PATMOS

- Monte Carlo method
- MPI/OpenMP parallelisation
- **80% CPU time** in a single hot spot (SIGMA1 kernel)
- Strong interest in **hybrid** architectures

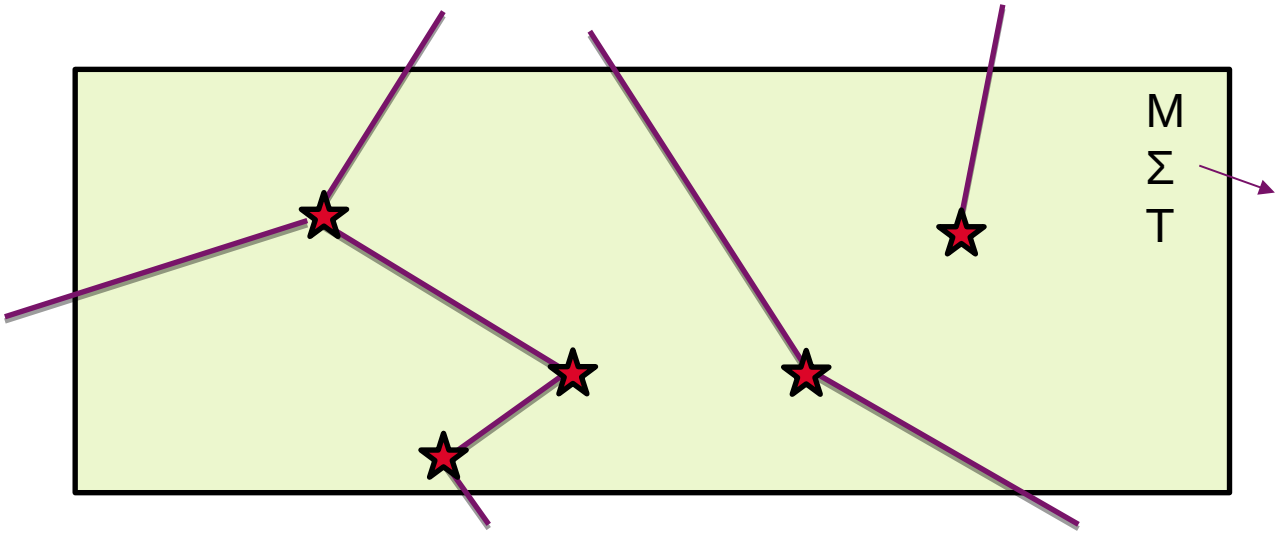

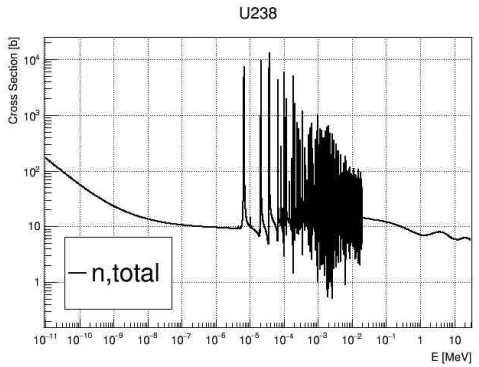

REX

- past studies for the optimisation of precision for SIGMA1
  - E. Brun et al., "A Study of the Effects and Benefits of Custom-Precision Mathematical Libraries for HPC Codes," in IEEE Transactions on Emerging Topics in Computing 9 (2021) 1467-1478
- parallel debugging is complicated
  - need for reproducibility
  - ideally: mono=para
- precision for the whole MC run: I am not even sure how to pose the problem, but it would be an interesting question

MONTE CARLO TRANSPORT ALGORITHM (2)

```
for each batch
  initialize particle state from source
  for each particle in batch
    while (particle alive)
      find material at particle position
      for each nuclide in material
        load total micro cross-section & atomic density
        accumulate macro cross section
      sample next interaction point
      sample collided nuclide
      for each reaction type in collided nuclide
        load partial micro cross-section
      sample reaction
      update particle state (direction, energy, alive…)
    end
  end
end
```

