PROCEEDINGS OF SCIENCE

COSMICA: a GPU-optimized code for solar modulation studies

L. Bacciu,[d] **G. Cavallotto,**[b,a] **S. Della Torre,**[a,b,*] **M. Grazioso,**[d,a] **M. Gervasi,**[a,c] **G. La Vacca ,**[a,c] **S. Rossi**[d] **and M. S. Nobile**[d]

[a] *INFN Sez. Milano-Bicocca, Piazza della Scienza 3, Milano, Italy*
[b] *ICSC - Centro Nazionale di Ricerca in HPC, Big Data and Quantum Computing, Via Magnanelli, 2, Casalecchio di Reno (BO), Italy*
[c] *Department of Physics "Giuseppe Occhialini",University of Milano-Bicocca, Piazza della Scienza 3, Milano, Italy*
[d] *Department of Environmental Sciences, Informatics and Statistics, Ca' Foscari University of Venice, Via Torino 155, Mestre (VE), Italy*
E-mail: stefano.dellatorre@mib.infn.it, giovanni.cavallotto@mib.infn.it, giuseppe.lavacca@mib.infn.it, leone.bacciu@unive.it, marco.nobile@unive.it, massimo.gervasi@mib.infn.it, matteo.grazioso@unive.it, sabina.rossi@unive.it

We present COSMICA, an opensource high-performance GPU-accelerated numerical code for modeling cosmic ray solar modulation, and its application to study CR diffusion parameters. Developed within the framework of the ICSC-Italian Research Center on High-Performance Computing, Big Data and Quantum Computing (Spoke-3), COSMICA is undergoing continuous software optimization to maximize efficiency on NVIDIA architectures. COSMICA is coupled with SDEGnO, another ICSC project, designed for the efficient parameter tuning, exploring the large parameter space in solar modulation studies. As a first physical use-case study, we exploit COSMICA to investigate Forbush Decreases (FDs), which are transient cosmic ray intensity reductions caused by interplanetary disturbances. The analysis leverages the high-precision daily measurements from AMS-02, which provide cosmic ray fluxes across a wide range of rigidities. The ability to simultaneously study not only protons but also helium isotopes offers complementary insights into charge- and mass-dependent transport effects. By analyzing FD events, we assess localized variations in diffusion parameters and their impact on cosmic ray transport. The results confirm the stability of the rigidity dependence of the diffusion tensor, supporting the use of FDs as probes of localized well constrained heliospheric conditions. The computational efficiency of COSMICA paves the way for large-scale simulations, systematic FD catalogue analysis and a more in-depth understanding of the parameter that regulates the solar modulation, together with their dependencies.

ArXiv ePrint: TBD
39th International Cosmic Ray Conference (ICRC2025)
15–24 July 2025
Geneva, Switzerland

*Speaker

PoS(ICRC2025)1285

1. Introduction

The passage of particles through an astrophysical plasma is described using the Fokker-Planck equation. In particular, models describing the passage through the heliosphere are based on the Parker Transport Equation (PTE), named after its author in 1965 [1]. In this formulation, cosmic rays experienced diffusion-dominated transport with adiabatic energy losses due to the expansion of the Solar Wind (SW). The full treatment of the PTE includes other processes like convection, due to solar wind expansion, and drift of GCR gyro-radius guiding centers, due to large-scale spatial gradients of the heliospheric magnetic field (HMF) [2–6]. The net effect of these processes is an energy and time-dependent reduction of cosmic ray intensity in the inner heliosphere regarding the local interstellar spectra that is known as Solar Modulation.

The importance of accurately describing the solar modulation process is becoming increasingly important as the precision of experiments increases. In fact, nowadays, space spectrometers like AMS–02 can produce time-dependent spectra with unprecedented accuracy [7–10] that, in turn, stimulate new studies on more fundamental questions (*e.g.*, see, Ref's [11–14]). In this context, the required computational effort could be a breaking point in the development of accurate models. In parallel, the development of GPU parallel computing hardware reaches its golden era, allowing us to have high-performance computers even in a laptop and, thus, several authors developed their own version of PTE solvers on GPU [15–19].

In this paper, we present COSMICA (COde for a Speedy Montecarlo Involving Cuda Architecture), a speedy and high-precision Monte Carlo simulator of CR modulation, which solves the system of Stochastic Differential Equations (SDE) equivalent to the PTE. In this implementation, we place performance on top priority while maintaining the numerical accuracy as high as possible.

