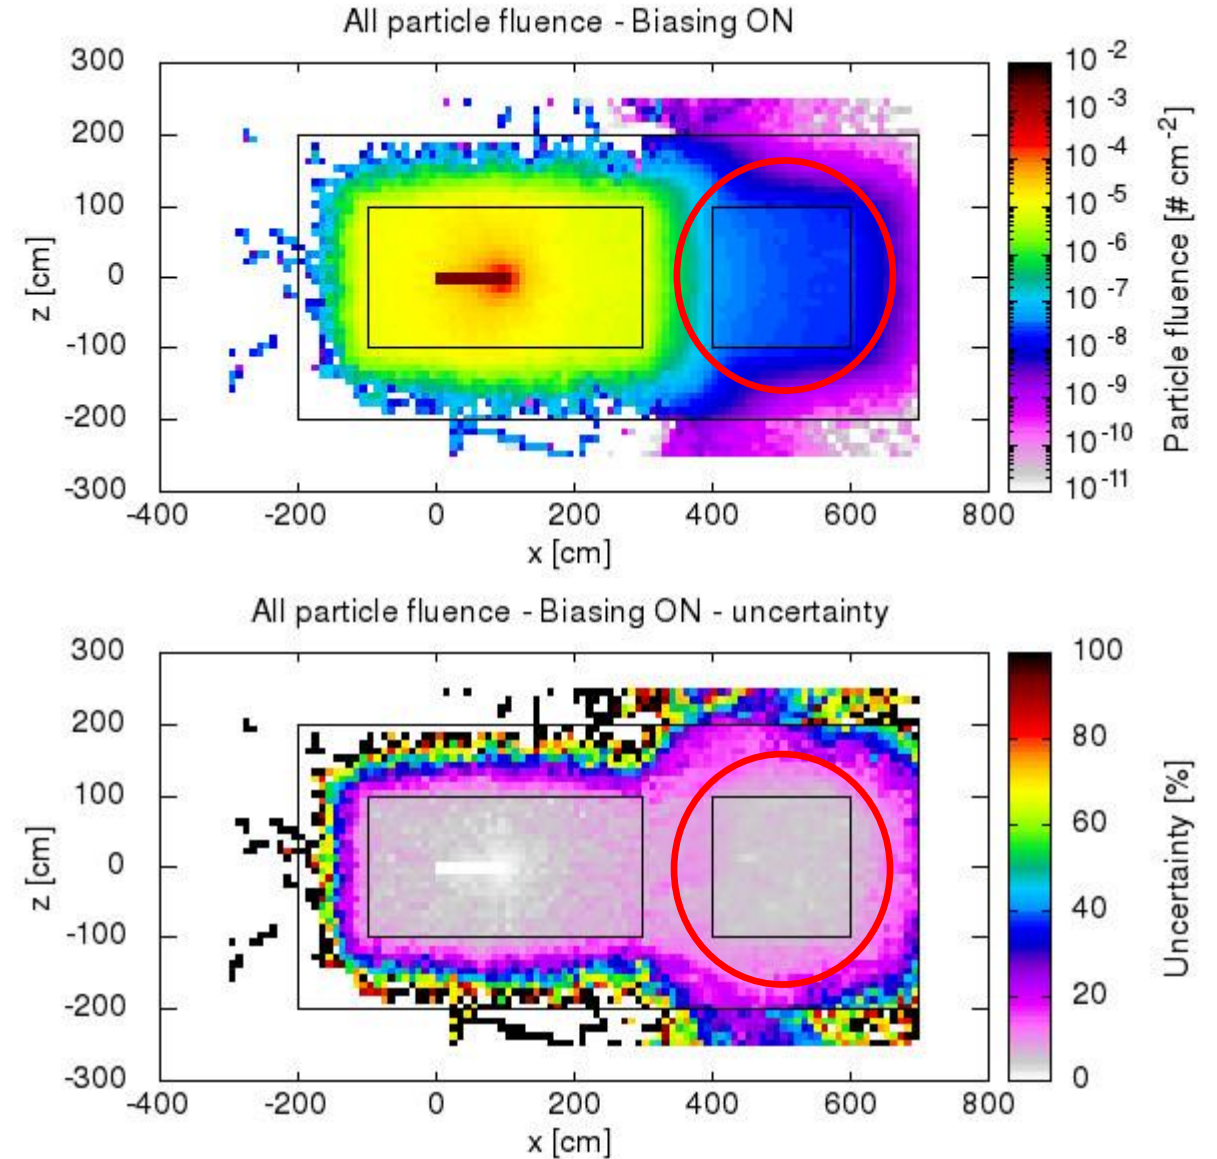
- Moving toward a lower importance (I) region
- Particles get a survival probability $P_s = I_2/I_1$
- *Weight* of surviving particles is increased to $w = 1/P_s = I_1/I_2 > 1$
- Weight of all surviving particles is equal to the weight of the original particles

