Radiation transport Monte Carlo massively parallel computations for fusion neutronics applications

EUROfusion

Arkady Serikov*, Roman Afanasenko, Yuefeng Qiu, Dieter Leichtle, Jin Hun Park

Karlsruhe Institute of Technology (KIT), Institute for Neutron Physics and Reactor Technology, Hermann-von-Helmholtz-Platz 1, 76344 Eggenstein-Leopoldshafen, Germany

**Principal Investigator of the MCHIFI (Monte Carlo High Fidelity) project, Email: arkady.serikov@kit.edu*

6th IFERC workshop on the usage of GPU-based system for fusion applicatons

June 24, 2025 - 9:00-12:15 (CEST) 16:00-19:15 (JST), Video Conference

This work has been carried out within the framework of the EUROfusion Consortium, funded by the European Union via the Euratom Research and Training Programme (Grant Agreement No 101052200 — EUROfusion). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Commission. Neither the European Union nor the European Commission can be held responsible for them.

Part I

Monte Carlo radiation transport massively parallel computations in the framework of the **MCHIFI** (Monte Carlo High Fidelity) project:

2012 – 2016: MCHIFI project foundation with the HPC resources of the F4E Broader Approach (BA) IFERC-CSC Helios supercomputer;

2016 – 2023: MCHIFI project continuation on the EUROfusion Marconi-Fusion HPC;

2024 – present time (June 2025):

- MCHIFI on the EUROfusion Leonardo Booster and Data Centric General Purpose (DCGP) partitions.
- MCHIFI on the EUROfusion Pitagora-CPU partition.

Physics of Monte Carlo radiation transport: parallel computations

Monte Carlo (MC) radiation transport runs on supercomputers

⇒ *Simulation of* <u>*independent*</u> *random pathways on microscopic level, i. e. tracking of individual particle histories from "birth" to "death"*

⇒ *Simulation can be computed on* <u>*parallel*</u> *multiprocessor systems*

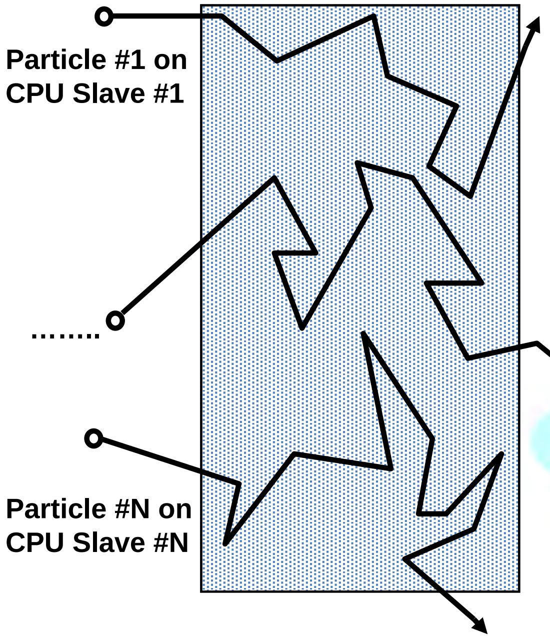

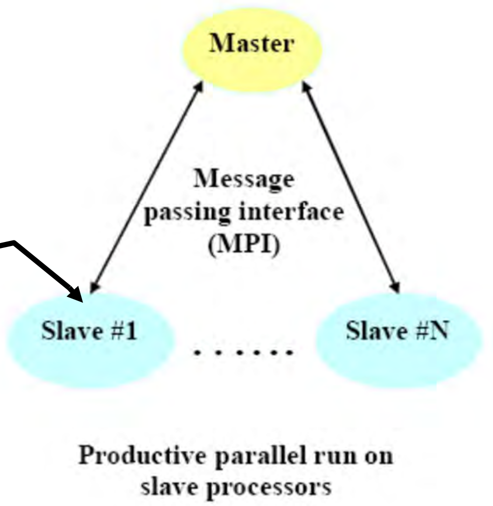

<u>Monte Carlo method</u> is most suitable computational technique for nuclear fusion applications. That is because of the following reasons:

- ***<u>Geometry</u>***: complex fusion devices can be modelled in 3D geometry without major geometry approximations
- ***<u>Data:</u>*** continuous energy representation as stored on evaluated data files in ENDF format
- ***<u>Calculation accuracy:</u>*** only limited by statistics and data uncertainties (no numerical approximations)

Used HPC systems:

- Europe HPC-FF system at JSC (Germany);
- CampusGrid OpusIB, SCC-KIT (Germany);
- Europe HC3 at SCC-KIT (Germany);
- Europe BWGrid (Germany)
- IFERC-CSC Helios (Japan)
- **EUROfusion Marconi-Fusion (CINECA, Italy**)
- **EUROfusion Leonardo Booster and DCGP (CINECA, Italy**)
- **EUROfusion Pitagora-CPU partition (CINECA, Italy**)

The maximum speed-up was found on **EUROfusion HPC Marconi-Fusion** supercomputer with **OpenMP / MPI parallelization** for non-biased MCNP5 run. The speed-up was **<u>2500 on 4096 cores</u>** (or ~850 on 1024 cores, and ~450 on 512 cores) – because the speed-up dependence on the number of cores, it is not linear due to overhead time spent for communications between computing nodes.

Slurm script on Leonardo dcgp submitted to run the MCNP (McDeLicious) job on 4 nodes, with 112 physical cores per node, total number of MPI tasks = 448, used as total number of cores in speed-up metrics of parallel performance.

```
#!/bin/bash
#SBATCH --nodes=4 # nodes
#SBATCH --ntasks-per-node=112 # tasks per node
# #SBATCH --cpus-per-task=1 # cores per task
# #SBATCH --gres=gpu:0 # GPUs per node
# #SBATCH --mem=494000 # mem per node (MB)
#SBATCH --time=24:00:00 # time limit (d-hh:mm:ss)
#SBATCH --account=FUA38_MCINS_0 # account ($ saldo -b  --dcgp )
#SBATCH --partition=dcgp_fua_prod

# #SBATCH --partition=boost_fua_prod # partition name
# #SBATCH --qos=boost_qos_fuadbg # quality of service

module load intel-oneapi-compilers/2023.2.1
module load intel-oneapi-mpi/2021.10.0
export DATAPATH=/leonardo_work/FUA38_MCINS_0/aserikov/C810mnycp00/MCNP_DATA

mpirun ./McDeLi_upd1.mpi i=Pb_7pinLi90_heat o=Pb_7pinLi90_heat_o1 runtpe=Pb_7pinLi90_heat_r1 xsdir=xsdir_2
src_cpu_Pb_7pinLi90_heat (END)
```

```
*****************************************************************************************

dump no.    1 on file Pb_7pinLi90_heat_r2    nps =          0    coll =          0
        0

     472 warning messages so far.
master starting     447 MPI slave tasks with       1 threads each  06/21/25 09:21:15
master set rendezvous nps =   100000000,  work chunks =    447    06/21/25 09:21:19
```

MCNP (McDeLicious) output: initialization of 447 MPI slave tasks (MPI ranks or “MCNP work chunks”). An MPI rank is an independent instance of the program running in parallel. Each rank processes a portion of the workload (e.g., a subset of particle histories in MCNP equaled ~223714 particles per MPI rank).

```
number of histories processed by each MPI task
          0     223713     223714     223713     223714     223714     223713     223714     223714     223713
     223714     223714     223713     223714     223714     223713     223714     223713     223714     223714
     223713     223714     223714     223713     223714     223714     223713     223714     223714     223713
     223714     223714     223713     223714     223713     223714     223714     223713     223714     223714
     223713     223714     223714     223713     223714     223714     223713     223714     223714     223713
     223714     223713     223714     223714     223713     223714     223714     223713     223714     223714
     223713     223714     223714     223713     223714     223714     223713     223714     223713     223714
     223714     223713     223714     223714     223713     223714     223714     223713     223714     223714
     223713     223714     223714     223713     223714     223713     223714     223714     223713     223714
```

MCNP5 parallel performance analysis

Computational wall-clock times have been measured in dependence on the number of slave CPUs used and the size of the job. Results are presented in the left figure in terms of speed-up and the right figure in terms of efficiency, respectively.
The speed-up (S_N) is defined as the ratio between the wall-time running on one processor (T_1) and the wall-time running on N slave processors (T_N), $S_N = T_1/T_N$. The efficiency (E_N) is estimated by the ratio between the speed-up (S_N) and the number of slave processors (N).

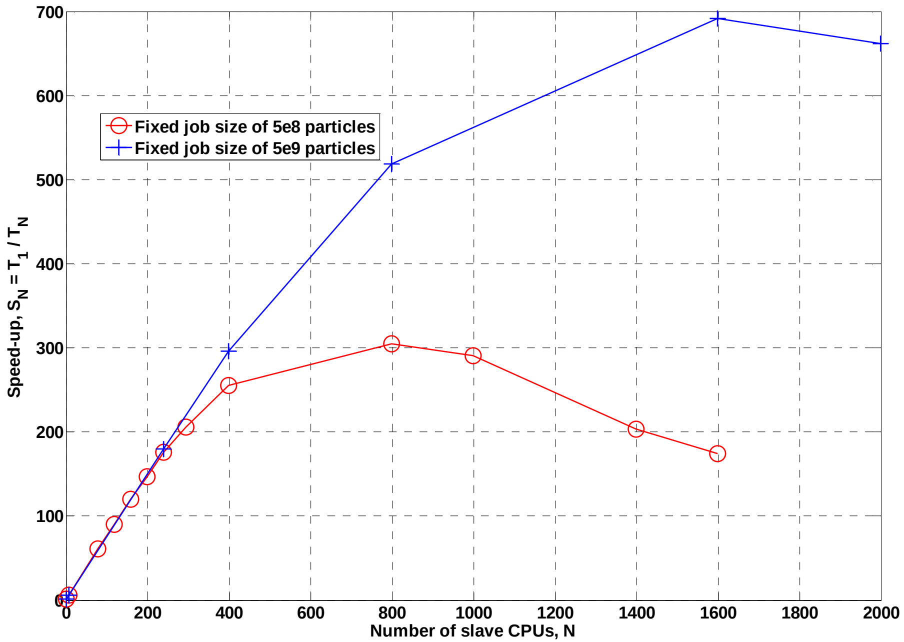

HPC-FF speed-up for two sizes of the MCNP jobs with 5e8 and 5e9 particles.

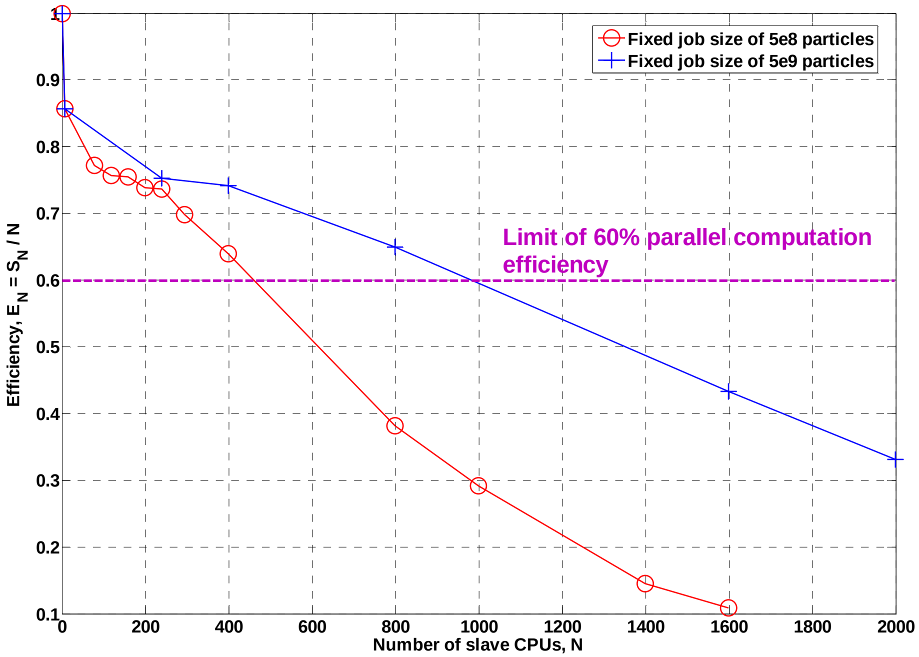

HPC-FF efficiency for two fixed MCNP5 job sizes of 5e8 and 5e9 particles.

The maximum speed-up was found on HPC-FF equaled ~700 with 1600 CPUs, this means that wall-clock time to run the same job on one CPU was 700 times longer than running on 1600 CPUs in parallel on the HPC-FF cluster at Juelich Supercomputing Centre (JSC).

MCHIFI project: fusion neutronics computations on HPCs of F4E BA and EUROfusion

- **MCHIFI** (**M**onte **C**arlo **Hi**gh **Fi**delity) project has been organized for massively parallel computations on the EUROfusion Marconi-Fusion HPC for the most urgent and computationally demanded fusion neutronics tasks.
- The MCHIFI project was founded in 2012 to use the IFERC-CSC Helios supercomputer in the framework of the F4E Broader Approach (BA) to serve the ITER neutronics tasks.

- MCNP5 tested on the F4E Broader Approach IFERC-CSC Helios: 2x8 Intel Sandy-bridge EP processors with 2.7 Hz and 64 GB RAM per node:
  - Excellent scalability of MPI/OpenMP parallel runs of MCNP5 code up to 1024 cores in analogue runs, no variance reduction.
  - Speed-up equals ~450 on 512 cores, and ~850 of speed-up for 1024 cores.
  - OpenMP/MPI hybrid, the satisfactory speed-up of more than 2500 on 4096 cores was achieved for not-biased MCNP5 calculations, as it is illustrated in Figure 1

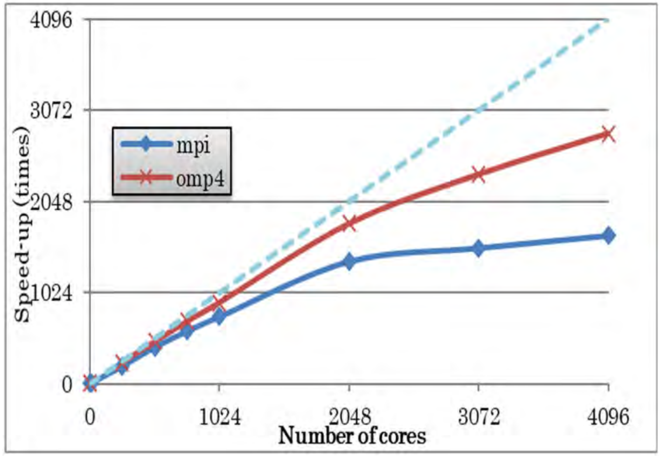

Figure 1. The MCNP5 speed-up on IFERC-CSC Helios supercomputer.

- McDeLicious tested on the EUROfusion HPC Marconi-Fusion with conventional partition (A3) based on INTEL Skylake with peak performance ~9.2 Pflops (2848 nodes). Each node is equipped with 2x24-cores Intel Xeon 8160 CPU (Skylake) at 2.10 GHz and 192 GB of RAM per node.
  - Speed-up MPI-parallel performance has been measured and presented in Figure 2 for the McDeLicious code for IFMIF-DONES radiation deep-penetration shielding tasks with variance reduction.

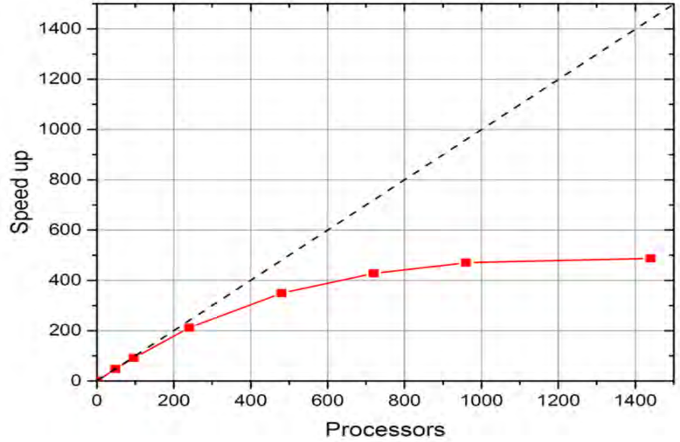

Figure 2. The speed-up of McDeLicious code on Marconi-Fusion HPC.

The optimal number of CPUs used in MCNP5/6 parallel calculations is dependent on complexity of the model. To improve the statistical errors of the MCNP5 results we are using the ADVANTG approach and the recently developed at KIT On-The-Fly (OTF) Monte Carlo variance reduction technique with dynamic Weight Window upper bounds, see Ref. [Yu Zheng, Y. Qiu, "Improvements of the on-the-fly MC variance reduction technique with dynamic WW upper bounds," *Nuclear Fusion* **62** (2022) 086036, https://doi.org/10.1088/1741-4326/ac75fc]

Radiation transport with the MCNP code

MCNP is a code for radiation transport calculations in 3D geometry. The abbreviation is translated as **M**onte **C**arlo **N**-**P**article.

Neutron, photon, electron, or coupled neutron/photon/electron transport can be performed by MCNP. The MCNP code was developed by X-5 Monte Carlo Team in Los Alamos National Lab. (LANL), USA.

Contributors to MCNP6.2

Jerawan Armstrong[1], Forrest B. Brown[1], Jeffrey S. Bull[1], Laura Casswell[1], Lawrence J. Cox[1], David Dixon[1], R. Arthur Forster[1], John T. Goorley[1], H. Grady Hughes[1], Jeffrey Favorite[1], Roger Martz[1], Stepan G. Mashnik[1], Michael E. Rising[1], Clell(CJ) Solomon[1], Avneet Sood[1], Jeremy E. Sweezy[1], Christopher J.Werner[1], Anthony Zukaitis[1], Casey Anderson[2], Jay S. Elson[2], Joe W. Durkee[2], Russell C. Johns[2], Gregg W. McKinney[2], Garrett E. McMath[2], John S. Hendricks[3], Denise B. Pelowitz[3], Richard E. Prael[4], Thomas E. Booth[4], Michael R. James[5], Michael L. Fensin[6], Trevor A. Wilcox[7], Brian C. Kiedrowski[8]

[1] XCP-3 Monte Carlo Methods, Codes, and Applications, Los Alamos National Laboratory
[2] NEN-5 Systems Design and Analysis, Los Alamos National Laboratory
[3] NEN-5 Systems Design and Analysis, Los Alamos National Laboratory, Guest Scientist
[4] XCP-3 Monte Carlo Methods, Codes, and Applications, Los Alamos National Laboratory, contractor
[5] C-NR Nuclear & Radiochemistry, Los Alamos National Laboratory
[6] W-13 Advanced Engineering Analysis, Los Alamos National Laboratory
[7] XTD –SS Safety & Surety, Los Alamos National Laboratory
[8] XCP-3 Monte Carlo Methods, Codes, and Applications, Los Alamos National Laboratory, Guest Scientist

History of the MCNP code development [1]

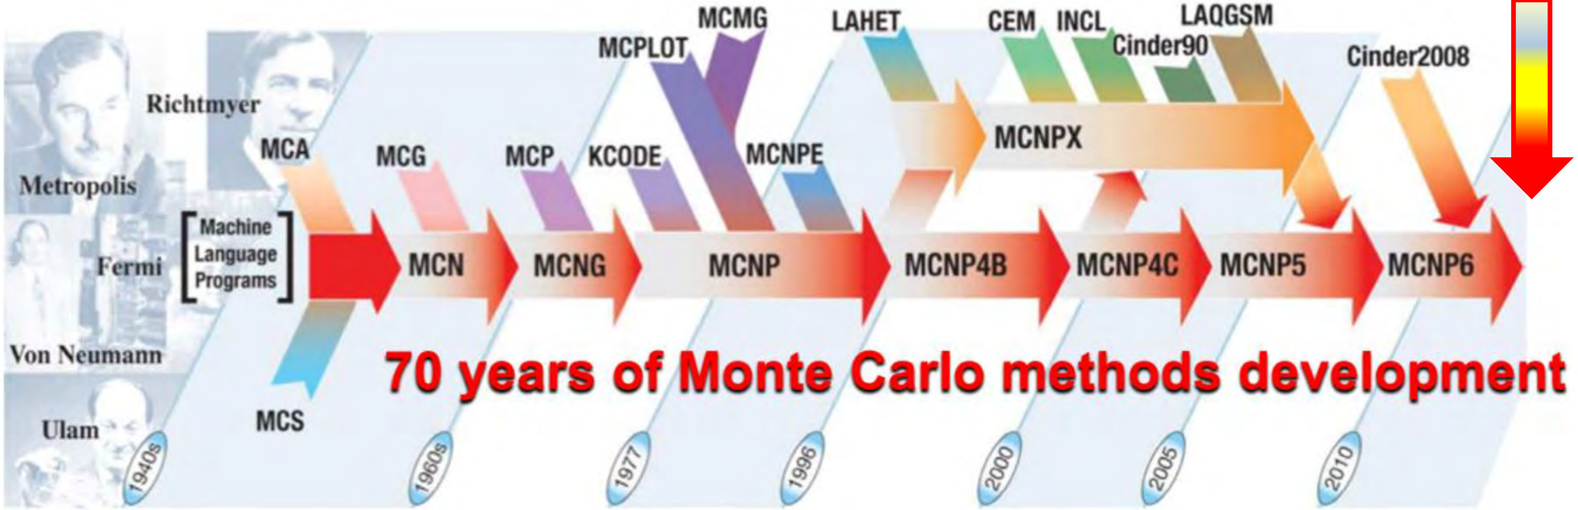

Reference:
[1] Avneet Sood, 2017. The Monte Carlo Method and MCNP-A Brief Review of Our 40 Year History, Presentation to the International Topical Meeting on Industrial Radiation and Radioisotope Measurement Applications Conference.