2. COSMICA Under the Lens

The numerical resolution approach used by COSMICA is based on the equivalence between the Fokker-Planck Equation with a set of Stochastic Differential Equations [20]. The modulated spectra are then evaluated given the GCR differential intensities at the boundary as the primary input, *i.e.*, outside the heliosphere, also known as the (very) Local Interstellar Spectrum (LIS). Numerical methods that solve PTEs employing the Monte Carlo integration of an equivalent set of SDEs are nowadays a standard method [16, 20–25] since this approach presents several advantages in terms of stability of the solution and modularity of the code, and allows expanding the model's functionality with relatively small changes in the code. The code is actually under refinement and a public release under the *GNU GPL v3* license is expected for winter 2025. One of the important features of the code is the possibility to solve within the same execution the spectra for different isotopes, different periods, or different heliospheric parameters. In fact, in most of the utilization scenarios, the PTE must be solved several times under similar parametrizations to reproduce experimental data. As a representative example, let us consider the simulation of a Bartel rotation-averaged spectrum for Helium. This ion is composed of ^{3}He and ^{4}He isotopes in a proportion where the ^{3}He / ^{4}He ratio is $\sim$ 0.10–0.16 [26] so that, for the most accurate evaluation, both elements must be considered in the computation of the helium-modulated spectra. Instead of running two separate programs, which means duplicating the needed resources or loading in memory the same datasets twice for

each solution, COSMICA can solve PTE for both isotopes simultaneously, providing final separate results for each case and allowing different post-processing analyses. COSMICA can also use the same approach, testing different parameter values with boosting in performances that increase while scaling up the complexity of the problem. This feature, in combination with the capability to run also in multi-GPU mode, allows COSMICA to use resources most efficiently, *e.g.*, reducing lag times due to memory copying and keeping the GPU fully occupied for the majority of the time. In the next section, we will see some performance plots and benchmark studies to test the capability of this code.

3. Code Performance and Validation

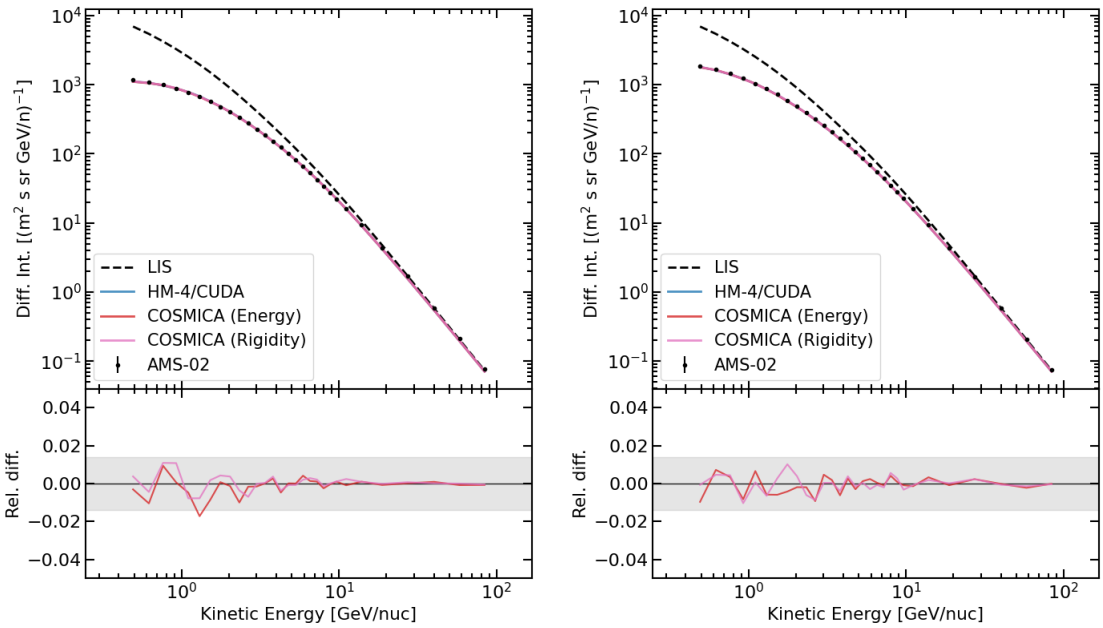

Figure 1: Modulated spectra for HelMod-4/CUDA (HM-4/CUDA) and for COSMICA, both solving the PTE in kinetic energy and in rigidity for November 2011 (left) and June 2017 (right). The bottom panel shows the relative difference regarding HelMod-4/CUDA.