Reduces CPU time



As a result

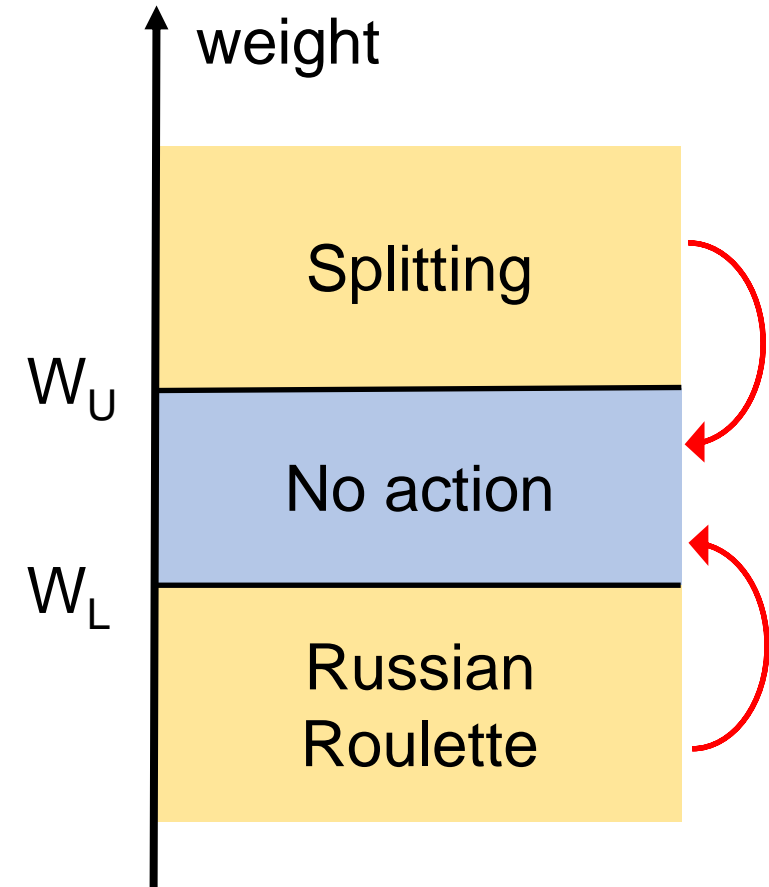
Good convergence
(for the same CPU time)




Weight windows

A particle of weight w is

- split on average into w/W_U replicas if $w > W_U$
- is subject to Russian Roulette if $w < W_L$





Background at the (very) forward physics LHC experiments



[M. Sabaté-Gilarte, FPF experiments meeting, May 2022]