MCNP6 versatility with OpenMP /MPI parallelization on CPUs

- MCNP6 performs continuous-energy transport of 36 different particle types, plus heavy-ion, fuel burnup, and high-fidelity delayed gamma emission. MCNP6 is written in Fortran 90, has been parallelized (OpenMP and MPI), and works on CPU-based platforms including PCs, workstations, Linux clusters, and supercomputers.
- MCNP6 has thousands of MCNP users worldwide.

Examples of the MCNP6.2 applications:

- **Neutronics and nuclear reactor design**
- **Nuclear criticality safety**
- **Fusion neutronics**
- **Transmutation, activation, and burnup in reactor and other systems**
- Nuclear safeguards
- Medical physics, especially proton and neutron therapy
- Design of accelerator spallation targets, particularly for neutron scattering facilities
- Investigations for accelerator isotope production and destruction programs, including the transmutation of nuclear waste
- Research into accelerator-driven energy sources
- Accelerator based imaging technology such as neutron and proton radiography
- Detection technology using charged particles via active interrogation
- Design of shielding in accelerator facilities
- Activation of accelerator components and surrounding groundwater and air
- High-energy dosimetry and neutron detection
- Investigations of cosmic-ray radiation backgrounds and shielding for high altitude aircraft and spacecraft
- Single-event upset in semiconductors from cosmic rays in spacecraft or from the neutron component on the earth’s surface
- Analysis of cosmo-chemistry experiments, such as Mars Odyssey
- Charged-particle propulsion concepts for spaceflight
- Investigation of fully coupled neutron and charged-particle transport for lower-energy applications
- Nuclear material detection
- Design of neutrino experiments

Particles transported by MCNP6. As listed in Chapter 2 of the MCNP® 6.2 USER'S MANUAL, October 27, 2017

Table 2-2. MCNP6 Particles

IPT*	Name of Particle	Symbol	Mass[1] (MeV)	Low Kinetic Energy Cutoff / Default Cutoff (MeV)	Mean Lifetime[1] (seconds) As treated by MCNP6	Actual (if different)
1	neutron (n)	N	939.56563	0.0 / 0.0	no decay	887.0
2	photon (γ)	P	0.0	1.e-6 / 1.e-3	1×1029	
3	electron (e)	E	0.511008	1.e-5 / 1.e-3	1×1029	
4	negative muon (μ)	|	105.658389	1.e-3 / 0.11261	2.19703×10-6	
5	anti neutron ( )	Q	939.56563	0.0 / 0.0	no decay	887.0
6	electron neutrino (ve)	U	0.0	0.0 / 0.0	1×1029	
7	muon neutrino (vm)	V	0.0	0.0 / 0.0	no decay	
8	positron (e^+) (See Note 1)	F	0.511008	1.e-3 / 1.e-3	1×10^{29}	
9	proton (p+)	H	938.27231	1.e-3 / 1.0	1×1029	
10	lambda baryon (Λ0)	L	1115.684	1.e-3 / 1.0	DOP†	2.632×10-10
11	positive sigma baryon (Σ+)	+	1189.37	1.e-3 / 1.26760	DOP†	7.99×10-11
12	negative sigma baryon(Σ)	-	1197.436	1.e-3 / 1.26760	DOP†	1.479×10-10
13	cascade; xi baryon (Ξ0)	X	1314.9	1.e-3 / 1.0	DOP†	2.9×10-10
14	negative cascade; negative xi baryon (Ξ)	Y	1321.32	1.e-3 / 1.40820	DOP†	1.639×10-10
15	omega baryon (Ω)	O	1672.45	1e-3 / 1.78250	DOP†	8.22×10-11
16	positive muon (μ^+)	!	105.658389	1.e-3 / 0.11261	2.19703×10^{-6}	
17	anti electron neutrino ($\overline{\nu}_e$)	<	0.0	0.0 / 0.0	1×10^{29}	
18	anti muon neutrino ($\overline{\nu}_m$)	>	0.0	0.0 / 0.0	no decay	
19	anti proton ($\overline{p}$)	G	938.27231	1.e-3 / 1.0	1×10^{29}	
20	positive pion (π+)	/	139.56995	1.e-3 / 0.14875	2.603×10-8	
21	neutral pion (π0)	Z	134.9764	0.0 / 0.0	8.4×10-17	
22	positive kaon (K+)	K	493.677	1.e-3 / 0.52614	1.2371×10-8	
23	kaon, short (K0S)	%	497.672	0.0 / 0.0	0.8926×10-10	
24	kaon, long (K0L)	^	497.672	0.0 / 0.0	5.17×10-8	
25	anti lambda baryon ($\overline{\Lambda}^0$)	B	1115.684	1.e-3 / 1.0	DOP†	2.632×10^{-10}
26	anti positive sigma baryon ($\overline{\Sigma}^+$)	_	1189.37	1.e-3 / 1.26760	DOP†	7.99×10^{-11}
27	anti negative sigma baryon ($\overline{\Sigma}^-$)	~	1197.436	1.e-3 / 1.26760	DOP†	1.479×10^{-10}
28	anti cascade; anti neutral xi baryon ($\overline{\Xi}^0$)	C	1314.9	1.e-3 / 1.0	DOP†	2.9×10^{-10}
29	positive cascade; positive xi baryon (Ξ^+)	W	1321.32	1.e-3 / 1.40820	DOP†	1.639×10^{-10}
30	anti omega ($\overline{\Omega}^-$)	@	1672.45	1e-3 / 1.78250	DOP†	8.22×10^{-11}
31	deuteron (d)	D	1875.627	1.e-3 / 2.0	1×1029	
32	triton (t)	T	2808.951	1.e-3 / 3.0	12.3 years	
33	helion (3He)	S	2808.421	1.e-3 / 3.0	1×1029	
34	alpha particle (α)	A	3727.418	1.e-3 / 4.0	1×1029	
35	negative pion (π^-)	*	139.56995	1.e-3 / 0.14875	2.603×10^{-8}	
36	negative kaon (K^-)	?	493.677	1.e-3 / 0.52614	1.2371×10^{-8}	
37	heavy ions‡	#	varies	1.e-3 / 5.0	1×10^{29}	

† DOP=Decayed on production

‡ The "#" symbol represents all possible heavy-ion types—in other words, any ion that is not one of the four light ions available in MCNP6.

** A list of heavy ions available for transport is provided in Appendix G.

McDeLicious code development history

McDeLicious is an extension to the **MCNP Monte Carlo code** with the ability to simulate the generation of source neutrons based on deuteron - lithium (**D-Li**) interaction processes

- 1999: **McDeLi** ***(****P. Wilson, Report FZKA 6218, 1999):*
  - An enhancement to MCNP-4a to sample the generation of d-Li source neutrons based on embedded analytical formulas representing direct deuteron striping (Serber model) and compound reactions.
- 2001: **McDeLicious** *(S.P.Simakov et al. J.Nucl.Mat.307-311(2002)1710, FZKA 6743)*
  - An enhancement to MCNP-4b,c to sample the d-Li source neutrons on the basis of tabulated double-differential d + 6,7Li cross-sections for deuteron energies up to 50 MeV (evaluated by *A. Konobeyev et al., NSE 139 (2001)1).*
- 2005: **McDeLicious-05** – compilation with MCNP-5 and use tabulated double-differential cross-sections from updated d + 6,7Li evaluation *(made by P. Pereslavtsev et al., J.Nucl.Mat.367-370(2007)1531).*
- 2011: **McDeLicious-11** - a new approach is implemented to enable direct sampling from the tabulated deuteron beam distribution data without using fitting functions. In this approach, the beam entry position is sampled from tabulated data representing the intensity distribution of the impinging deuteron beam – (*S. P. Simakov et al., "Status of the McDeLicious approach for the D-Li neutron source term modeling in IFMIF neutronics calculations," Fusion Sci. Technol., 62 (2012), pp. 233-239*)
- 2017: **McDeLicious-17** – the actual version of McDeLicious upgraded to MCNP version 6.1.0, an extension of the MCNP Monte Carlo code with the capability to simulate the deuterium-lithium neutron source on the basis of evaluated d + 6,7Li cross-section data. This code has been tested and confirmed to generate identical source particle data as the previous version McDeLicious-11 – (Y. Qiu et al., "IFMIF-DONES HFTM neutronics modeling and nuclear response analyses," Nuclear Materials and Energy, 15 (2018), pp. 185-189)