To validate COSMICA, we implemented in it the physical description of HelMod-4/CUDA [19, 24]. Particular emphasis was placed on the prevention of branch divergence and the reduction of registry pressure during the implementation process. These are the most dangerous threats while implementing physical description: an increase in registry pressure will lead to a reduction of threads that can be allocated within a block and, in turn, a suboptimal allocation of resources within the GPU; branch diverging may appear when *if* statements lead to computing different formulas according, *e.g.*, to position, energy or time. Since the GPU executes the program in *warps* of 32 threads, executing the same instruction in all threads, in case of branch divergence, some threads are in an idle state and only a fraction are actually computing, resulting in a loss of performance. To ensure that this optimization in rewriting the physical description does not introduce any bias, we compared our results with those provided by HelMod-4/CUDA. Two sample results are presented

in Fig. 1. In these figures, we compute the modulated spectra for HelMod-4/CUDA (HM-4/CUDA) and for COSMICA, both solving the PTE in kinetic energy and in rigidity. The bottom panel shows the relative difference among the various versions that behave well within numerical uncertainties of $(N_{pp})^{0.5}$ ~1.4% where $N_{pp} = 5024$ is the number of simulated pseudo particles for each energy bin. Comparing HelMod-4/CUDA and COSMICA, we found the latter is largely faster [27] with still some room for improvements in multi-GPU mode.

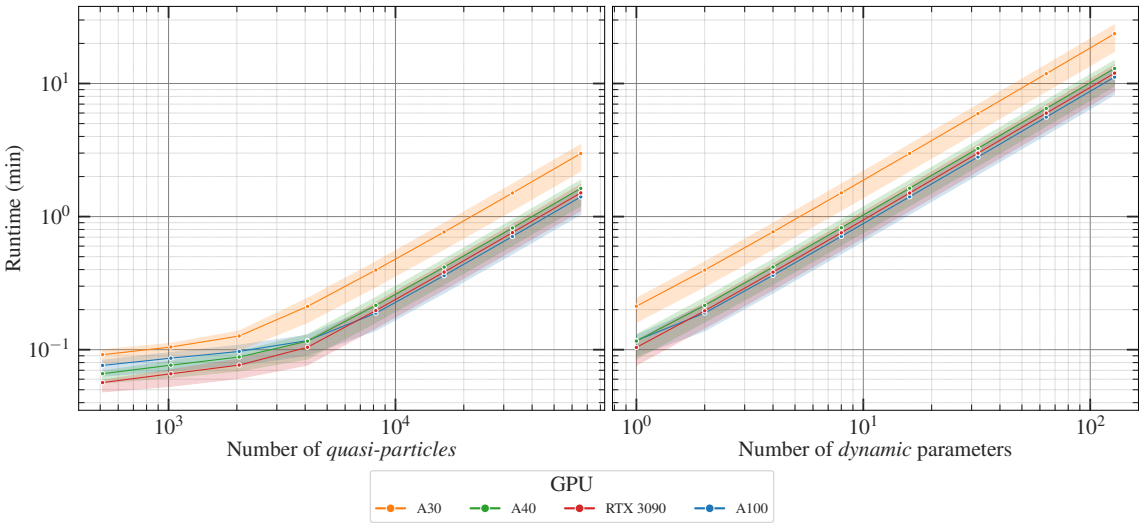

Figure 2: Scaling analysis of the COSMICA multi-GPU implementation versus HelMod-4/CUDA for proton + deuteron simulations (23 bins, 1–11 GV) across different GPU architectures and simulation periods (Carrington rotations: 2117, 2155, 2168, 2181, 2192, 2209). **(Left)** Runtime as a function of *quasi-particles* count for different periods. **(Right)** Runtime as a function of dynamic parameters count (K_0) for different periods with fixed *quasi-particles*. All results are shown for different simulation periods, with shaded areas indicating the 50% percentile interval.

We tested COSMICA on different NVIDIA hardware, in particular, we tested the GPU models: A100, A30, A40, and RTX 3090. The models A30 and A100 have the same chipsets and share the same Compute Capability (the hardware features and supported instructions), but the A100 has a larger and faster memory and almost twice as many cores as the A100 (6912 vs. 3584). The models RTX 3090 and A40 also share the same chipset but have a more modern compute capability. The RTX 3090 is a consumer gaming GPU, characterized by a much faster boost clock (1900 MHz), but 2.4% fewer cores and a higher TDP (maximum amount of heat the GPU is designed to dissipate under normal operation) of 350W (against 300W of the A40). The comparison is reported in Fig. 2. In more detail, in Fig. 2-right, we compare the execution time of COSMICA as a function of the number of parameters we want to test in the same simulation. Please note that, normally, a simulation computes one parameter set each; thus, testing more parameters increases linearly the execution times since these must be computed sequentially (on a single GPU). COSMICA also allows testing multiple parameter sets, allowing an optimization in this procedure that is on the order of 10000% faster regarding the not optimized approach. In Fig. 2-left, we compare the execution time of COSMICA as a function of the number of pseudo-particles simulated in each energy bin. It is interesting to note that performances are quite similar, confirming that the problem is more

computing-bound rather than memory-bound. Moreover, great performances can be obtained even with GPUs not designed to be installed on HPC infrastructure, which makes the code suitable without modification for scaling the approach from local search and test on a desktop to heavy parameter searches on a multi-GPU HPC infrastructure.