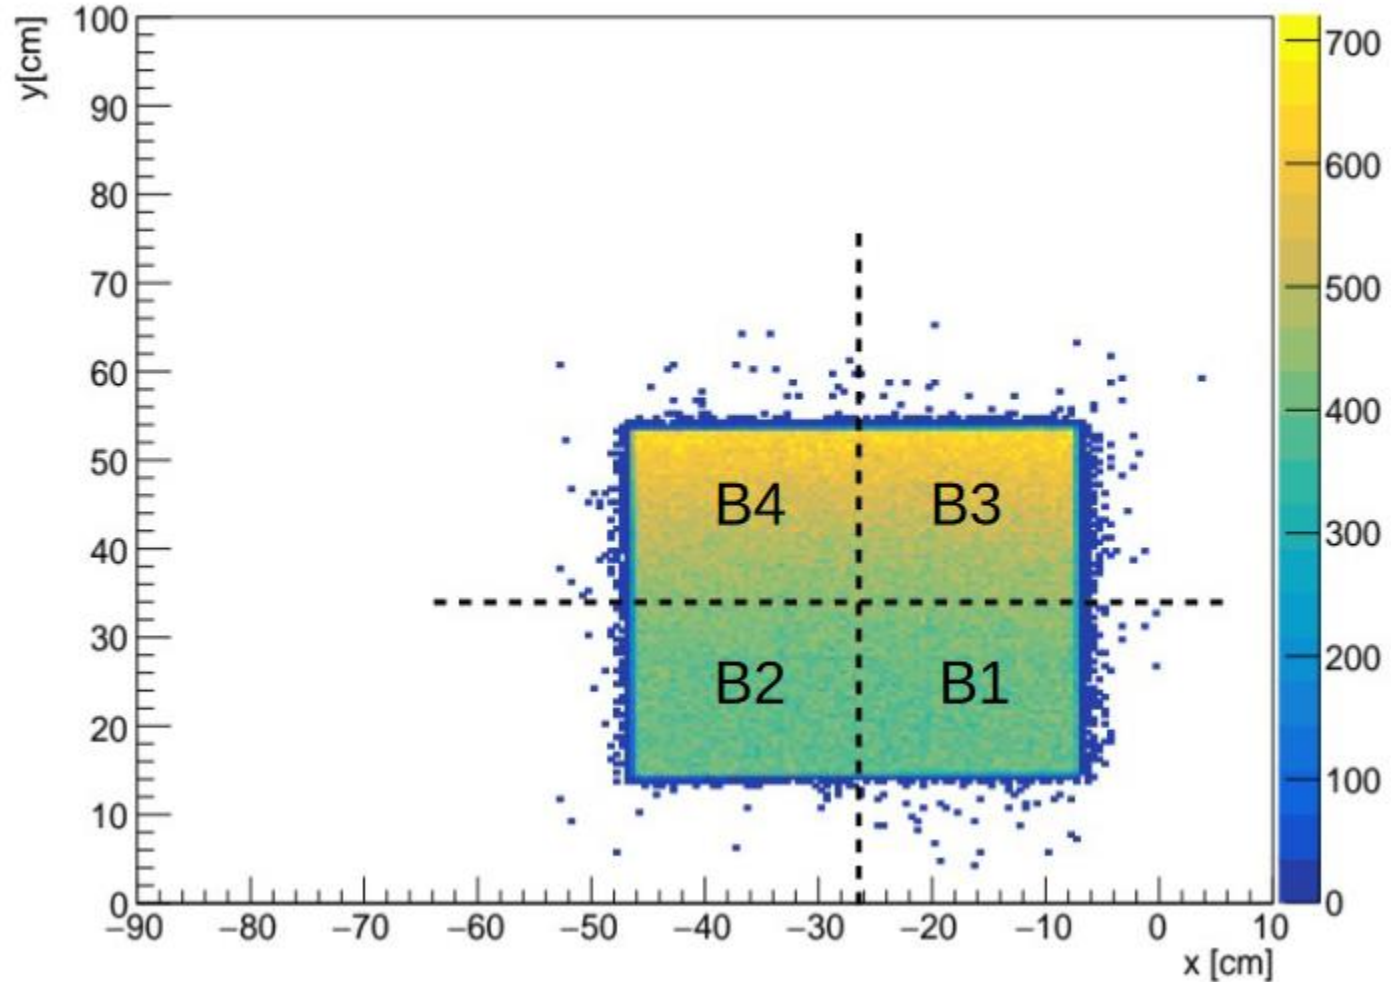
Muons at SND@LHC [1/2]



[M. Sabat -Gilarte, 12th SND@LHC
Collaboration meeting, Feb 2023]

Muons at SND@LHC [2/2]