Simple Model: cylindrical solid Li inside Aluminum capsule

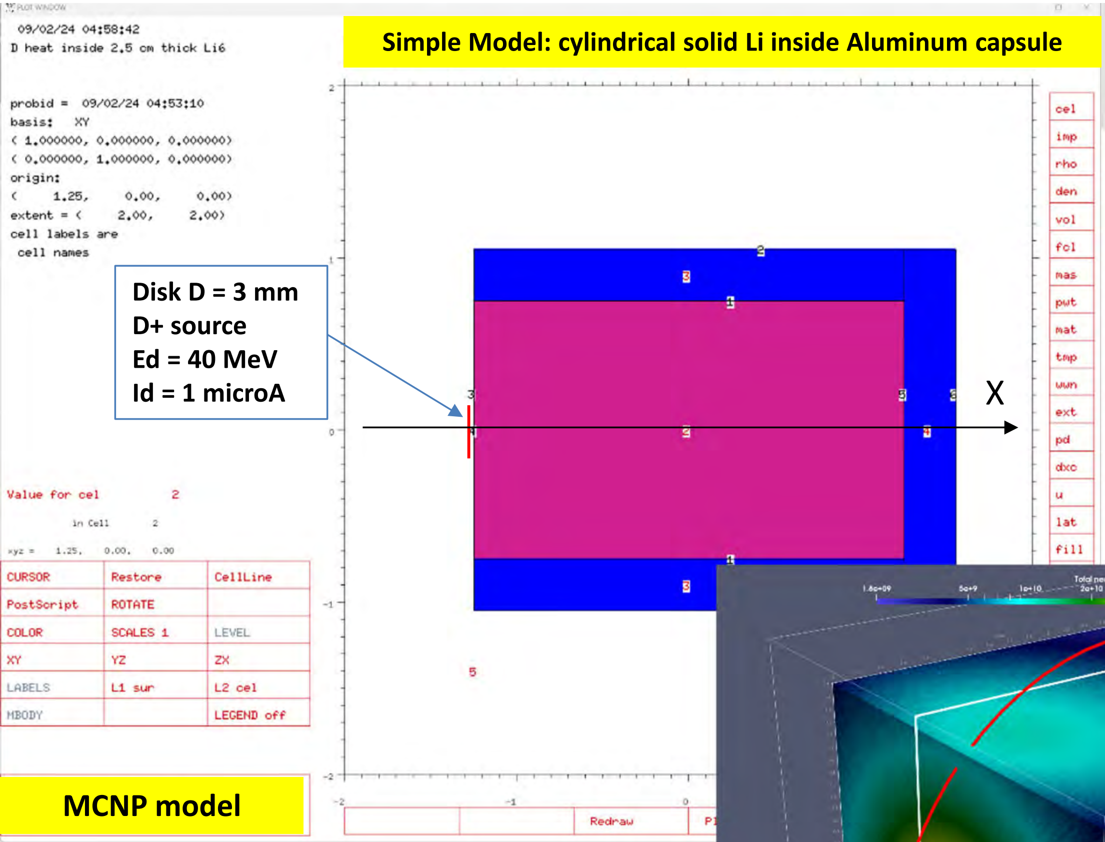

MCNP model

V1: Isotropic

uniformly distributed disk D+ source

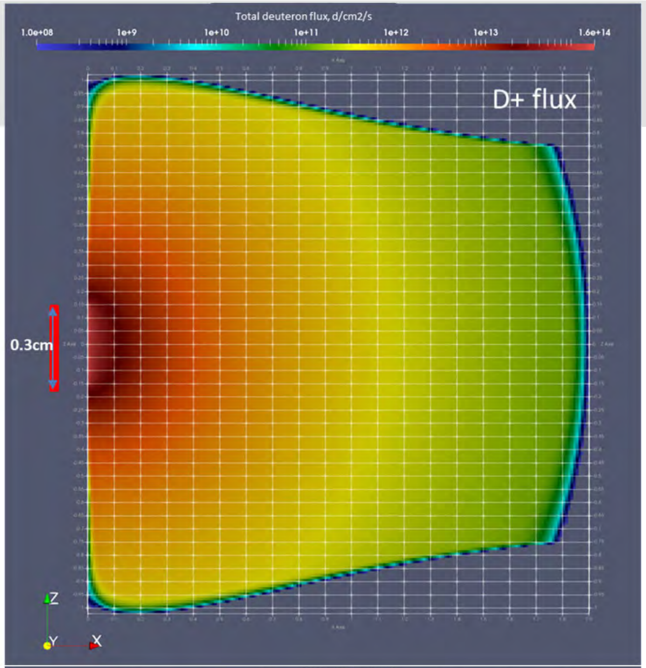

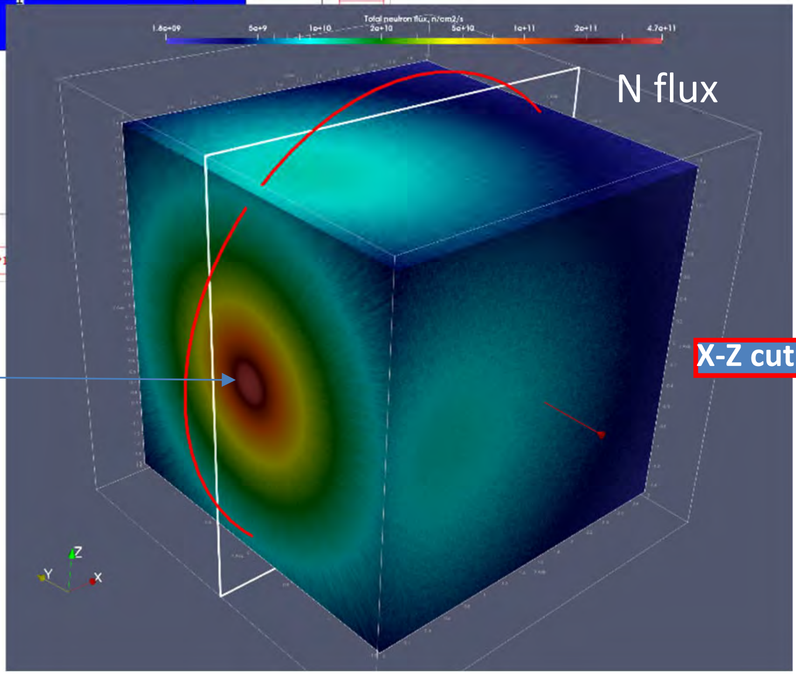

X-Z cut

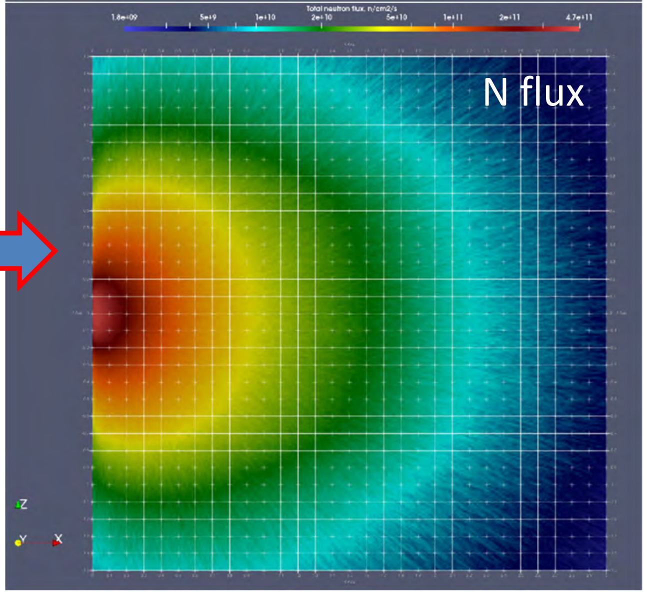

V1: **<u>Isotropic</u>**
uniformly distributed
disk D+ source

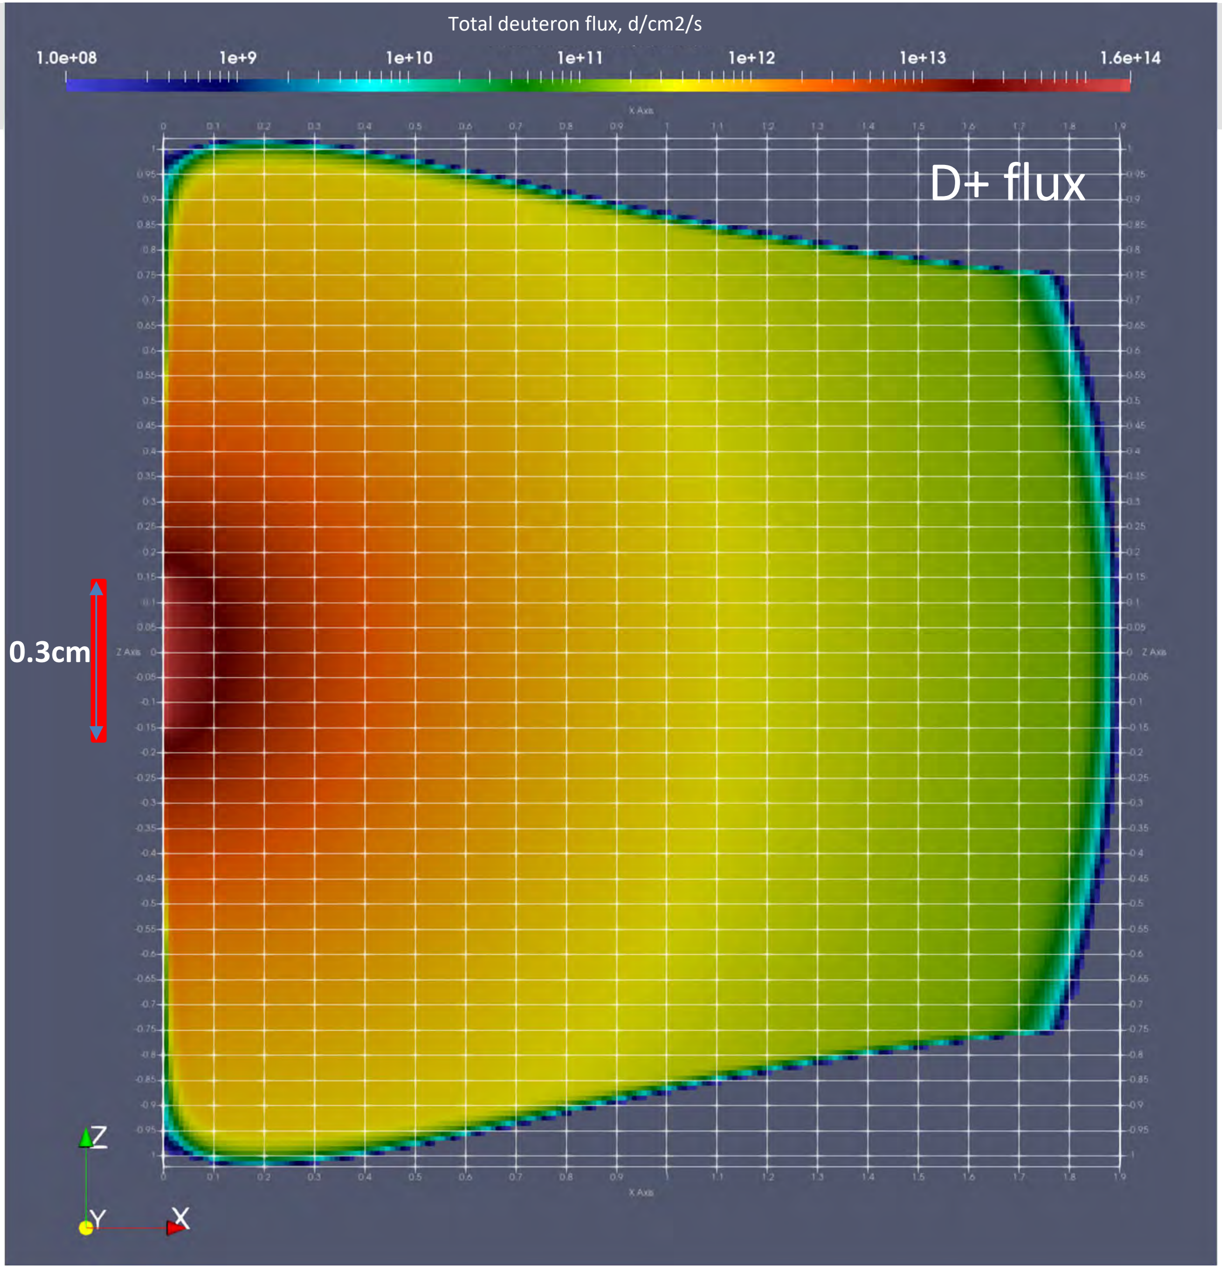

V2: Monodirectional along X-axis disk D+ source, Ed = 40 MeV, Id = 1 microA

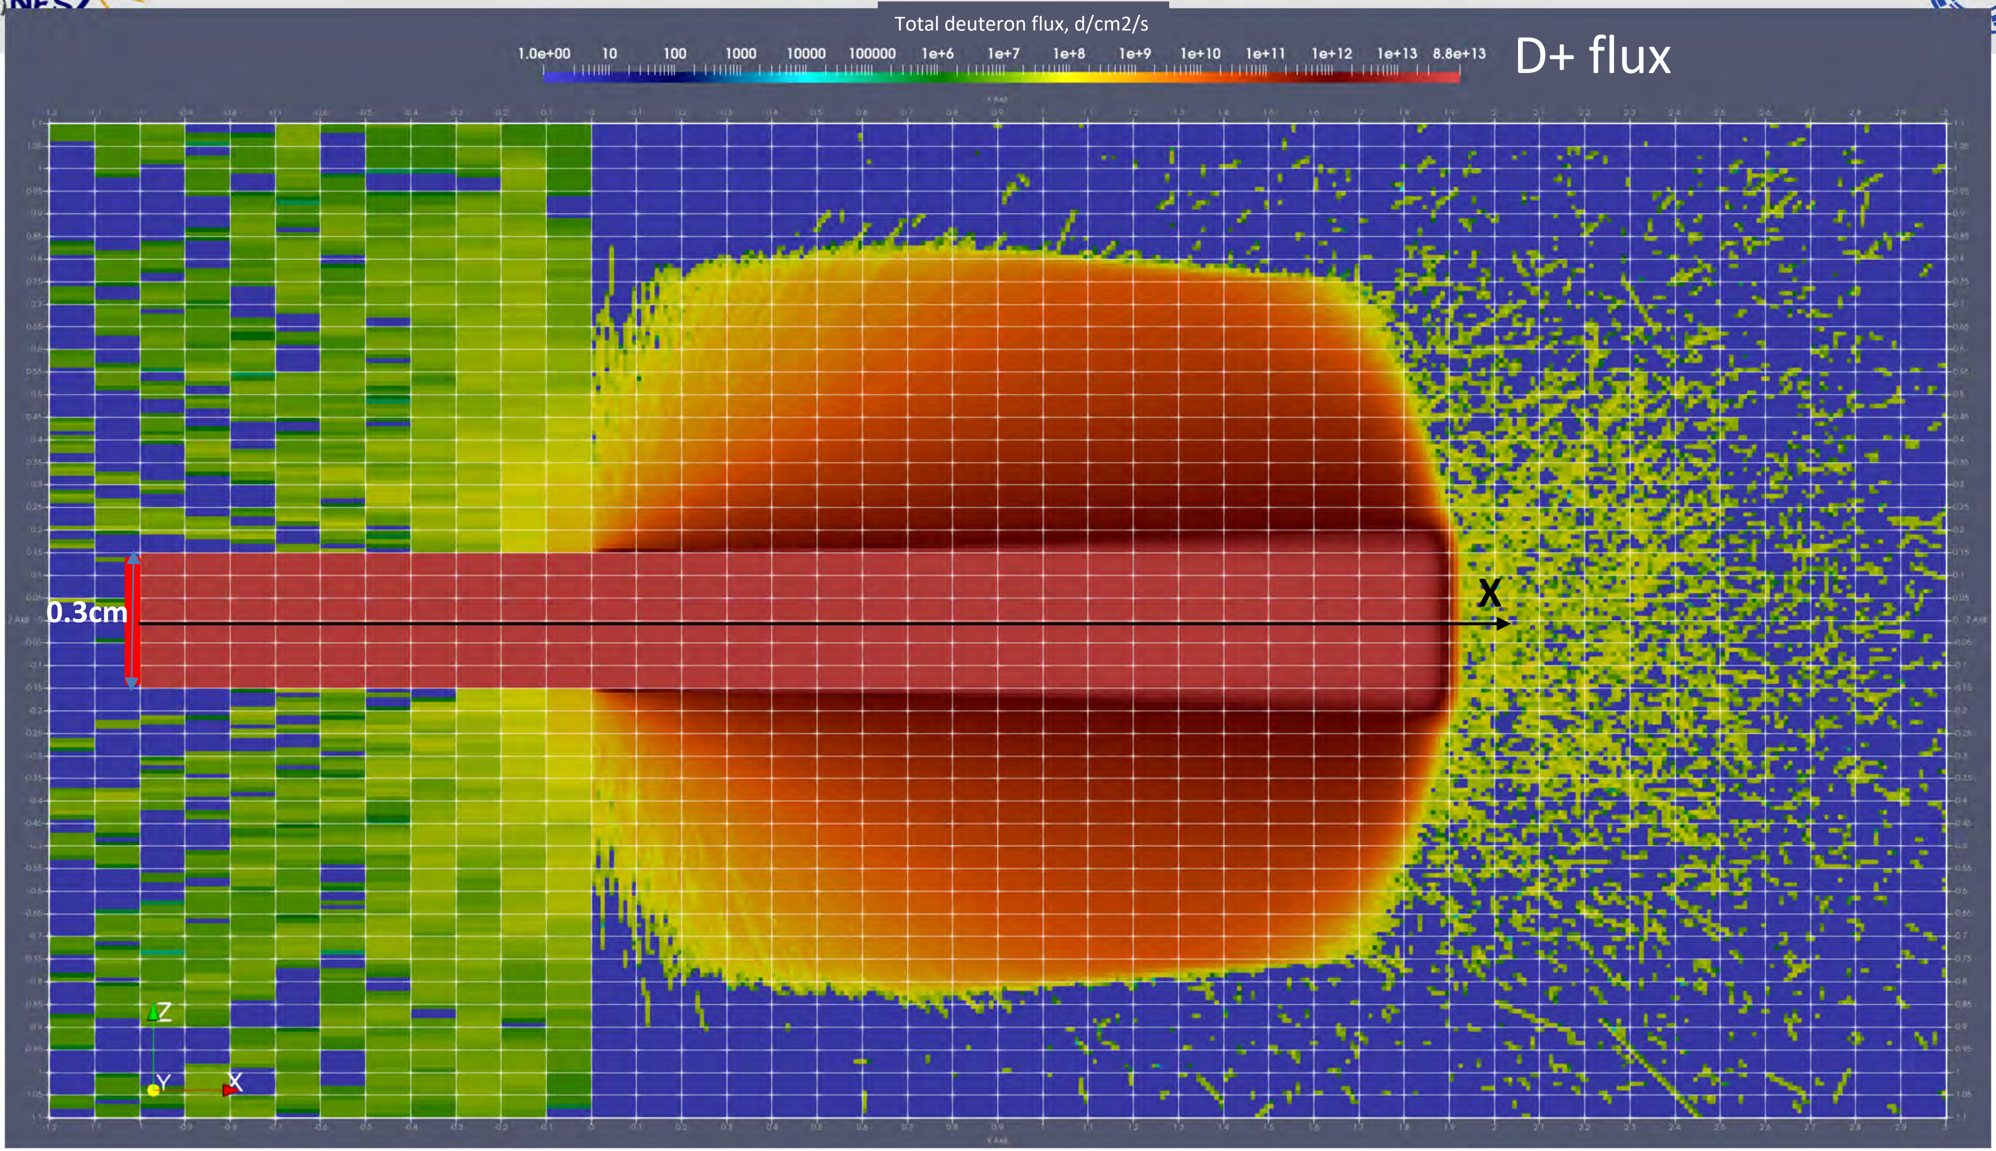

Deuteron slowing down and energy deposition in Li-jet

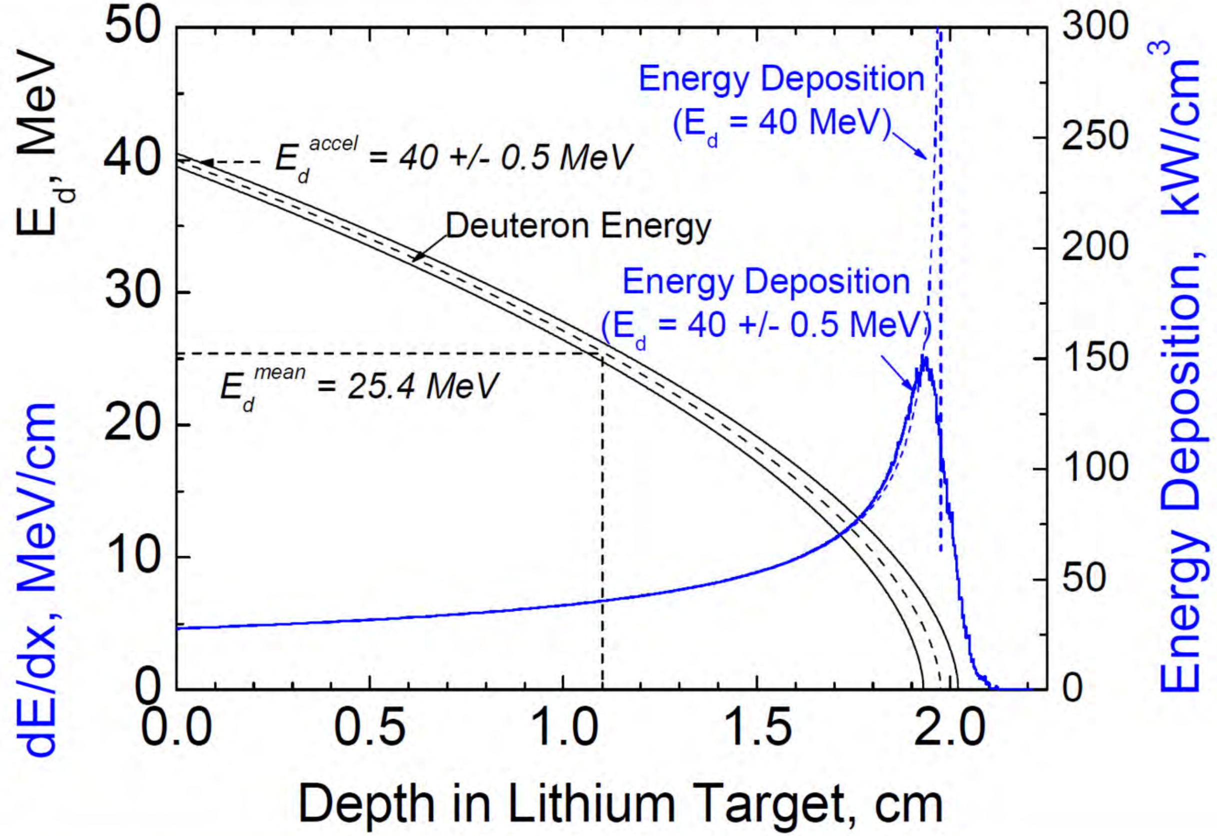

- Deuteron track length reach 2.1 cm;

- Peak Energy Deposition is 150 kW/cc at the depth of 2.0 cm (at the end of d-track)

- Average energy deposition in Li jet = (40 MeV x 250 mA=10,000 kW)/(20x5x2 cc) =50 kW/cc

Ref: S.P. Simakov, McDelicious Workshop, FZK/IRS, *Institut für Reaktorsicherheit, Forschungszentrum Karlsruhe, 13-14 March 2008*

Deuteron track depth (longitudinal, Ld) dependence on the D+ energy (Ed)

Ld=19.8 mm for Ed=40.0 MeV

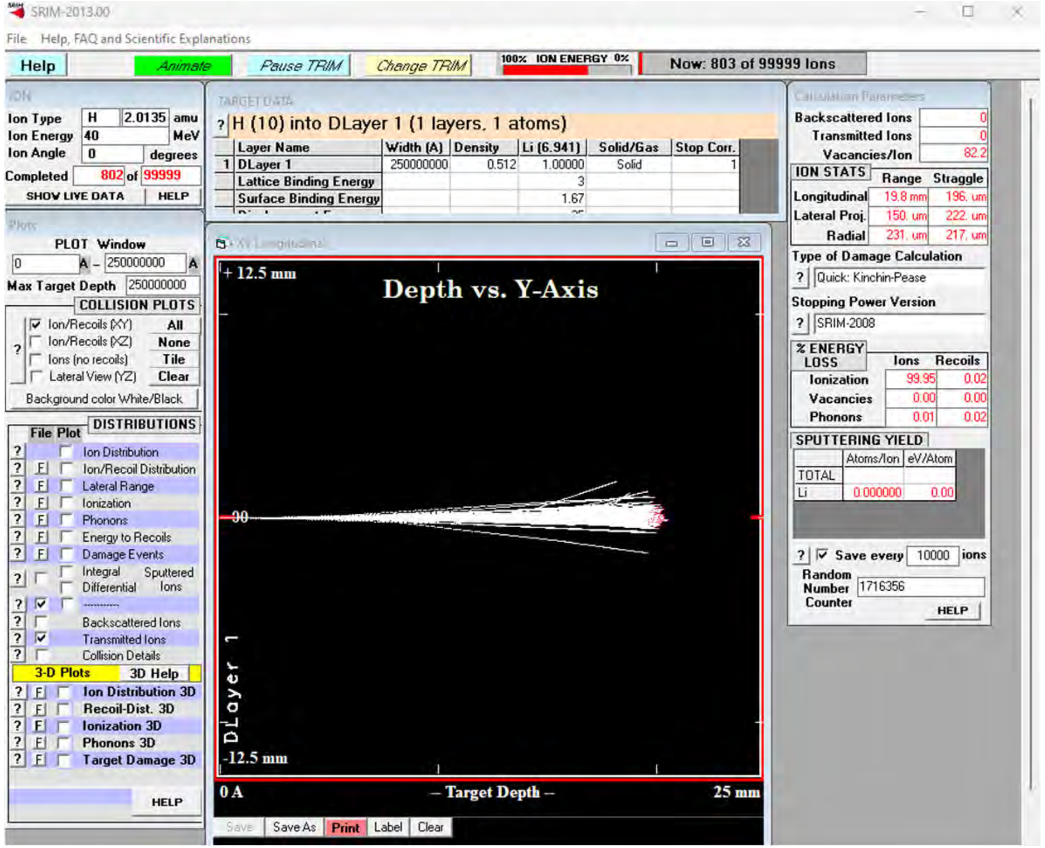

Ld=20.3 mm for Ed=40.5 MeV

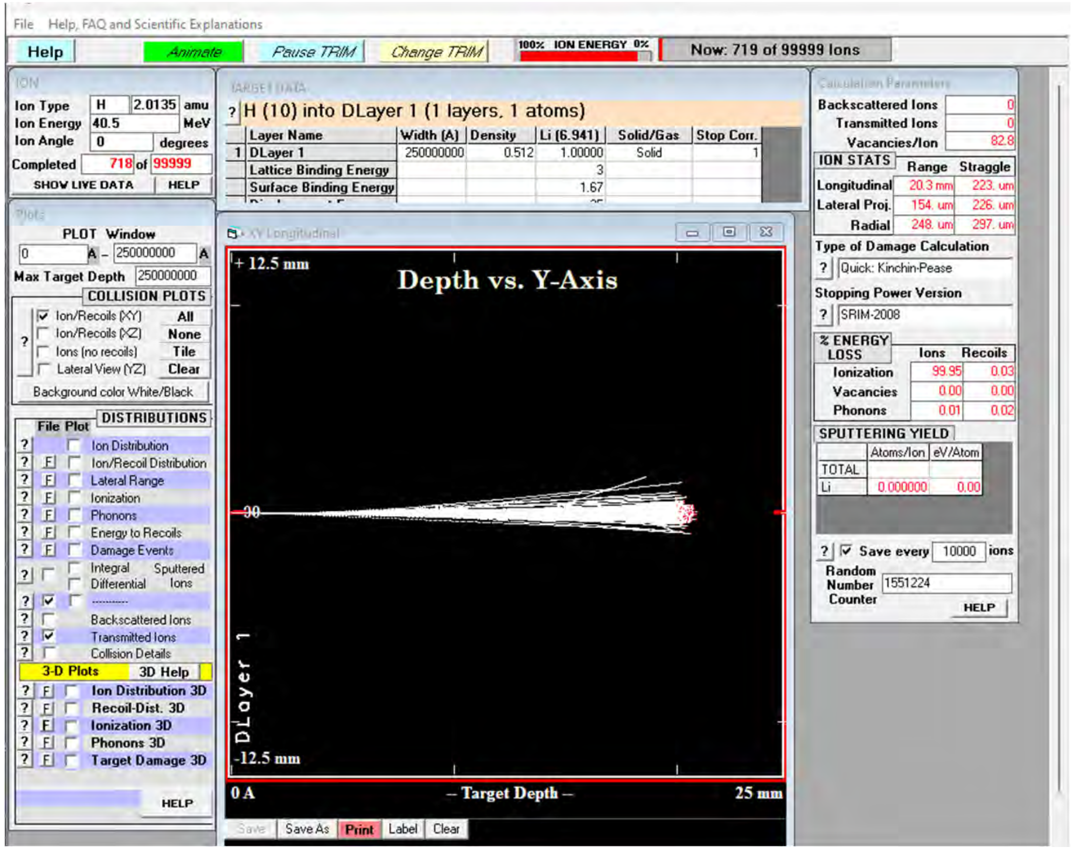

Part II

Application of the EUROfusion HPC (Marconi-Fusion, Leonardo) resources in the **MCHIFI** project for solving the IFMIF-DONES radiation shielding tasks

DONES Programme, IFMIF-DONES España to build the DONES facility

The International Fusion Materials Irradiation Facility—DEMO Oriented NEutron Source (IFMIF-DONES) aims to evaluate and validate the structural and functional materials for developing DEMO-type reactors. To achieve this ambitious goal, several projects have been promoted in recent years, which together form the DONES Programme.