4. Conclusions

We presented COSMICA as the newcomer within the panorama of solar modulation numerical solvers. COSMICA is designed to reduce the execution time without compromising the accuracy and stability of the solution. Moreover, the design of the code allows performing parameter sweep studies in a very efficient way that will reduce the computation time in future modulation studies. The code will be released as open source under the *GNU GPL v3* license.

References

[1] E.N. Parker, *The passage of energetic charged particles through interplanetary space*, *Plan. Space Sci.* **13** (1965) 9.

[2] J.R. Jokipii, E.H. Levy and W.B. Hubbard, *Effects of particle drift on cosmic-ray transport. I - General properties, application to solar modulation*, *Astrophys. J.* **213** (1977) 861.

[3] J.R. Jokipii and B. Thomas, *Effects of drift on the transport of cosmic rays. IV - Modulation by a wavy interplanetary current sheet*, *Astrophys. J.* **243** (1981) 1115.

[4] J.R. Jokipii and E.H. Levy, *Effects of particle drifts on the solar modulation of galactic cosmic rays*, *Astrophys. J. Lett.* **213** (1977) L85.

[5] M. Potgieter and S. Ferreira, *Modulation of cosmic rays in the heliosphere over 11 and 22 year cycles: a modelling perspective*, *Advances in Space Research* **27** (2001) 481.

[6] N.E. Engelbrecht, F. Effenberger, V. Florinski, M.S. Potgieter, D. Ruffolo, R. Chhiber et al., *Theory of Cosmic Ray Transport in the Heliosphere*, *Spac. Sc. Rev.* **218** (2022) 33.

[7] M. Aguilar, L.A. Cavasonza, G. Ambrosi, L. Arruda, N. Attig, F. Barao et al., *Periodicities in the Daily Proton Fluxes from 2011 to 2019 Measured by the Alpha Magnetic Spectrometer on the International Space Station from 1 to 100 GV*, *Phys. Rev. Lett.* **127** (2021) 271102.

[8] M. Aguilar, L.A. Cavasonza, G. Ambrosi, L. Arruda, N. Attig, F. Barao et al., *Properties of Daily Helium Fluxes*, *Phys. Rev. Lett.* **128** (2022) 231102.

[9] M. Aguilar, L.A. Cavasonza, G. Ambrosi, L. Arruda, N. Attig, C. Bagwell et al., *Temporal Structures in Electron Spectra and Charge Sign Effects in Galactic Cosmic Rays*, *Phys. Rev. Lett.* **130** (2023) 161001.

[10] M. Aguilar, G. Ambrosi, H. Anderson, L. Arruda, N. Attig, C. Bagwell et al., *Temporal Structures in Positron Spectra and Charge-Sign Effects in Galactic Cosmic Rays*, *Phys. Rev. Lett.* **131** (2023) 151002.

[11] M. Di Mauro, M. Korsmeier and A. Cuoco, *Data-driven constraints on cosmic-ray diffusion: Probing self-generated turbulence in the Milky Way*, *Phys. Rev. D* **109** (2024) 123003 [2311.17150].

[12] Y. Yang, X. Luo, X. Song, W. Xu and M.S. Potgieter, *A numerical study of unusual flux decreases for cosmic ray protons and electrons observed by Alpha Magnetic Spectrometer in 2017*, *Astron. Astrophys.* **694** (2025) A197.

[13] S. Della Torre, G. Cavallotto, G. La Vacca and M. Gervasi, *General properties of charged particle diffusion in heliosphere inferred from forbush decreases*, *Advances in Space Research* (2025) .

[14] D. Maurin, L. Audouin, E. Berti, P. Coppin, M. Di Mauro, P. von Doetinchem et al., *Precision cross-sections for advancing cosmic-ray physics and other applications: a comprehensive programme for the next decade*, *arXiv e-prints* (2025) arXiv:2503.16173 [2503.16173].

[15] P. Dunzlaff, R.D. Strauss and M.S. Potgieter, *Solving Parker's transport equation with stochastic differential equations on GPUs*, *Computer Physics Communications* **192** (2015) 156.