Two approaches:

- Binary search on pre-tabulated XS
- On-the-fly Doppler broadening of XS

MONTE CARLO TRANSPORT ALGORITHM (3)

Binary search on pre-tabulated XS

Neutron cross sections σ(E)

- They are stored as **long tables** of unequally spaced values (E_i, σ_i) with linear reconstruction for $E_{i-1}<E<E_i$
  - ^{1}H: $6x10^{2}$ grid points
  - ^{238}U: $1.5x10^{5}$ grid points
  - "unified" grid: $3x10^{6}$ grid points

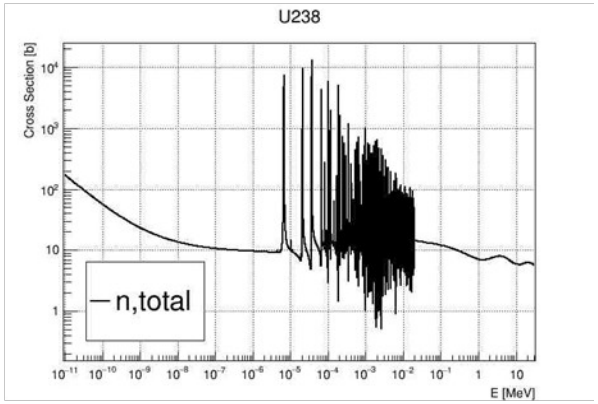

- 80% of computing time in reactor physics calculations is spent in the evaluation of σ(E)
  - This time is mostly spent in the **binary search** to locate E in the energy grid

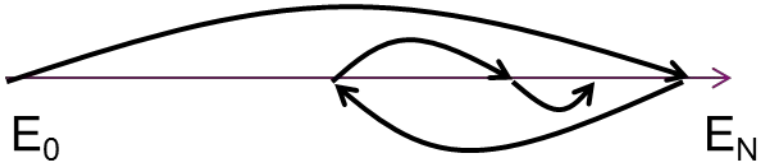

  - The energy array is accessed *randomly*
  - The **floating point unit is mostly idle** because of *cache misses* in the binary search
  - The **vector unit does nothing**

Alternatives: unified grid, hash tables, …

MONTE CARLO TRANSPORT ALGORITHM (4)

On-the-fly Doppler broadening

Cross sections at temperature

- Microscopic cross sections depend on the temperature

$$\sigma(E,T) = \int \sigma(E,0) M(E,T)$$

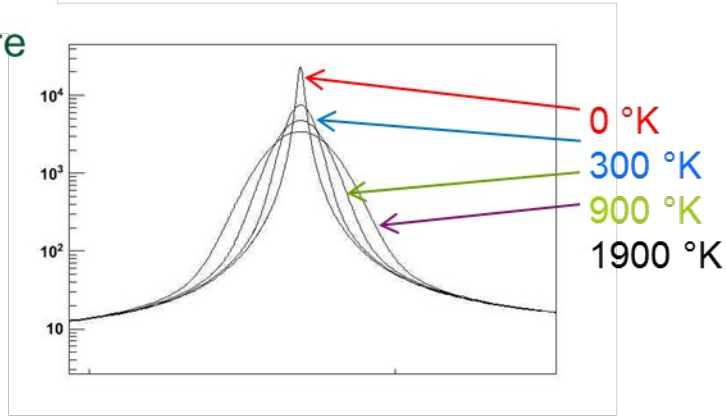

- Traditionally, **Doppler broadening** is **pre-calculated** and loaded in memory for **a few temperatures**
  - For 200 isotopes, cross sections weight more than *1Gigabyte per temperature*: loading hundreds of temperatures is **not** a viable option

- Doppler broadening must be carried out "**on the fly**", that is during the simulation, at each transport flight
  - A few options for the "on the fly" temperature treatment are available today, each with pros and cons
  - We have focused on the "sigma1" method: the most accurate, the *slowest* but with **good potential for vectorization**

Alternatives: multipoles (OpenMC, thesis : Y. Wang and T. Freiman), TMS (Serpent)