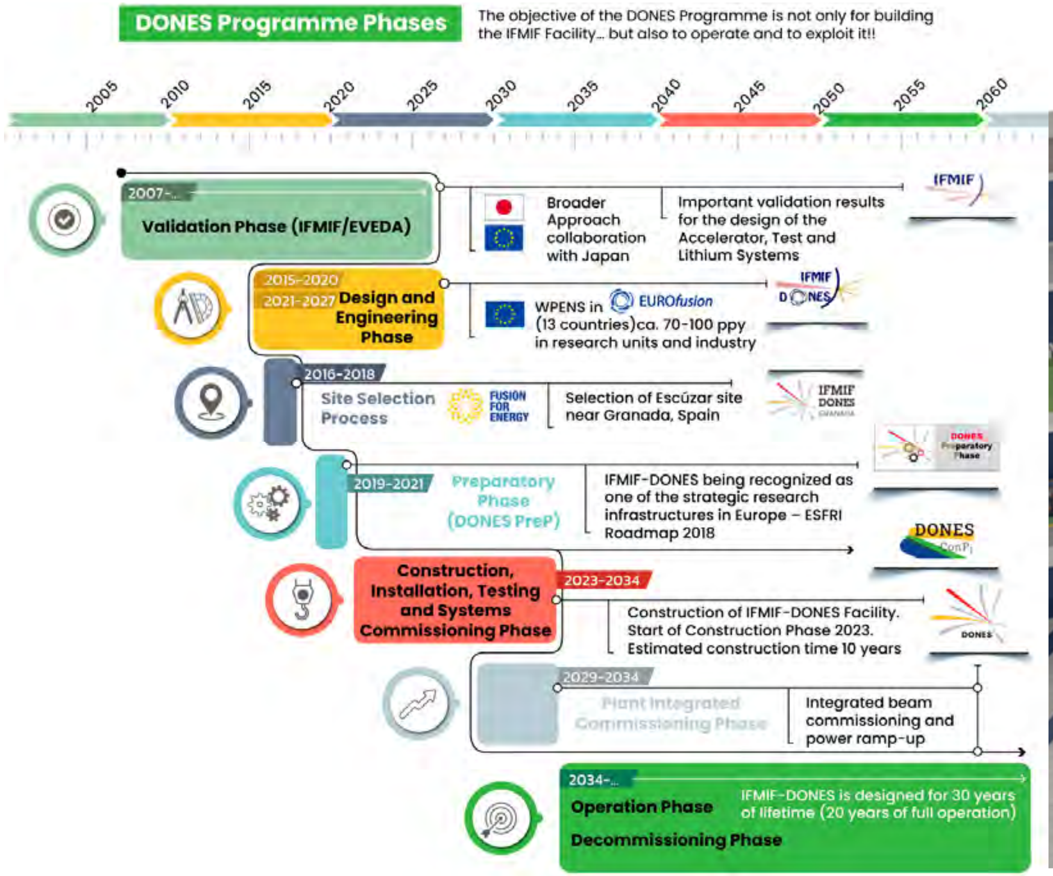

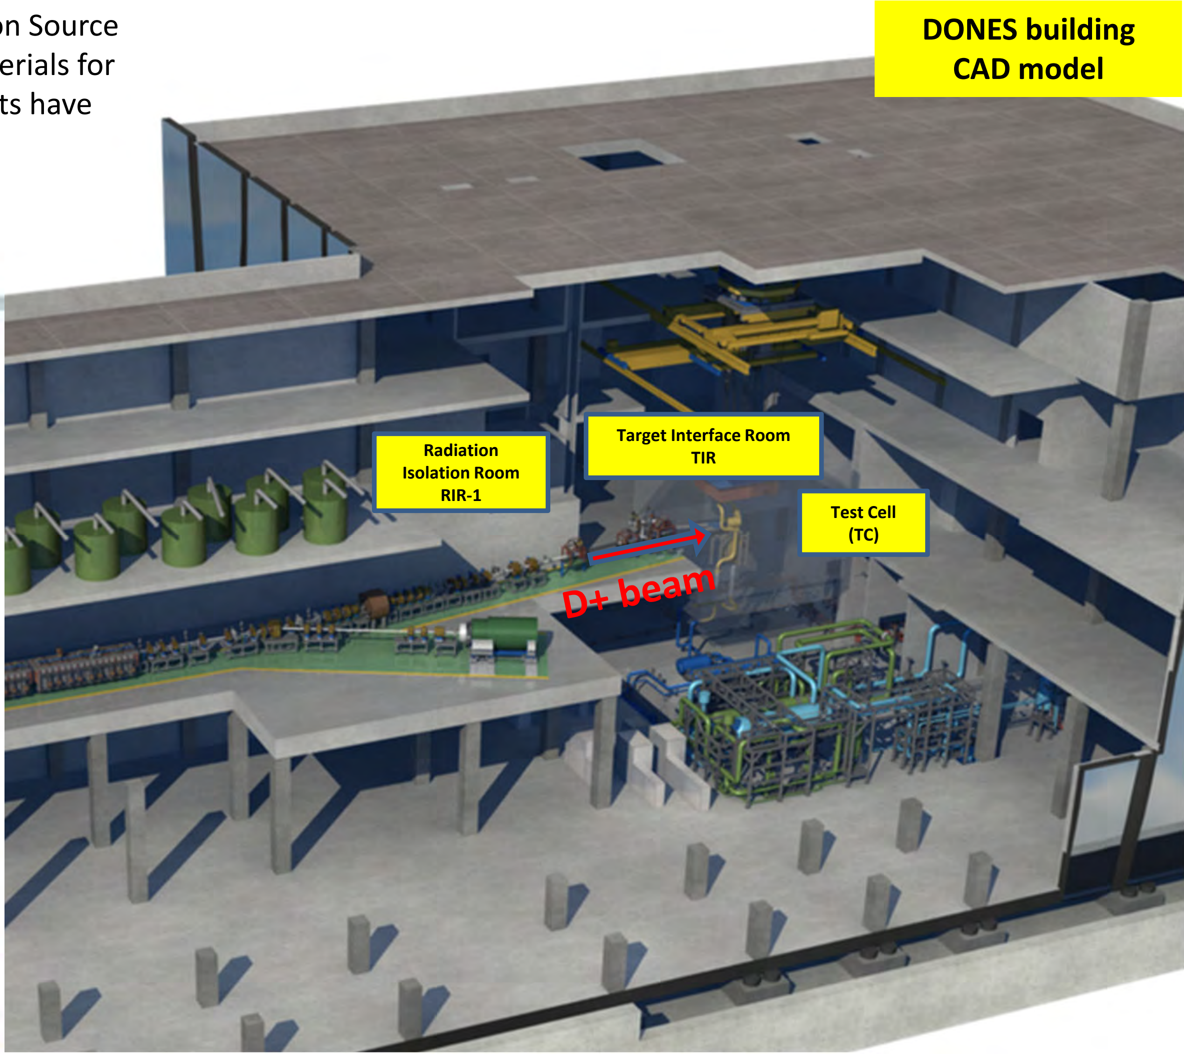

IFMIF-DONES neutronics simulations

- The CAD model of IFMIF-DONES building is prepared (simplified and decomposed) for the CAD-to-MCNP conversion using the codes: **McCad (INR-KIT developed) or SuperMC (developed by FDS-team, China)**
- **McDeLicious-17 code package developed at INR-KIT** – an MCNP6 code modification for deuteron-lithium (d-Li) nuclear reactions in Li of IFMIF-DONES Test Cell. The beam of deuteron ions accelerated up to 40 MeV with current of 125 mA impinges the liquid Li target delivering 5 MW power. The Li target volume is 5×20×2.5 cm^3

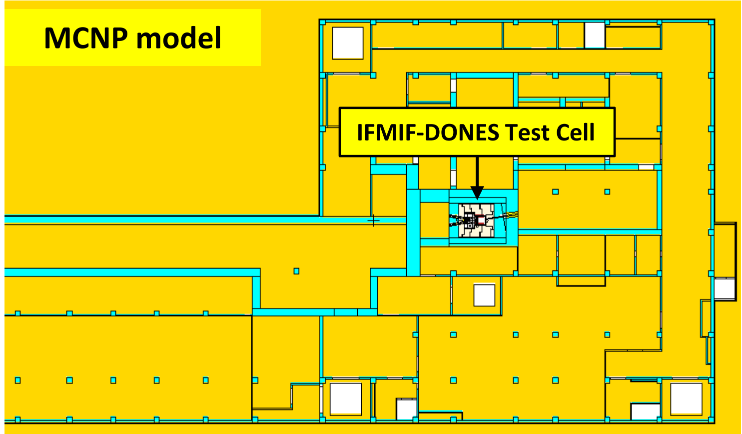

DONES building model horizontal cut at the beam level.

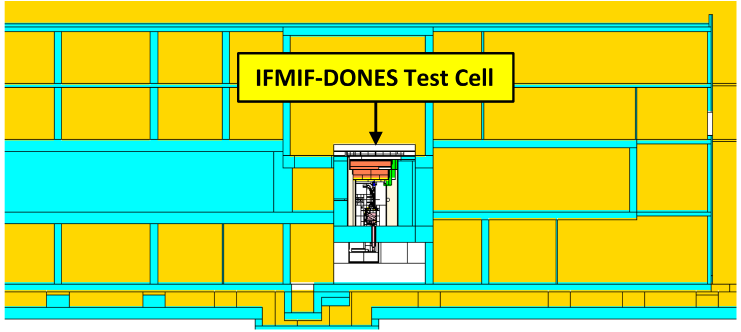

DONES building model vertical cut at the target center

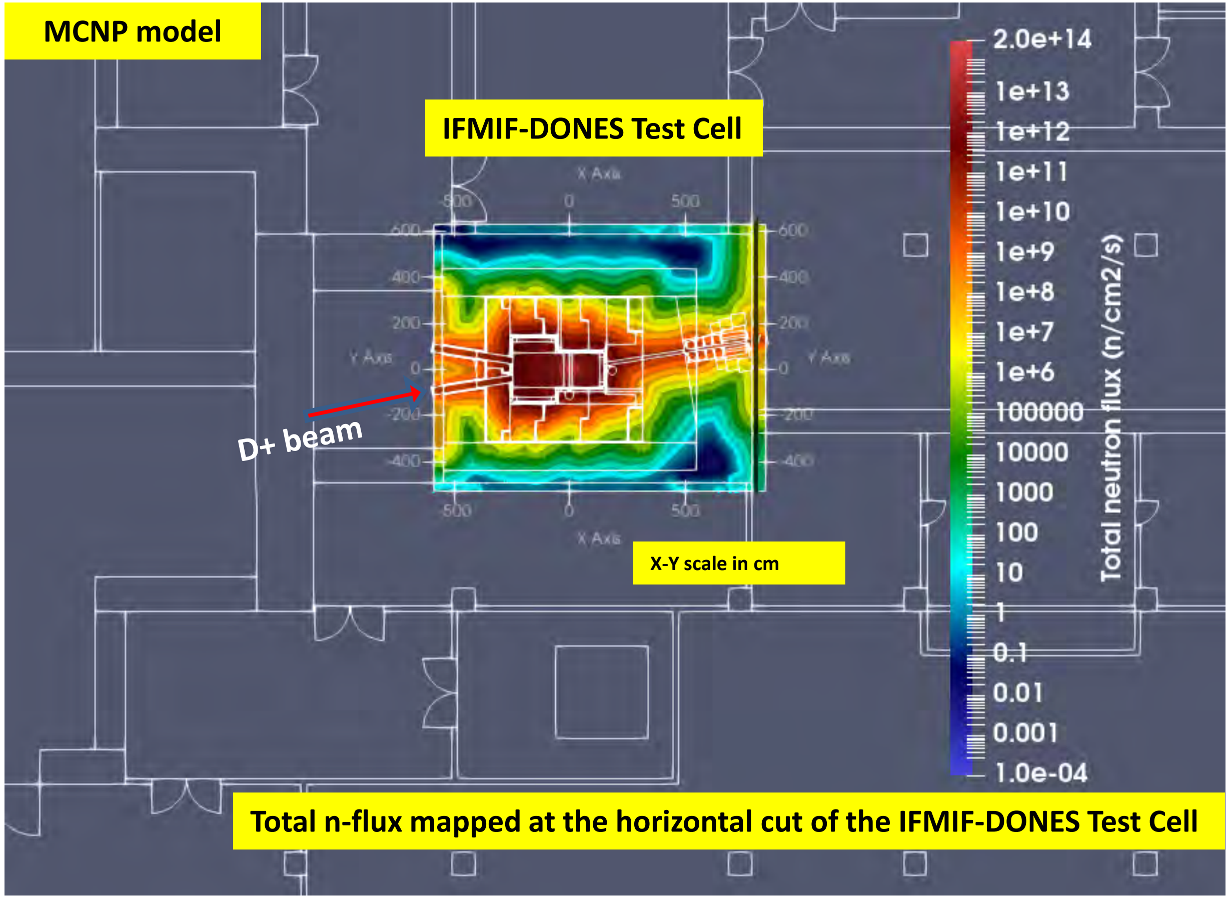

Total n-flux mapped at the horizontal cut of the IFMIF-DONES Test Cell

Neutronics geometry of the accelerator systems

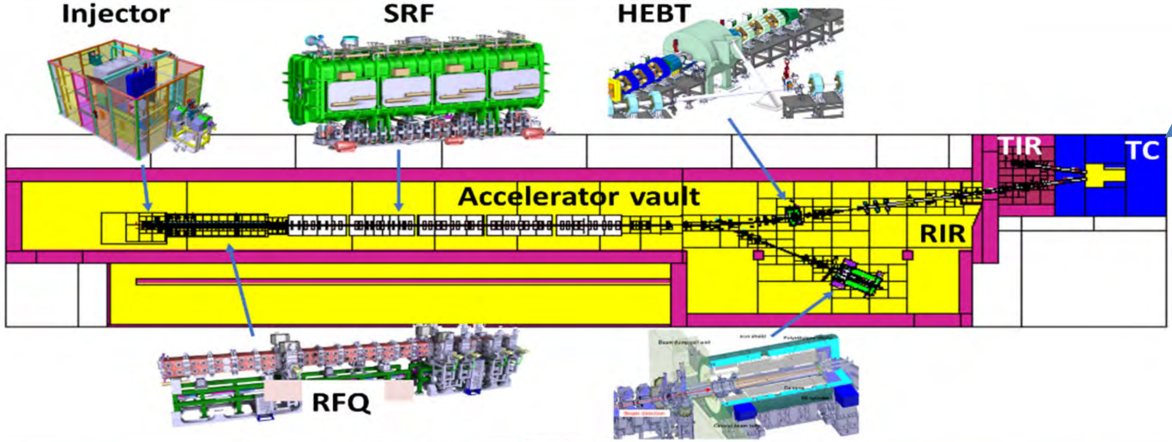

Test Cell (TC) houses the Target Assembly (TA) and the High Flux Test Module (HFTM)