SciFi tracks at Wall 2

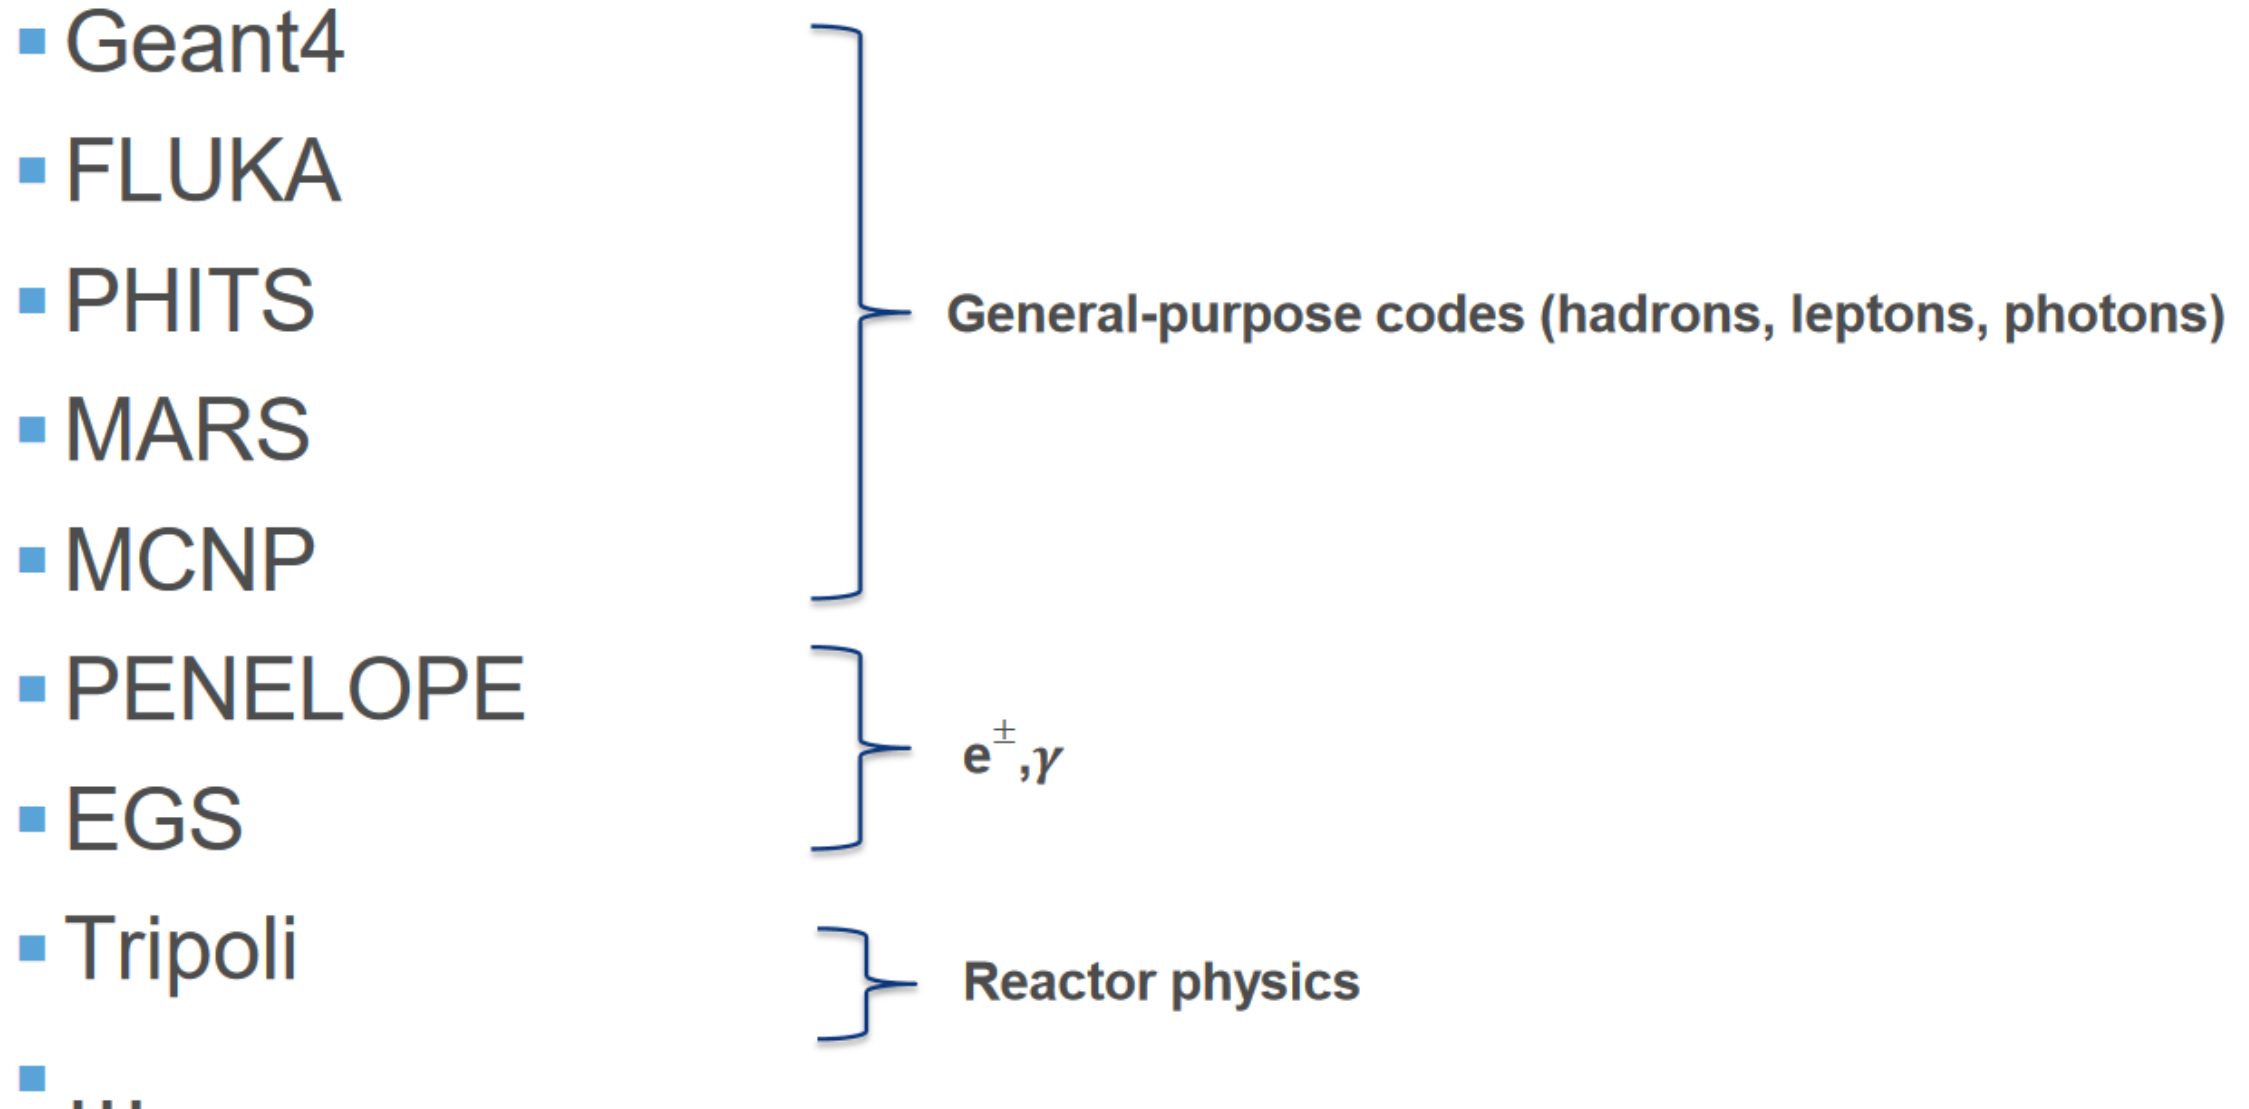


MC true muon flux in SciFi, all area:
 1.7×10^4 tracks fb/cm²

**Data and MC expectation agree
within 10%!**

[S. Ilieva, 12th SND@LHC
Collaboration meeting, Feb 2023]

MC codes for simulation of radiation transport



BACKUP

The mean free path and its distribution

- Let n particles per unit time and surface impinge normally on a thin material slab of width ds with a density of N scattering centers per unit volume, each having a cross sectional area σ .
- Number of particles that interacted: $dn = n N \sigma ds$.
- The interaction probability in ds : $dn/n = N \sigma ds$
- Let $p(s)$ be the distribution of path lengths to the next interaction.
- The probability that the next interaction is within ds of s is

$$p(s) = [1 - \int_0^s ds' p(s')] (N \sigma) = \int_s^{inf} p(s') (N \sigma) ds'$$
- The solution of this diff eq is

$$p(s) = (N \sigma) e^{-s(N \sigma)}$$
- Thus, the path length to the next interaction follows an exponential distribution. The average distance to the next interaction is:

$$\langle s \rangle = 1/(N \sigma) = \lambda,$$
 i.e., we recover the expression of the mean free path given above.