[16] A. Vogt, N.E. Engelbrecht, R.D. Strauss, B. Heber, A. Kopp and K. Herbst, *The residence-time of Jovian electrons in the inner heliosphere*, *Astron. Astrophys.* **642** (2020) A170.

[17] M. Solanik, P. Bobik and J. Genci, *Cosmic rays modulation in heliosphere models on GPU*, in *Proceedings of 37th International Cosmic Ray Conference — PoS(ICRC2021)*, vol. 395, p. 1320, 2021, DOI.

[18] M. Solanik, P. Bobik and J. Genči, *Geliosphere - parallel cpu and gpu based models of cosmic ray modulation in the heliosphere*, *Computer Physics Communications* **291** (2023) 108847.

[19] M.J. Boschini, G. Cavallotto, S. Della Torre, M. Gervasi, G. La Vacca, P. Rancoita et al., *Fast and accurate evaluation of deep-space galactic cosmic ray fluxes with helmod-4/cuda*, *Advances in Space Research* **74** (2024) 4302.

[20] P. Bobik, M.J. Boschini, S. Della Torre et al., *On the forward-backward-in-time approach for monte carlo solution of parker's transport equation: One-dimensional case*, *Journal of Geophysical Research: Space Physics* **121** (2016) 3920.

[21] F. Effenberger, H. Fichtner, K. Scherer and I. Büsching, *Anisotropic diffusion of Galactic cosmic ray protons and their steady-state azimuthal distribution*, *Astron. Astrophys.* **547** (2012) A120.

[22] A. Kopp, I. Büsching, R.D. Strauss et al., *A stochastic differential equation code for multidimensional Fokker-Planck type problems*, *Comput. Phys. Commun.* **183** (2012) 530.

[23] L.L. Zhao, G. Qin, M. Zhang and B. Heber, *Modulation of galactic cosmic rays during the unusual solar minimum between cycles 23 and 24*, *Journal of Geophysical Research (Space Physics)* **119** (2014) 1493.

[24] M.J. Boschini, S. Della Torre, M. Gervasi, G. La Vacca and P.G. Rancoita, *The helmod model in the works for inner and outer heliosphere: From ams to voyager probes observations*, *Adv. Space Res.* **64** (2019) 2459 .

[25] K.D. Moloto, N.E. Engelbrecht, R.D. Strauss, D.M. Moeketsi and J.P. van den Berg, *Numerical integration of stochastic differential equations: A parallel cosmic ray modulation implementation on Africa's fastest computer*, *Advances in Space Research* **63** (2019) 626.

[26] M.J. Boschini, S. Della Torre, M. Gervasi, D. Grandi, G. Jóhannesson, G. La Vacca et al., *Spectra of he isotopes and the 3he/4he ratio*, *The Astrophysical Journal* **981** (2025) 85.

[27] G. Cavallotto, S.D. Torre, G.L. Vacca and M. Gervasi, *Cosmica: a novel parallel gpu code for cosmic rays propagation in heliosphere*, in *2025 33rd Euromicro International Conference on Parallel, Distributed, and Network-Based Processing (PDP)*, pp. 420–427, 2025, DOI.

Acknowledgments

This activity is supported by Fondazione ICSC, Spoke 3 Astrophysics and Cosmos Observations. National Recovery and Resilience Plan (Piano Nazionale di Ripresa e Resilienza, PNRR) Project ID *CN*_00000013.
The project *"SDEGnO – Optimization and performance testing of CUDA-(multi)GPU-accelerated codes for the automatic parameterization of physical models based on SDE Monte Carlo algorithms"* is funded through the *"Bando a Cascata"* issued by Spoke 3 Astrophysics and Cosmos Observations, under the National Recovery and Resilience Plan (*Piano Nazionale di Ripresa e Resilienza*, PNRR, Project ID CN_00000013, Spoke Leader: Istituto Nazionale di Astrofisica – INAF, CUP C53C22000350006), and is carried out under the Agreement Rep. 531/2024 Prot. 283696/2024 between INAF and Ca' Foscari University of Venice.
We acknowledge the ICSC for awarding this project access to the EuroHPC supercomputer LEONARDO, hosted by CINECA (Italy).
This work was also supported by the high-performance computing infrastructure developed under the project "CONVECS", funded by the PR Veneto FESR 2021-2027 program, Priority 1 – Specific Objective 1.1 – Action 1.1.2.
MG, SDT and GLV are supported by INFN and ASI under ASI-INFN Agreement No. 2019-19-HH.0 and its amendments and by ASIF implementation agreement No. 2021–36-HH.0 involving ASI and Milano-Bicocca University.