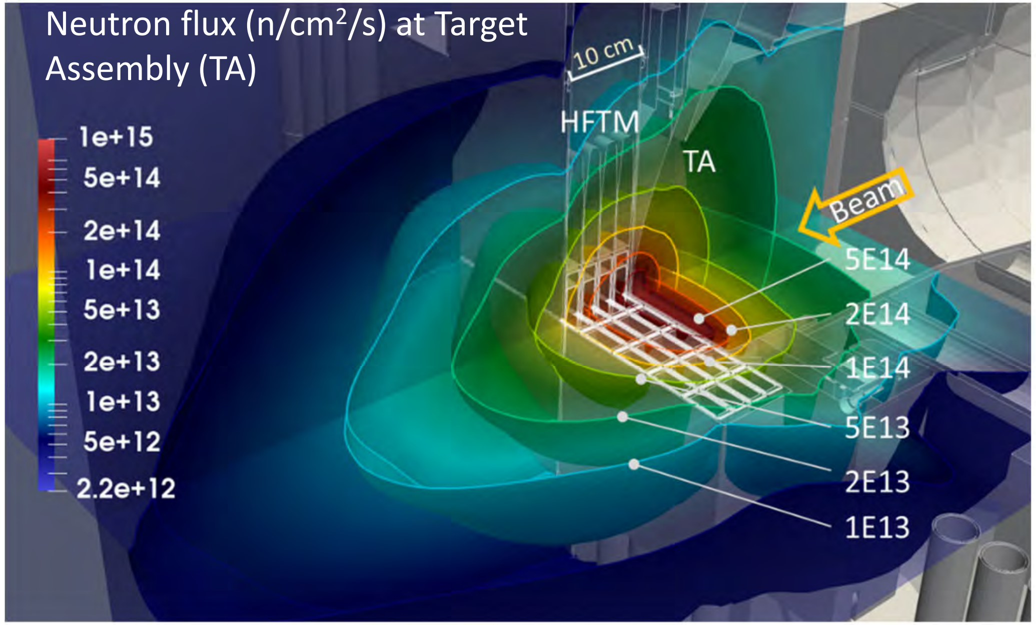

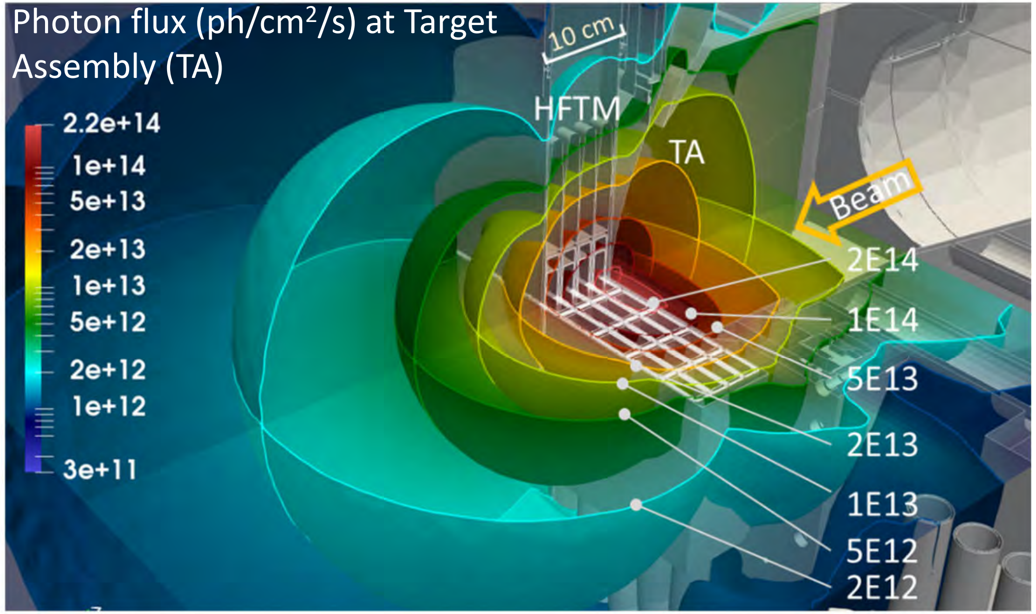

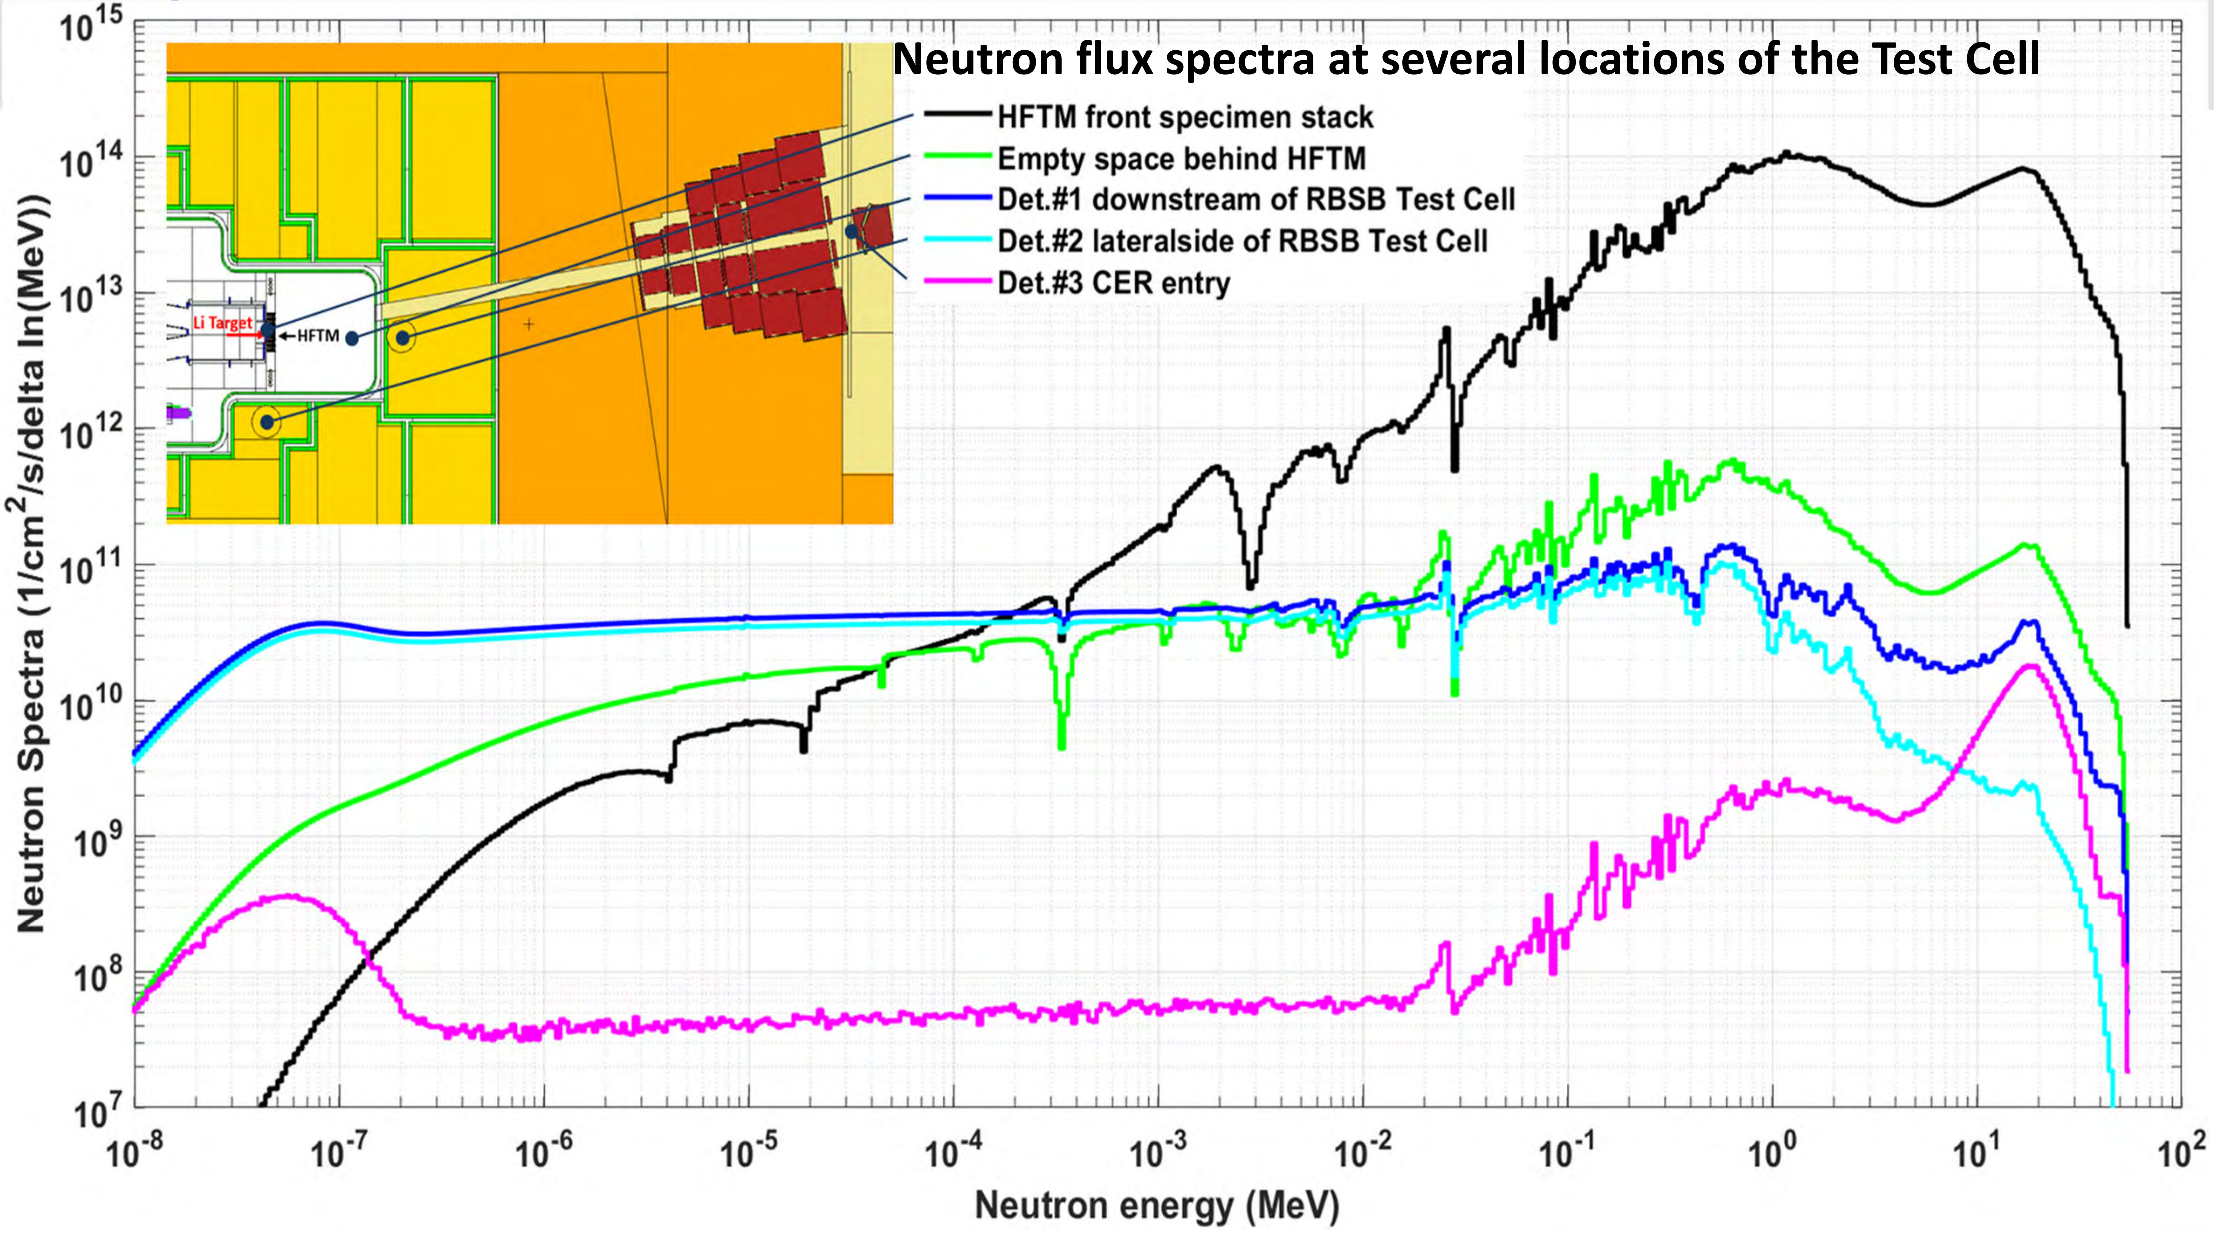

Neutron flux spectra at several locations of the Test Cell
HFTM front specimen stack
Empty space behind HFTM
Det.#1 downstream of RBSB Test Cell
Det.#2 lateralside of RBSB Test Cell
Det.#3 CER entry
Li Target
HFTM
Neutron Spectra (1/cm²/s/delta ln(MeV))
Neutron energy (MeV)

MCHIFI: Development of the On-The-Fly (OTF) MC variance reduction technique

- **OTF-GVR**: On-The-Fly Global Variance Reduction:
  - Weight windows mesh (WWM) is a common method used for MC shielding calculation.
  - OTF performs "on-the-fly" iterations to get a global flux map and a weight-window mesh (WWM).
  - OTF uses novel dynamic WW upper bound method to solve the neutron streaming and "long-history" particles
  - Compared with ADVANTG, the Figure-of-Merit in OTF is raised by a **factor of 20**

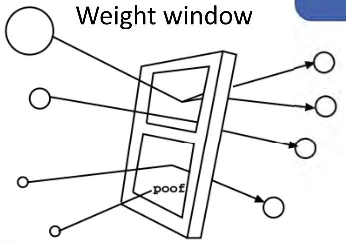

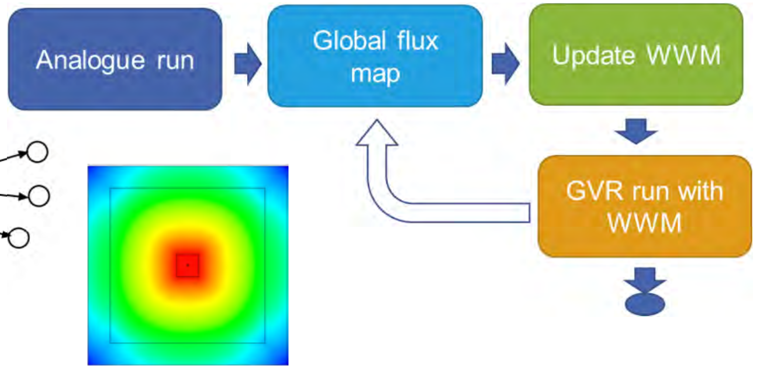

On-the-fly Global weight window mesh generation

OTF-GVR:

$$n(\vec{r}) \approx m(\vec{r})\bar{w}(\vec{r})$$

$$\bar{w}(\vec{r}) = \phi(\vec{r})/\max(\phi(\vec{r})) \Rightarrow \bar{w}(\vec{r}) = c \times \phi(\vec{r})/\max(\phi(\vec{r}))$$

Definition of "c" to avoid "long-history" by limiting the n, p particles splitting in the OTF run in Ref. [Yu Zheng, Yuefeng Qiu, et al., "An improved on-the-fly global variance reduction technique by automatically updating weight window values for Monte Carlo shielding calculation", Fusion Eng. Des. 147 (2019) 111238, https://doi.org/10.1016/j.fusengdes.2019.06.011]

Comparison of the MC statistical precision of the neutronics results indicated by relative error maps: MCNP vs. OTF

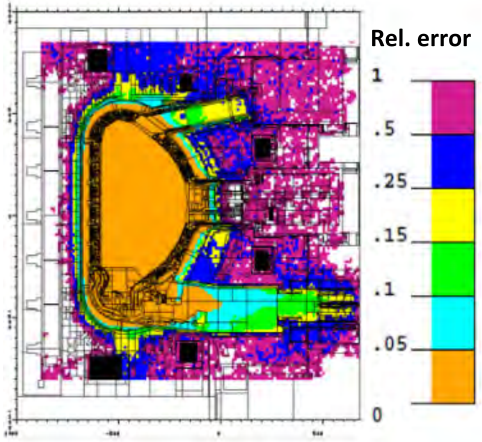

Analogue MCNP

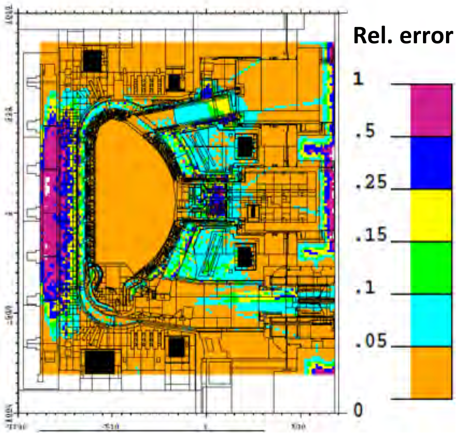

MCNP - ADVANTG (ORNL) WWM

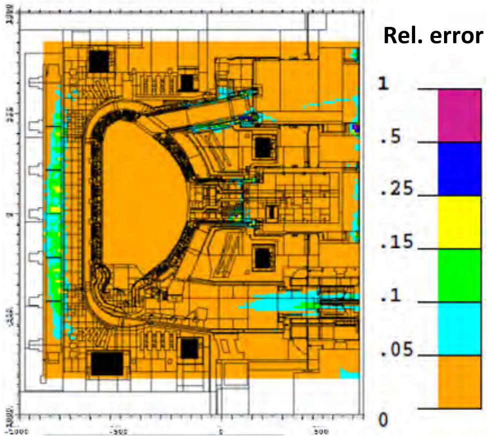

OTF-GVR

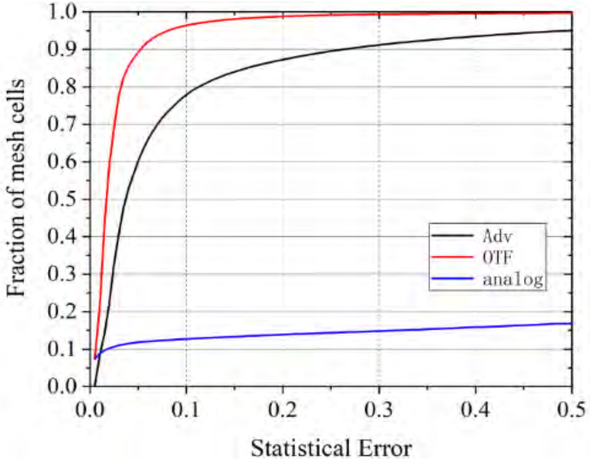

Percentage of mesh cells and rel. error

Ref.: [Yu Zheng et al 2022 Nucl. Fusion 62 086036, https://doi.org/10.1088/1741-4326/ac75fc]

Deuteron beam energy deposition in the Li jet at the TA d-Li footprint area

- D+ ion beam stops in the lithium jet delivering a total power of 5 MW on a volume of 5×20×2.5 cm^3, with d-Li footprint area of 5x20 cm^2.
- Deuterons lose their energy in Li by interactions with Li electron clouds and nuclei – all the processes have been taken into account in the MCNP6 energy deposition calculations with the TMESH card.
- For calculation of deuteron beam energy deposition in Li at the d-Li footprint area, **transport of neutrons, photons, deuterons, and protons** – 4 particles have been transported with the MCNP6 mode **n p d h**

Horizontal cut of the MCNP6 geometry at d-Li footprint

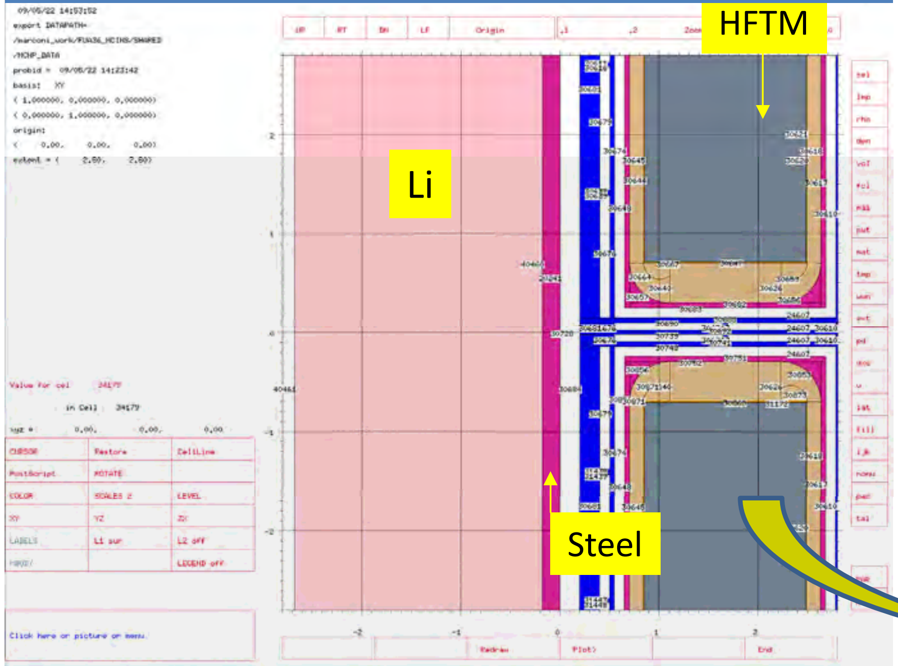

Energy deposition, W/cc

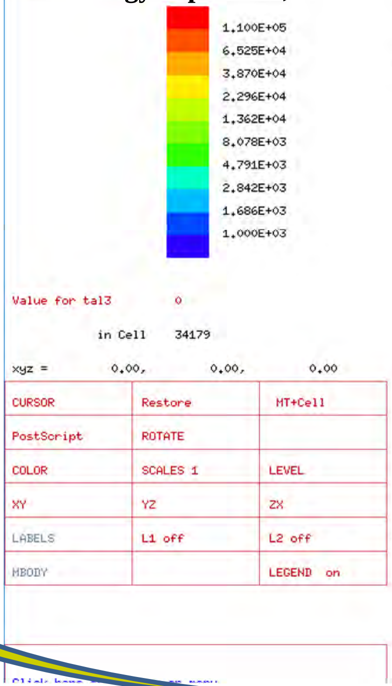

MCNP6 horizontal cut of the D+ beam energy deposition at the d-Li footprint area with heat peak of 110 kW/cc

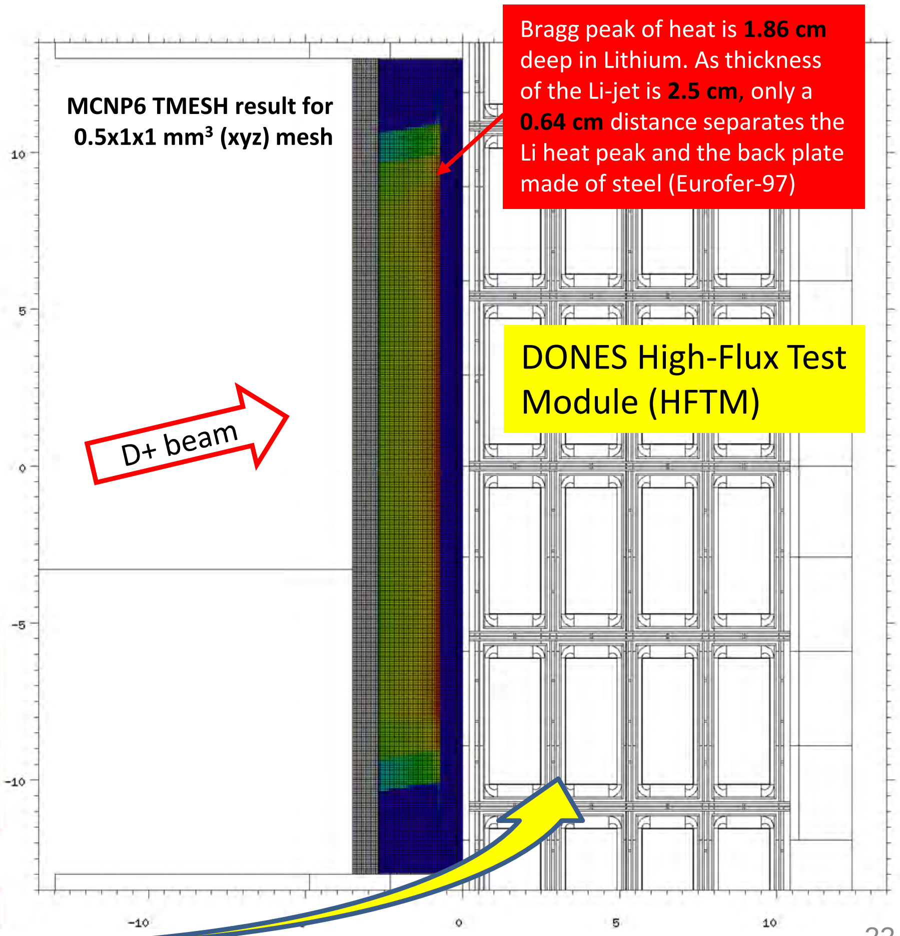

D+ ion beam energy deposition in Li target with Li(d, xn) neutron source in IFMIF-DONES

D+ beam profile (IFMIF/EVEDA)

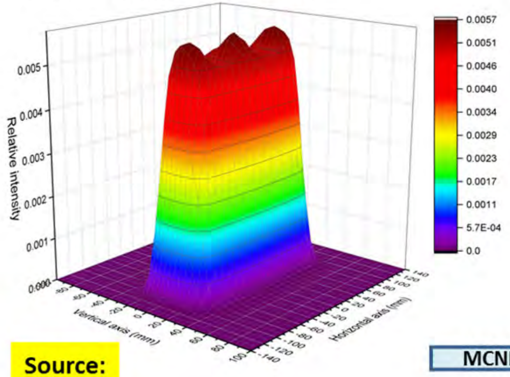

Source:

MCNP6 transport

Result:

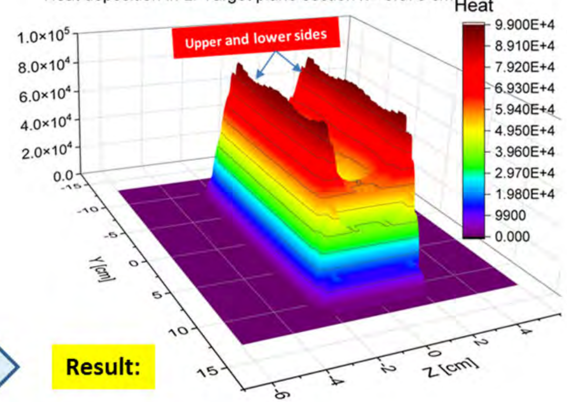

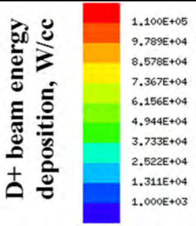

Heat at peak is 110 kW/cc
Integral D+ heat 4.8 MW

Notice: Bragg peak at Lithium thickness of **1.86 cm** corresponds to **-0.84 cm** coordinate in the MCNP model geometry with X=0 at the TA back plate and X=-2.7cm is the front point of Li at Z=0 central plane: **1.86cm - 2.7cm = -0.84cm**

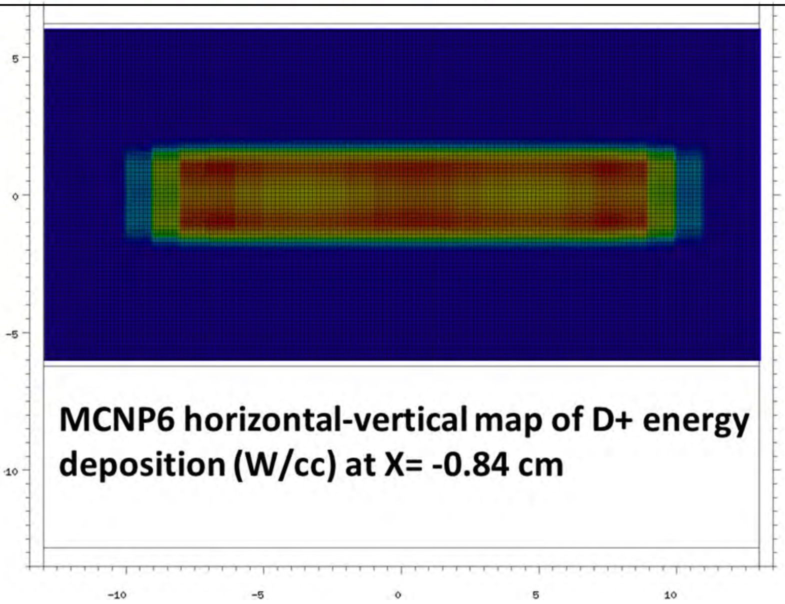

MCNP6 horizontal-vertical map of D+ energy deposition (W/cc) at X= -0.84 cm

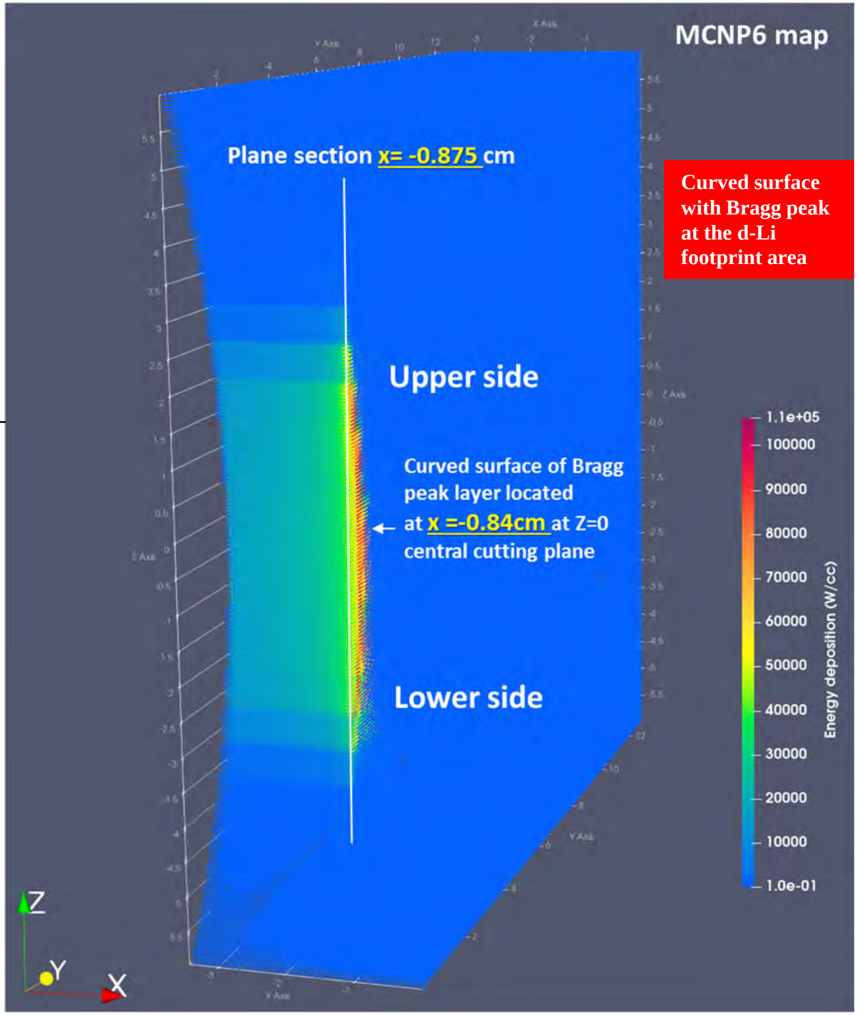

MCNP modeling of the d-Li source Target Assembly (TA) in DONES

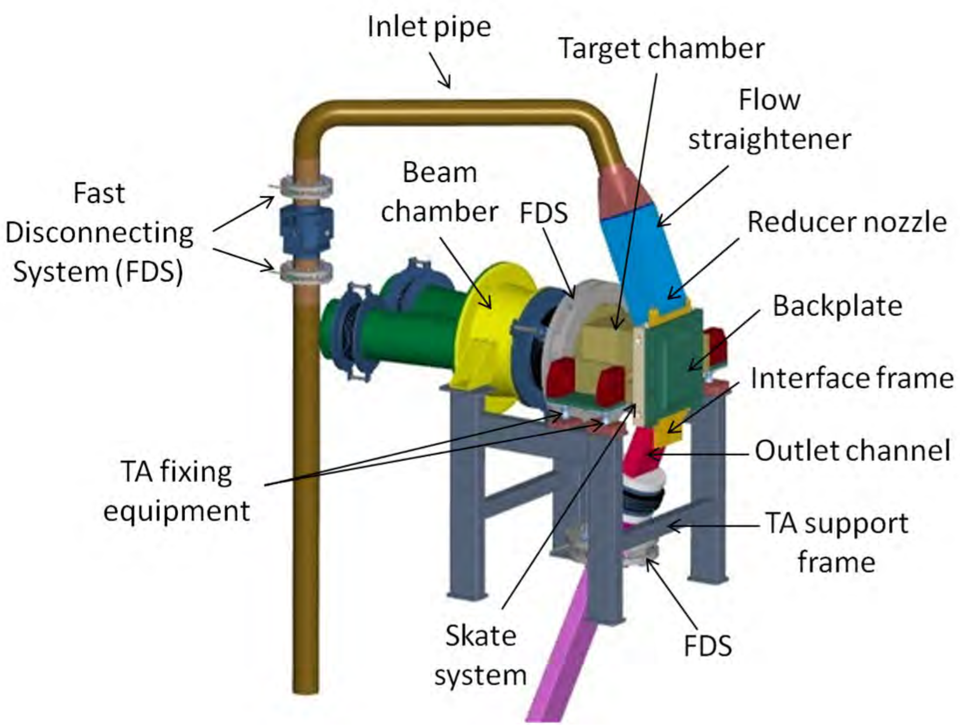

DONES Target Assembly (TA) components

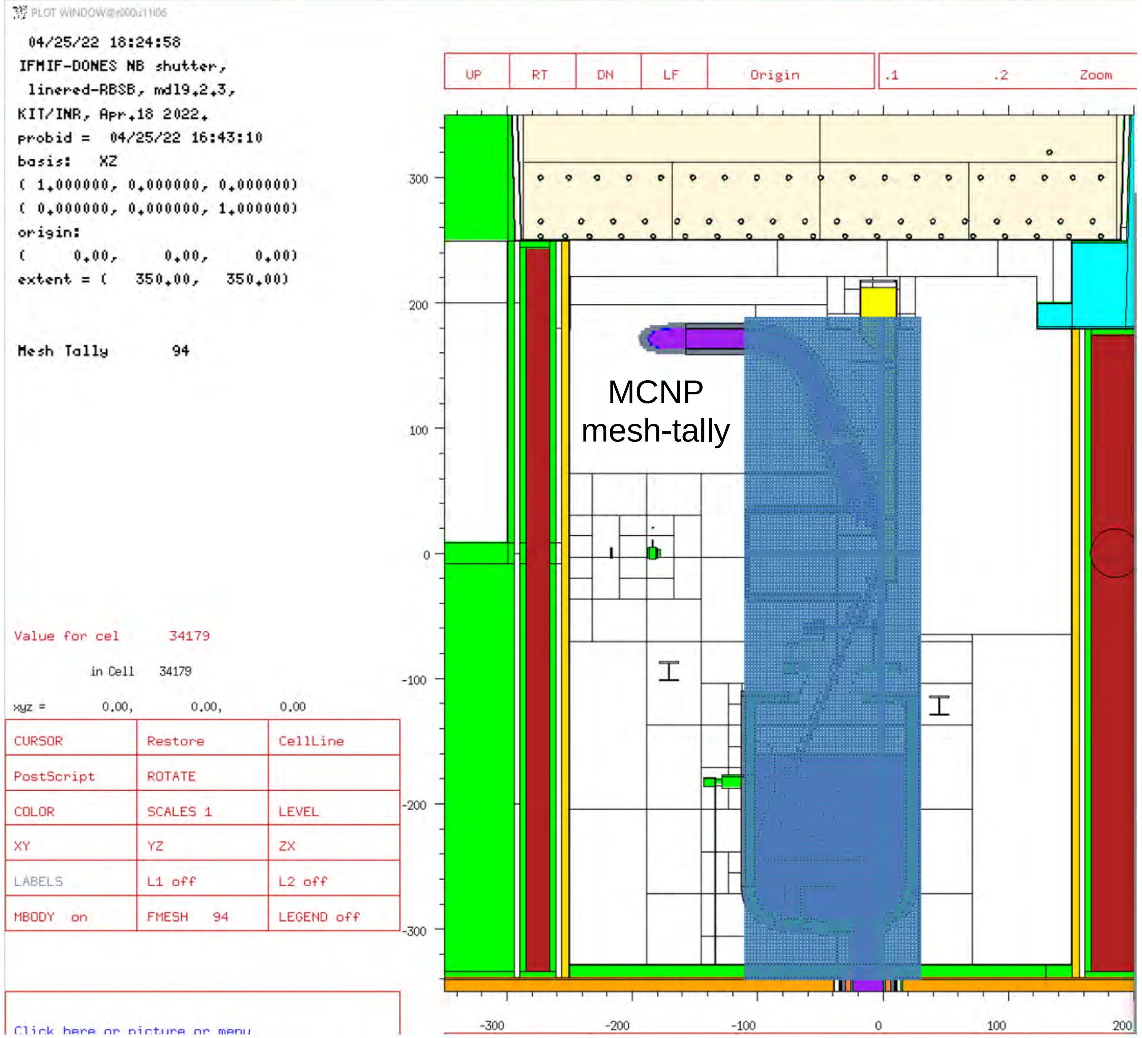

MCNP model vertical cut of the DONES TA covered with mesh-tally

Nuclear heat density (W/cc) in the TA materials of the MCNP model

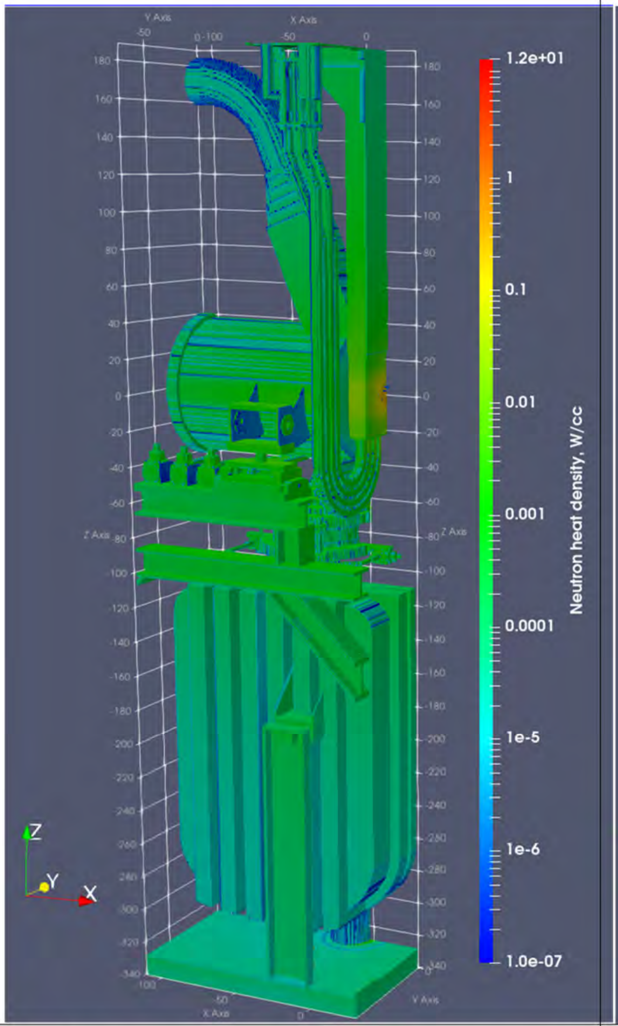

Fig. 1. Neutron heat density (W/cc) in actual materials of the MCNP model – look from the outside.

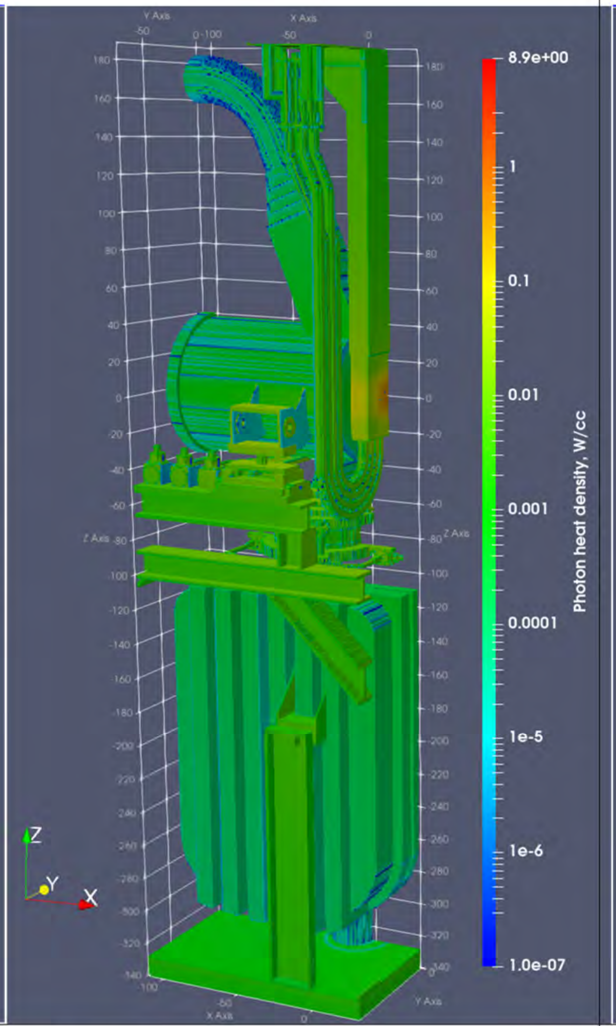

Fig. 2. Photon heat density (W/cc) in actual materials of the MCNP model – look from the outside.

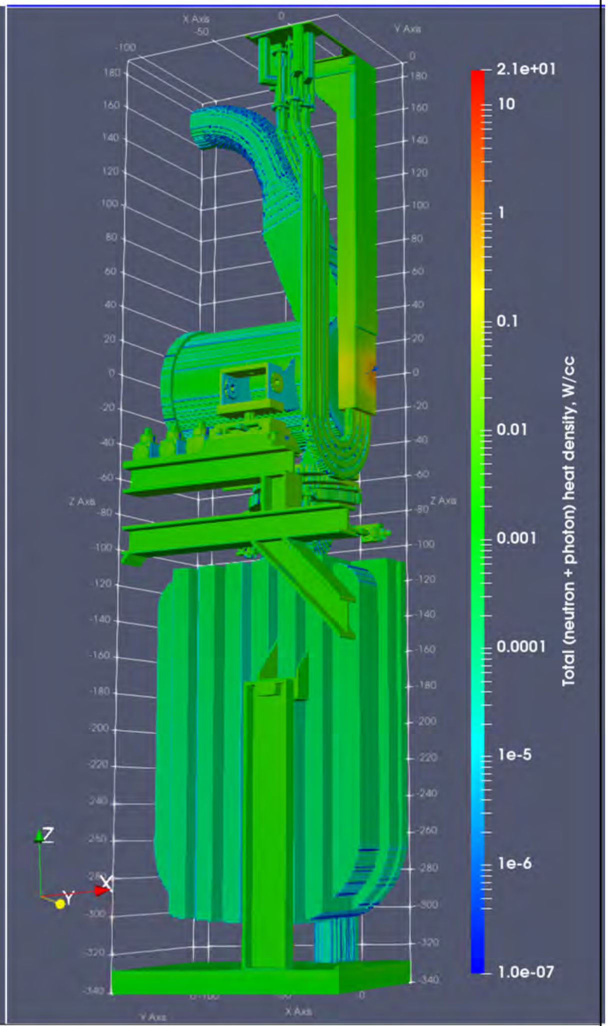

Fig. 3. Total (neutron + photon) heat density (W/cc) in actual materials of the model – look from the outside.

TA materials:

Steel SS316L material density 7.93 g/cc

EUROFER steel with density 7.87 g/cc

Lithium (Li) with impurities, its density is 0.512 g/cc.

Heating in Li jet at the area of deuteron footprint requires inclusion of the heat contributions of charged particles.

Power balance for the D+ beam energy released in Test Cell (TC) and internal components

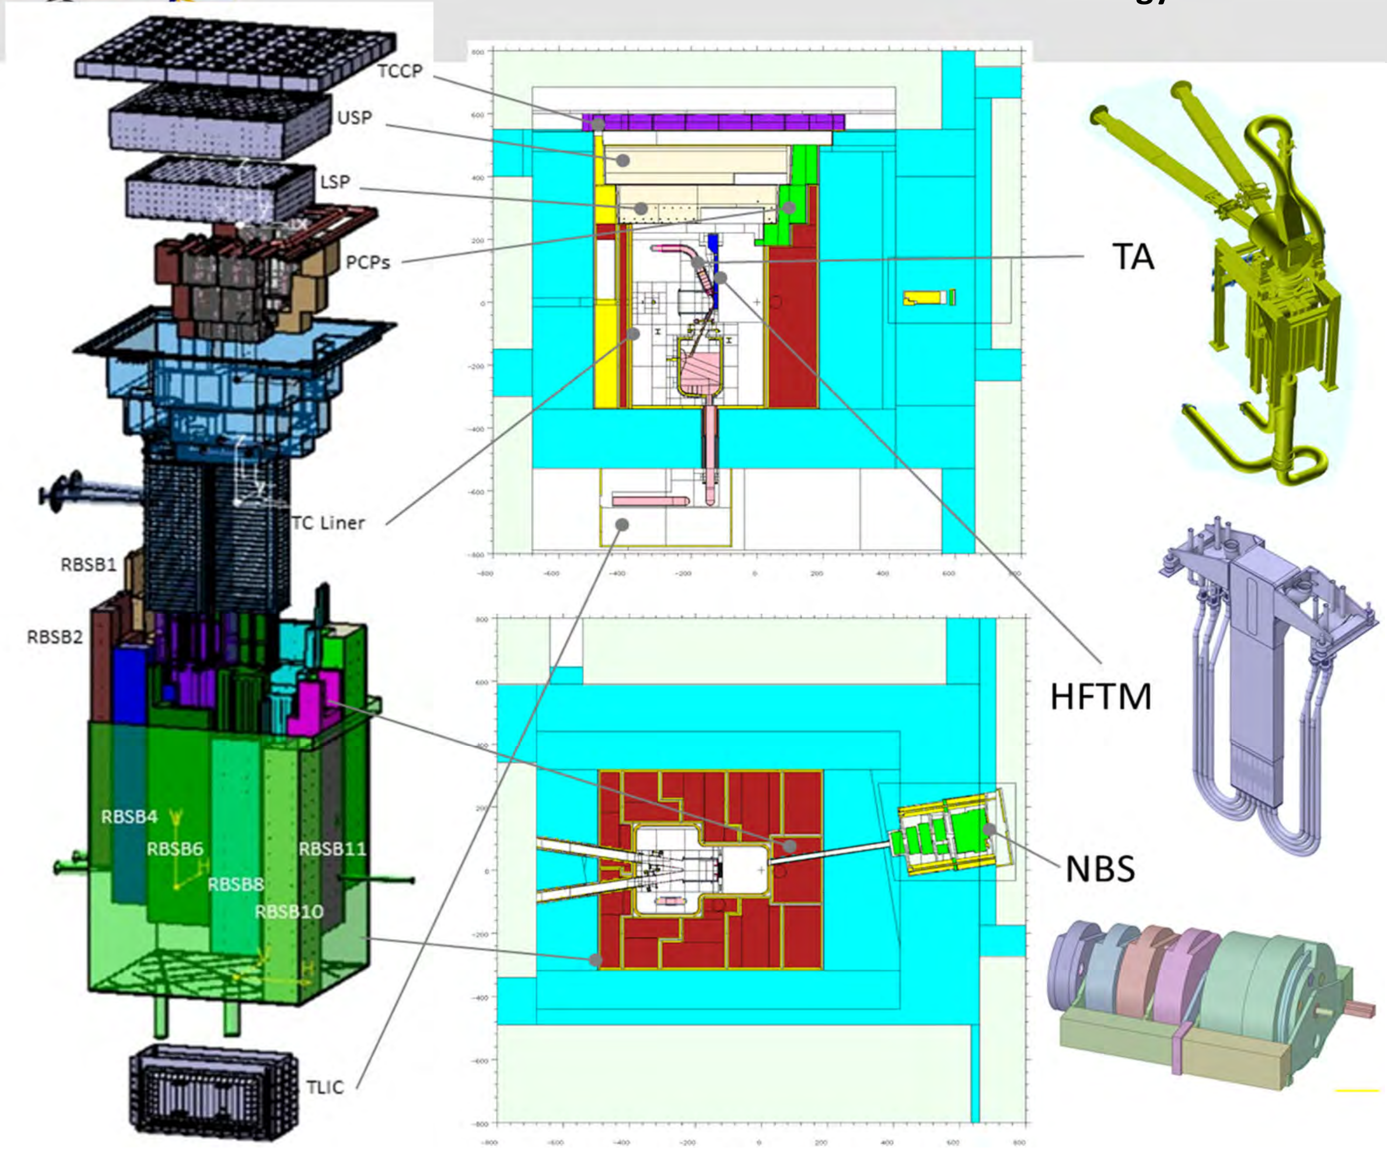

Type of heat power	Heat, kW
Input heat power of D+ beam delivered by the IFMIF-DONES one accelerator beam current of I=125 mA	**5000**
D+ heat released in Li Target	**4858.8**
Neutron + Photon heat released in the TC components (numbered and displayed in the next slide):	
1) TC liner	**15.2**
2) Removable Biological Shielding Blocks	**77.3**
3) Bucket liner	**0.2**
4) Bucket	**1.6**
5) Piping and Cabling Plugs (PCP)	**2.6**
6) Lower Shielding Plug (LSP)	**9.9**
7) Upper Shielding Plug (USP)	**0.01**
The sum of 7 TC components:	**107**
Target Assembly (TA) structural parts	**17.3**
High Flux Test Module (HFTM)	**16.9**
Integral neutron and photon heat in all considered TC components	**~141.2**
Output: total integral heat released by D+, neutrons, and photons:	**4858.8+141.2 = 5000**

Integral nuclear (neutron + photon) heating in the Test Cell (TC) components

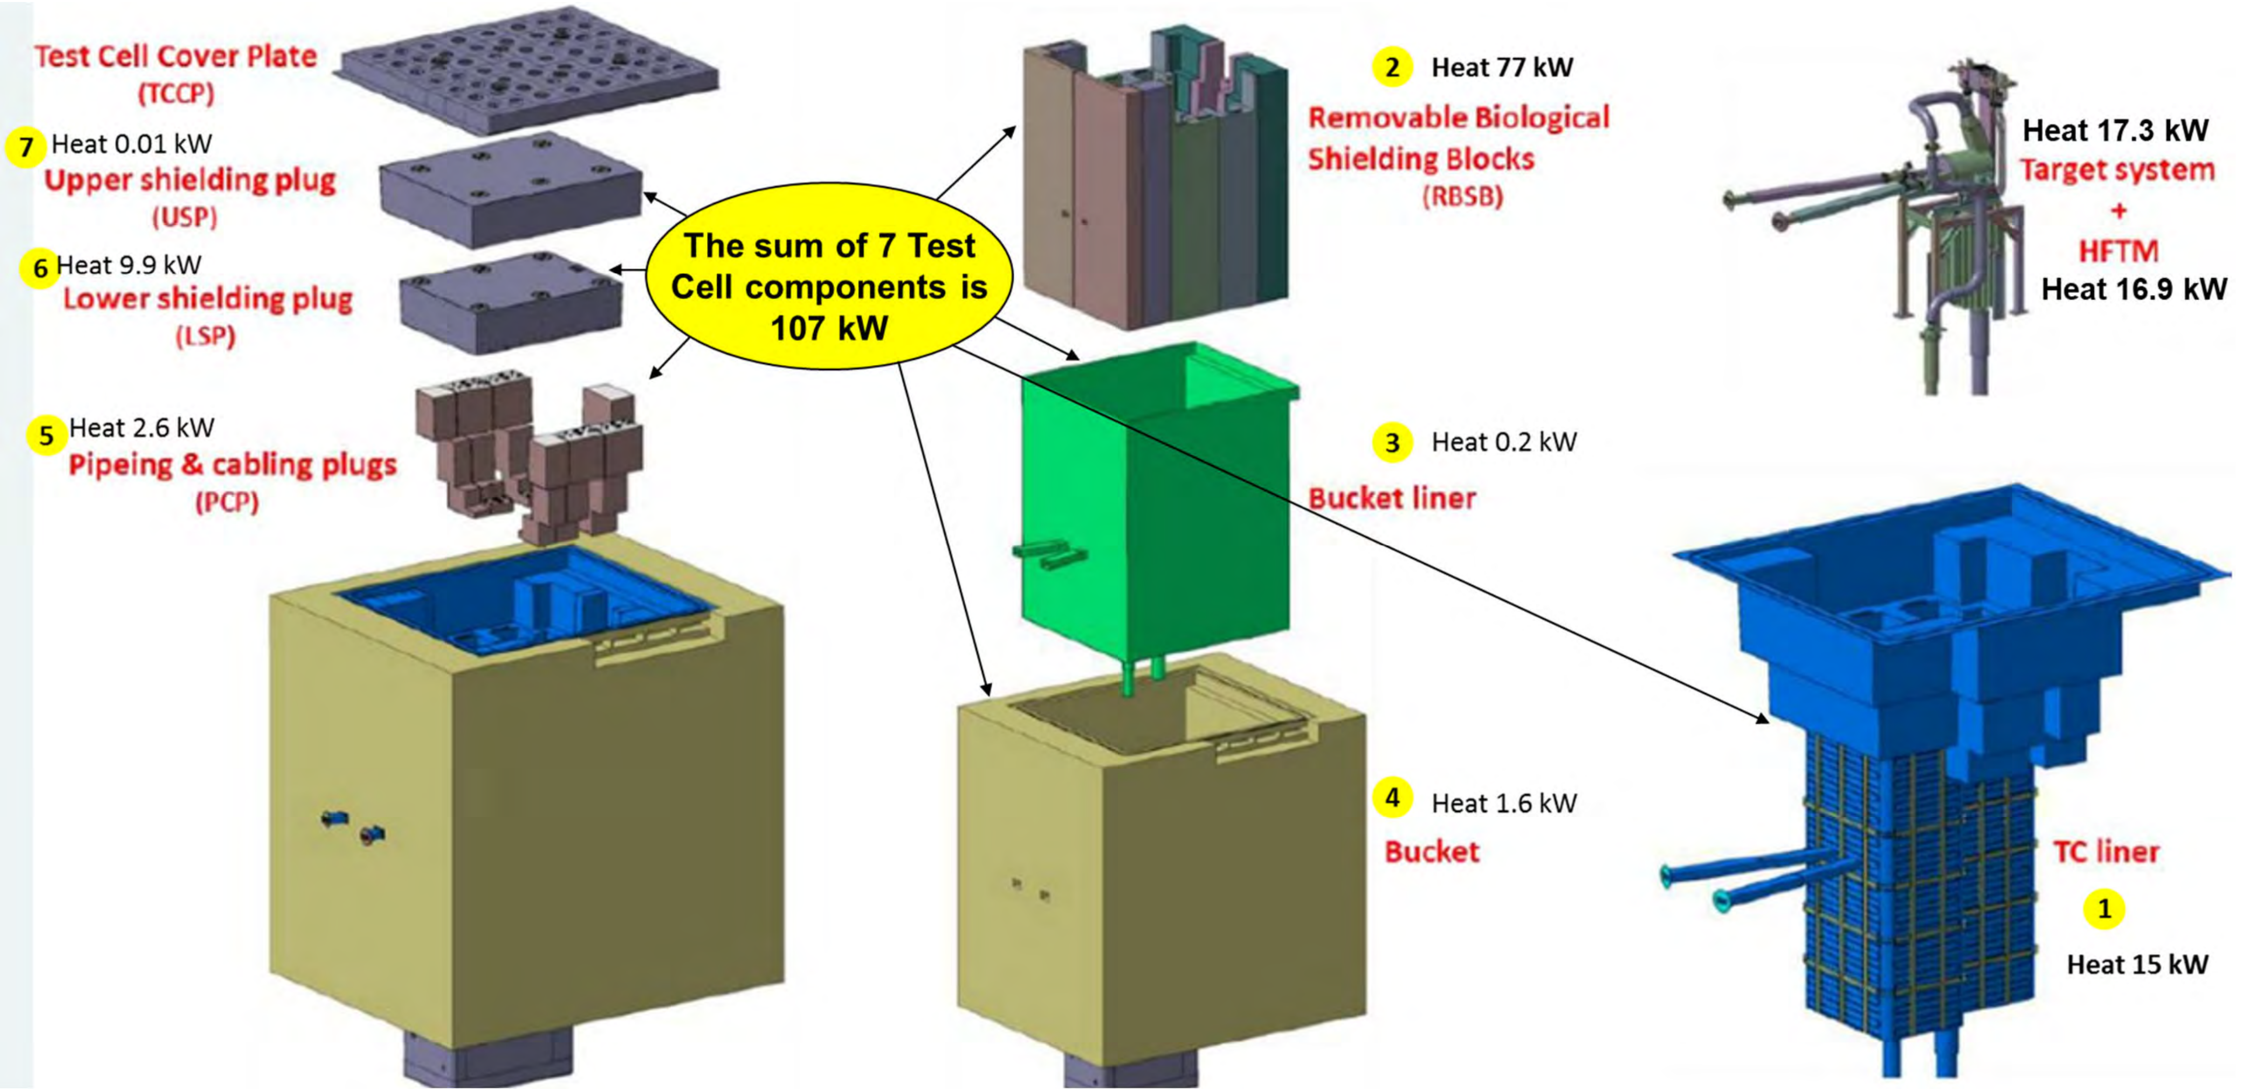

Part III

How to use the GPU partitions of Leonardo Booster and Pitagora for the MC radiation (neutrons, deuterons, protons, photons, electrons) transport?

- MCNP follows the particle history track from production to absorption or geometry escape.
- A cascades of secondary particles are generated stochastically during the particle interaction.
- MCNP computes all secondaries recursively until all of them terminate (absorption or escape), while all primary particles are kept in the CPU memory stack. The MCNP is banking primary particles and processing secondary particles recursively until all the multiplied branched processes of the initial particles are ended. The CPU has the LIFO (Last-In, First-Out) stack for recursive particle tracking. The GPU has a very small stack (<10KB vs. CPU's MBs). The GPU stack is emulated with high-latency global memory. **MCNP's recursion depth** exceeds GPU capabilities. The most promising example of the **GPU MC codes (OpenMC)** eliminates recursion entirely. Instead of history-based sampling, OpenMC applies event-based parallelism.

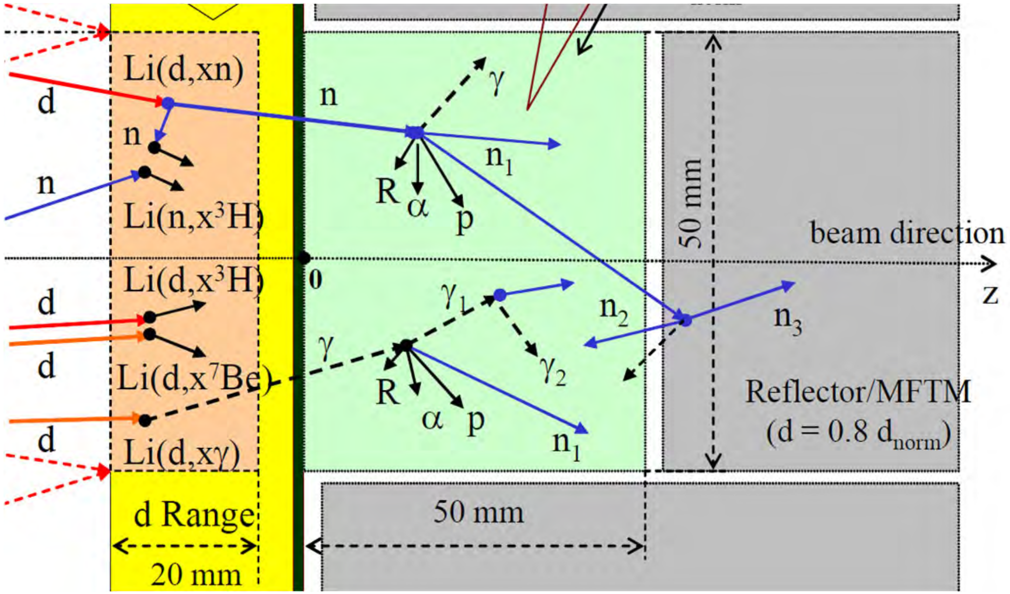

Feature	CPU Stack	GPU "Stack"
Hardware Acceleration	Dedicated SP/BP registers	Emulated via global memory
Access Latency	1–10 cycles (L1 cache)	300–500 cycles
Typical Size per Thread	~8MB (Linux default)	1–2KB (CUDA default)
Deep Recursion Support	Yes (1000+ levels)	No (practically <10 levels)
Context Switch Cost	Cheap (register save/restore)	Extremely expensive

First Step: Event-Based Parallelism

<u>OpenMC Event Kernels</u>

- Particle Initialization
- Calculate Cross Sections (Fuel)
- Calculate Cross Sections (non-Fuel)
- Advance Particle
- Cross Surface
- Collision
- Particle Death

Example of an event kernel launch via OpenMP

```
#pragma omp target teams distribute parallel for
for (int i = 0; i < n_particles; i++) {
 int particle_idx = calculate_fuel_xs_queue[i].idx;
 Particle& p = particles[particle_idx];
 p.event_calculate_xs_execute();
}
```

- All main **event kernels** in OpenMC have been offloaded to device
- Some kernels are very large:
  - Deep call stacks
  - Functions scattered over many files
  - O(1000's) lines of code per kernel

Ref. [1] John Tramm (Argonne National Laboratory), “OpenMC GPU Developments”, OpenMC Application to Tokamak Neutronics Analysis Meeting on May 28-30, 2025

Early 2021: First correct results achieved with naïve GPU OpenMC port

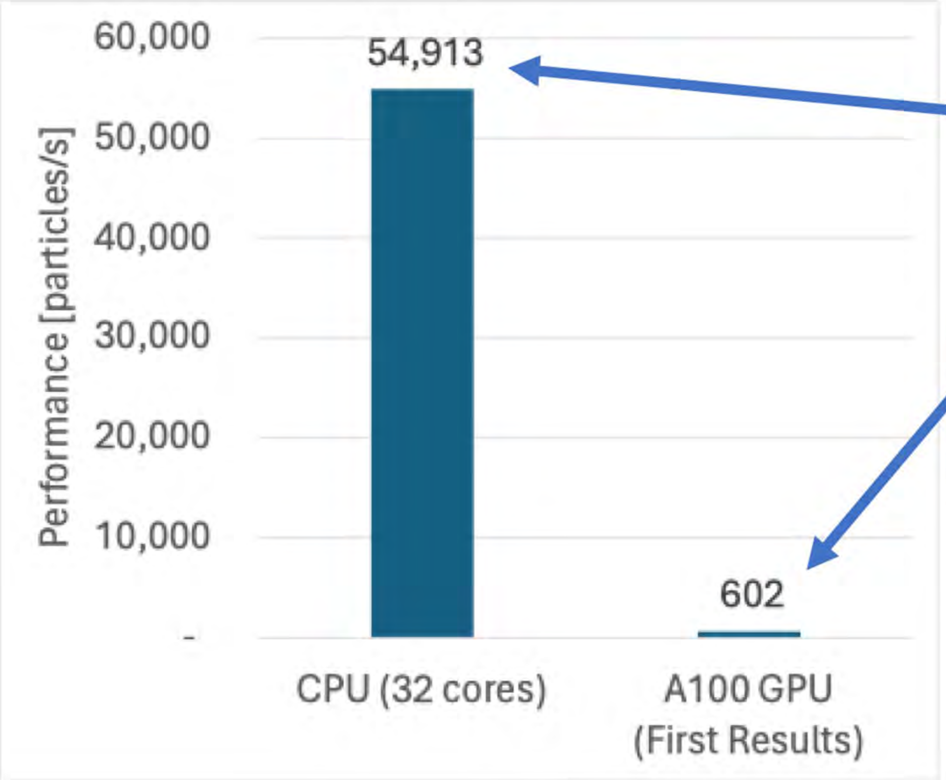

Starting point was not good! A100 GPU slower than a single CPU core!

How to move forward?

1. Interface/co-design with OpenMP compiler teams
2. Implement known GPU algorithmic optimizations (from Shift, PRAGMA)
3. Develop additional new algorithmic & implementation optimizations

Ref. [1] John Tramm (Argonne National Laboratory), "OpenMC GPU Developments", OpenMC Application to Tokamak Neutronics Analysis Meeting on May 28-30, 2025

Fall 2021: LLVM Compiler Optimization Successes

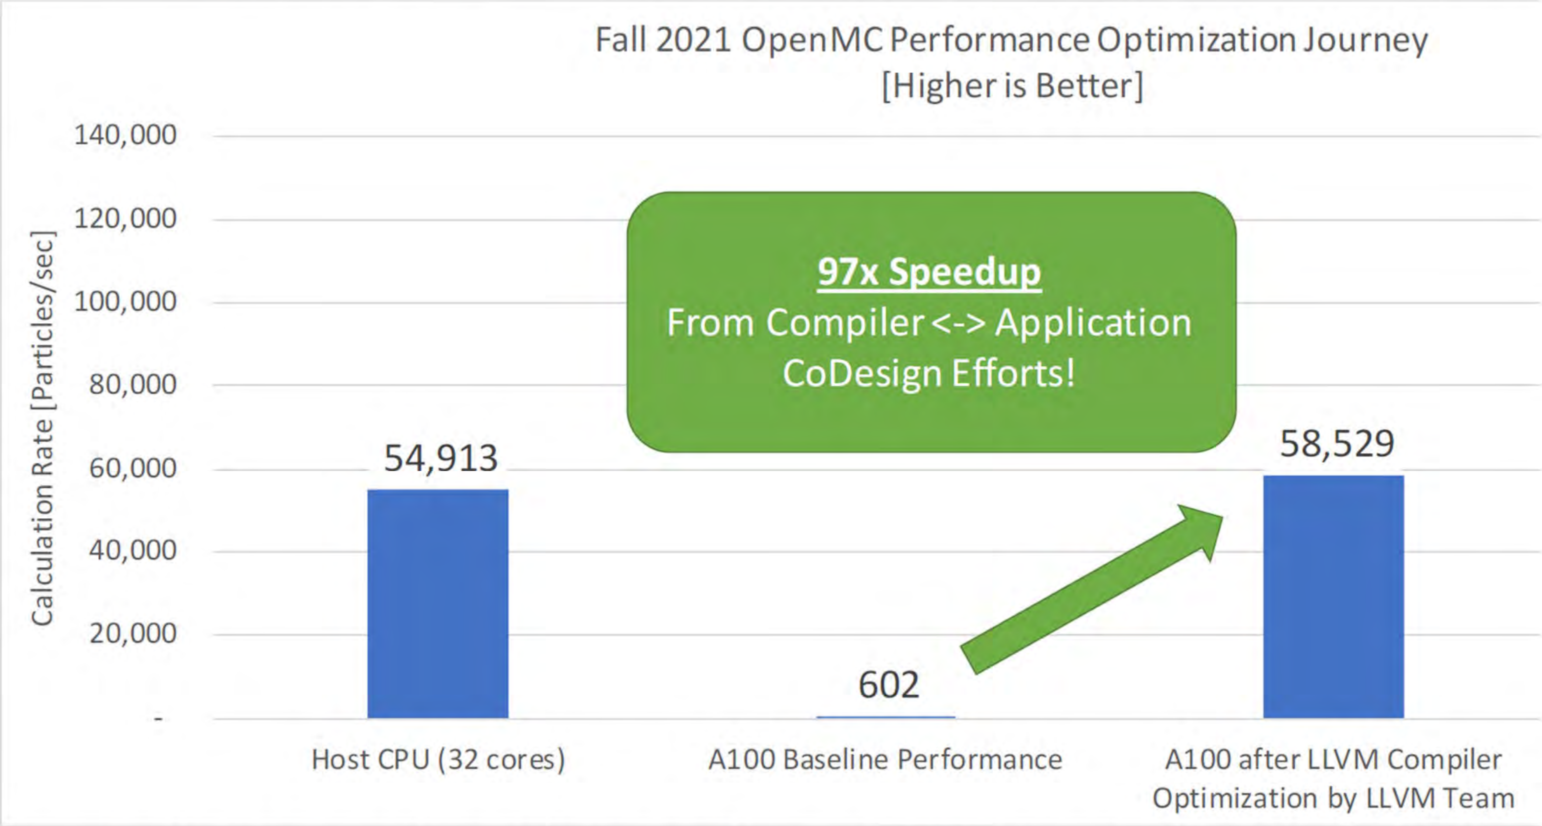

- Close engagement with LLVM team (mainly **Johannes Doerfert @ LLNL**) resulted in some problems being identified in LLVM Clang.
- Issues were fixable via small OpenMP runtime optimizations and hotfixes.

Ref. [1] John Tramm (Argonne National Laboratory), "OpenMC GPU Developments", OpenMC Application to Tokamak Neutronics Analysis Meeting on May 28-30, 2025

Event-Based Algorithm, Sorting, and MicroXS Cache Removal Massively Improve GPU Performance

But Reduce CPU Performance

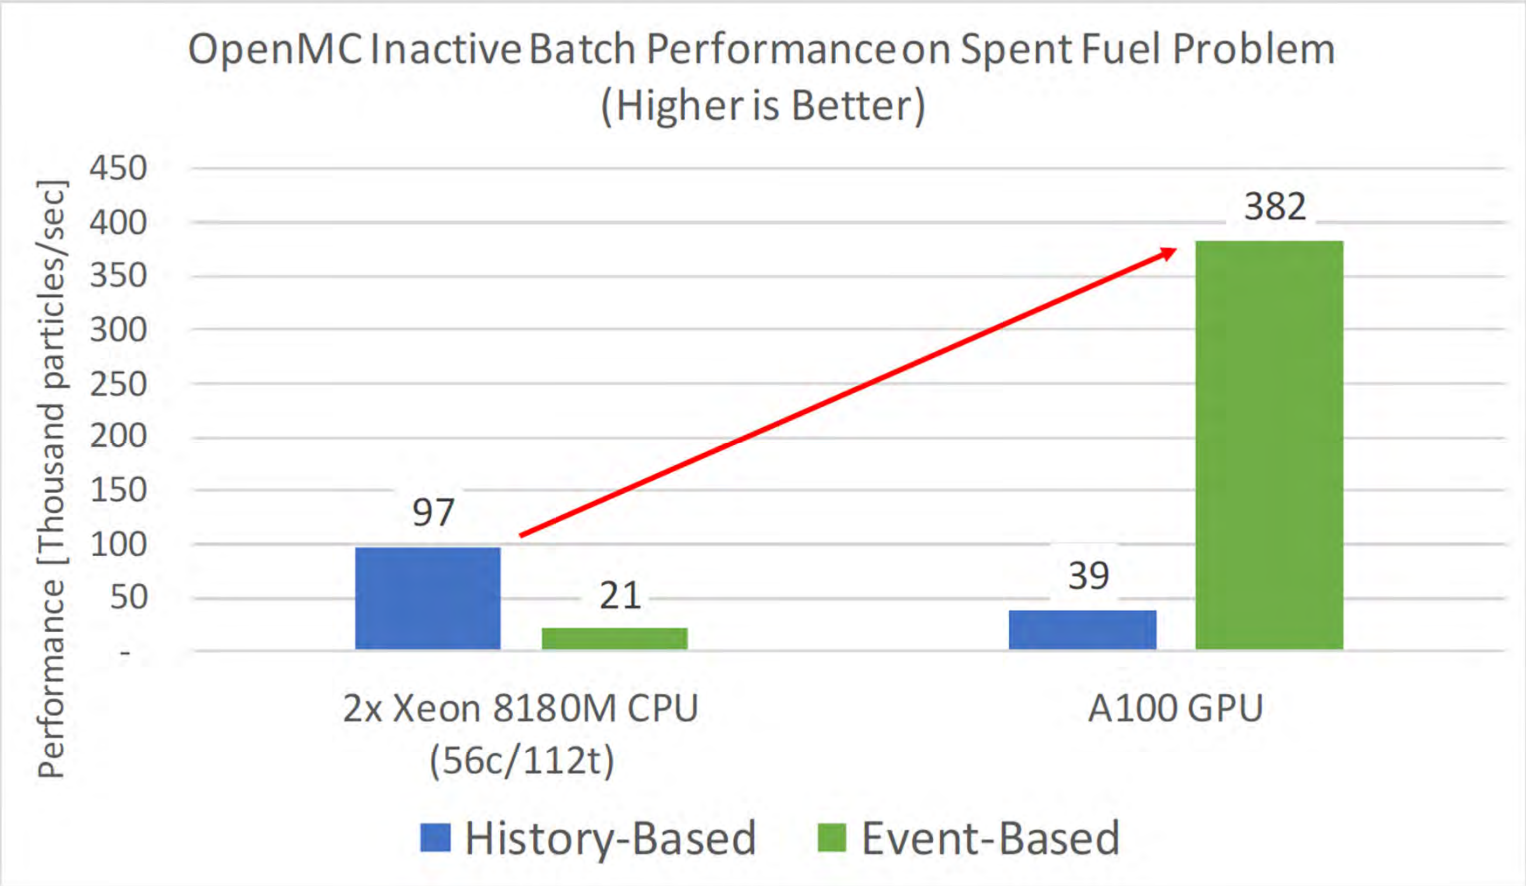

Great performance with event-based transport
+
removal of micro XS cache
+
energy sorting

Ref. [1] John Tramm (Argonne National Laboratory), "OpenMC GPU Developments", OpenMC Application to Tokamak Neutronics Analysis Meeting on May 28-30, 2025,

Argonne NATIONAL LABORATORY

Improvements Shown at SNA + MC 2024

- Addition of tally support
- GPU stream optimization
- Further code fine-tuning
- On-device particle sorting via vendor library
- Support for fixed source simulation problems
- **Scaling ability on supercomputing architectures from all three GPU vendors**

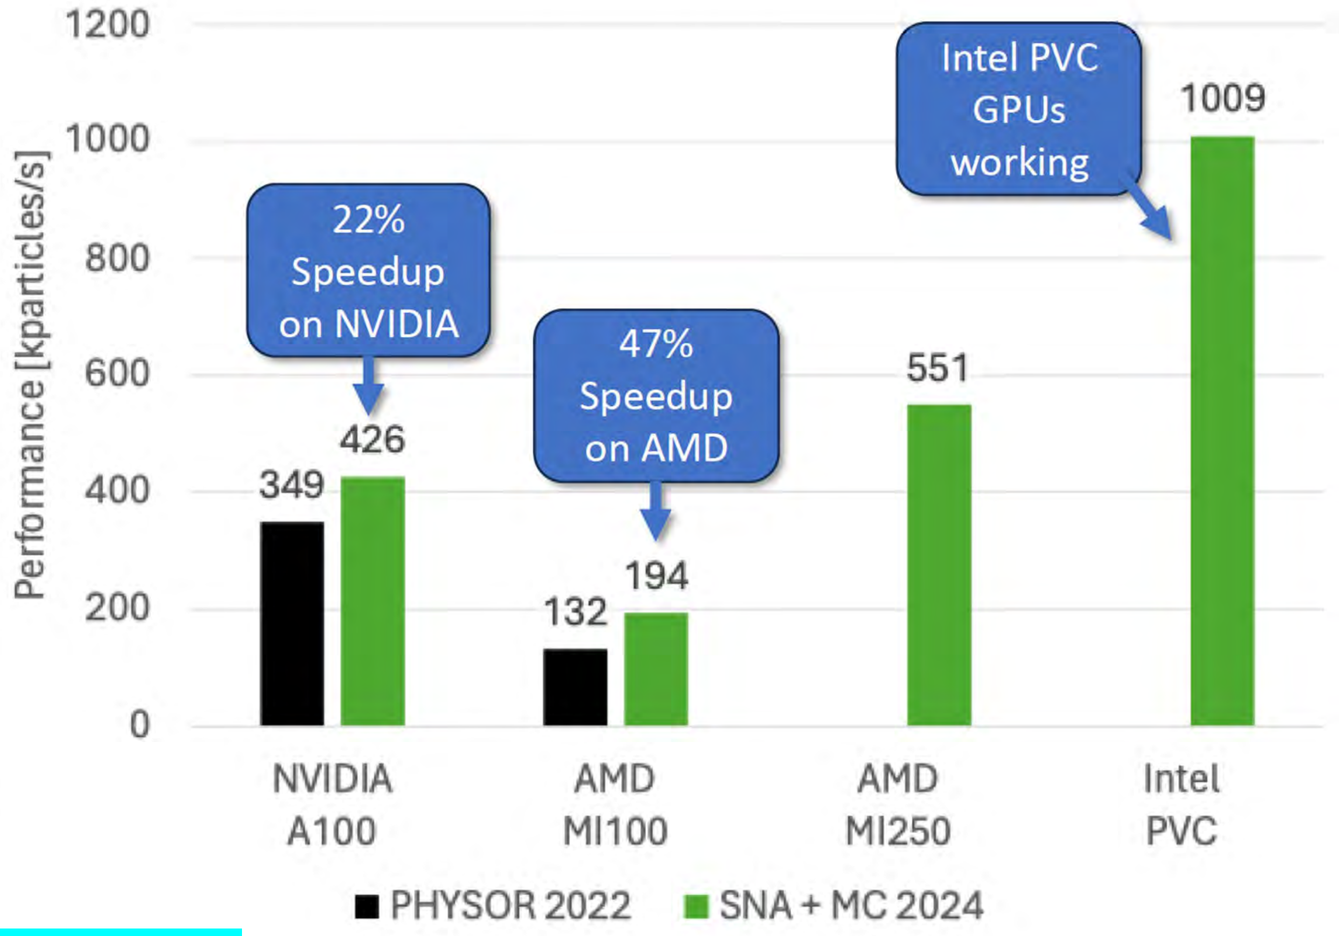

Ref. [1] John Tramm (Argonne National Laboratory), “OpenMC GPU Developments”, OpenMC Application to Tokamak Neutronics Analysis Meeting on May 28-30, 2025

Conclusions-1

- **Serving most demanded Monte Carlo radiation transport computations for fusion large facilities, the MCHIFI (Monte Carlo High Fidelity) HPC** project has been established in 2012 to use the IFERC-CSC Helios supercomputer in the framework of the F4E Broader Approach (BA) for ITER neutronics tasks.
- Since 2016, MCHIFI project was deployed on the computational resources of the Marconi-Fusion HPC facility offered by EUROfusion and operated by ENEA/CINECA. At this moment, MCHIFI is running on 7th cycle of Marconi-Fusion HPC. In our application for the 9th cycle, we have justified the number of core-hours requested to fulfill not only the IFMIF-DONES shielding tasks, but also numerous Shutdown Dose Rate (SDR) calculations for the JET NEXP SDR experiment of the C38 2019-2020 DD and the C41 2021 DT2 campaigns of JET, and using the data of JET DT3 campaign to be available at the end of 2023.
- A large number of the Marconi-Fusion HPC nodes are needed in the for the IFMIF-DONES large-scale complicated models to run the Monte Carlo (MC) radiation transport parallel computations with the MCNP and McDeLicious codes.
- The methodology improvement is demonstrated in the recent development of the On-The-Fly (OTF) modification of the MCNP code. The OTF Global Variance Reduction (OTF-GVR) is the state-of-the-art code for the IFMIF-DONES radiation shielding tasks characterized by neutrons' deep penetration. We have used OTF-GVR for radiation transport through the 6.4 m shielding between the IFMIF-DONES Test Cell and its Complementary Experiments Room (CER).
- **Developed at KIT Monte Carlo radiation transport CAD-based methodology** can reproduce the d-Li neutron & photon source at the Li target and extremely strong radiation attenuation in heterogeneous IFMIF-DONES geometry:
  - Neutron flux attenuation by 18 orders of magnitude (from 2e14 to 2e-4 n/cm^2/s inside Tritium room);

Conclusions-2

- **The interactions of deuterons with lithium target** for the energies relevant to fusion applications, particularly Ed=40 MeV in IFMIF-DONES facility, are most accurately described with the McDeLicious code in its actual version McDeLicious-17, as an extension of the MCNP6.1 Monte Carlo radiation transport code. The McDeLicious code has been validated & verified in experimental and computational benchmarks.
- Using the D+ beam settings, McDeLicious samples neutrons and photons using evaluated d+ $^{6,7}Li$ data.

- The simple model of D+ interactions with cylindrical solid Li inside Aluminum capsule allows to investigate the D+ flux attenuation, track length, Ed attenuation, and D+ energy deposition. This work presented simple model with two settings of the D+ sources:
  - V1: Isotropic uniformly distributed disk
  - V2: Monodirectional directed source defined at a disk

- The (d-Li) reactions defined in McDeLicious-17 have been studied for the IFMIF-DONES facility. The beam of deuteron ions accelerated up to 40 MeV with current of 125 mA impinges the liquid Li target delivering 5 MW power. The presented results include distributions of D+ energy deposition, neutron and photon fluxes and heating.

- The integral heating calculations in IFMIF-DONES Test Cell (TC) components reveals that D+ energy deposition in liquid Li at thin Bragg peak with a D+ beam footprint area of 20x5 cm^2 contributes 97% of total heating in the whole Test Cell volume. The 5 MW heat power of D+ beam delivered by the IFMIF-DONES is released by 97% in liquid lithium.

Outcomes and Future Activities

- **At present, only the MCNP6 code with OpenMP/MPI parallelization on the CPU partitions can be used for neutronics applications supporting the design development of the main fusion facilities such as ITER, IFMIF-DONES, EU DEMO, VNS, and JET. The management of the listed projects accepts only the neutronics results computed with the reference MCNP code. All other Monte Carlo neutronics codes are considered complementary. Other neutronics codes can be used for scoping and parametric analyses.**
- **MCNP's architecture is the result of decades of evolution in CPU-based Monte Carlo simulation. In the MCNP parallel jobs on CPUs, the parallelization is performed on the level of particle histories. Porting it to leverage GPUs would require extensive rewriting and restructuring of its core algorithms to better align with the vectorized and parallel processing models that GPUs demand.**
- **The event-based radiation transport implemented in the OpenMC code opens perspectives on application of the GPU partitions of Leonardo Booster and Pitagora for the Monte Carlo radiation transport. In the framework of the EUROfusion MCHIFI project, we are planning to test the potential of the OpenMC code parallelization on GPUs.**
- **However, the neutronics results produced by the OpenMC code can not be used alone in the major projects (ITER, IFMIF-DONES, EU DEMO, VNS, and JET) without the reference MCNP results. The cross-checks with MCNP are required.**
- The OpenMC code is under development by the fusion neutronics community. Jin Hun Park (KIT, MCHIFI project) in "Neutronics Benchmark of European DEMO and VNS using MCNP and OpenMC for the Eurofusion Programme", OpenMC Application to Tokamak Neutronics Analysis Meeting on May 28-30, 2025, Argonne National Laboratory (ANL) in collaboration with EUROfusion proposed the following areas of work on further OpenMC development:
- MCNP-to-OpenMC source converter
- Lattice structure flexibilities - MCNP-to-OpenMC converter doesn't support
- Boundary (only reflective boundary??) issue with source
- MCNP and OpenMC TBR results are different
- WW calculation performance
- photon transport with weight window (for ADVANTG, OpenMC MAGIC + on-the-fly) performance
- Surface tally for NWL comparison
- Shutdown dose rate methodology on main openmc distribution
- CAD & unstructured mesh base model benchmark studies
- Normalized WW, materials definition in the OpenMC models.