# Calculation of the High-Energy Neutron Flux for Anticipating Errors and Recovery Techniques in Exascale Supercomputer Centres

Hernán Asorey[1,2*] and Rafael Mayo-García[3]

[1*] Medical Physics Department, Comisión Nacional de Energía Atómica, Centro Atómico Bariloche, Av. E. Bustillo 9500, San Carlos de Bariloche, R8400, RN, Argentina.

[2] Instituto de Tecnologías en Detección y Astropartículas (ITeDA), Comisión Nacional de Energía Atómica, Centro Atómico Constituyentes, Av. Gral Paz 1499, Villa Maipú, B1650, Buenos Aires, Argentina.

[3] Technology Department, Centro de Investigaciones Energéticas, Medioambientales y Tecnológicas (CIEMAT), Av. Complutense 40, Madrid, 28040, Madrid, Spain.

*Corresponding author(s). E-mail(s): hernanasorey@cnea.gob.ar;

**Abstract**

The age of exascale computing has arrived and the risks associated with neutron and other atmospheric radiation are becoming more critical as the computing power increases, hence, the expected Mean Time Between Failures will be reduced because of this radiation. In this work a new and detailed calculation of the neutron flux for energies above ~~1 GeV~~50 MeV is presented. This has been done by using state-of-the-art Monte Carlo astroparticle techniques and including real atmospheric profiles at each one of the next 23 exascale supercomputing facilities location. Atmospheric impact in the flux and seasonal variations were observed and characterised, and the barometric coefficient for high-energy neutrons at each site were obtained. With these coefficients, potential risks of errors associated with the increase in the flux of energetic neutrons, such as the occurrence of single event upsets or transients, and the corresponding failure-in-time rates, can be anticipated just by using the atmospheric pressure before the assignation of resources to critical tasks at each exascale facility. For more clarity, examples about how the rate of failures is affected by the cosmic rays are included, so administrators will better anticipate which more or less restrictive actions could take for overcoming errors.

**Keywords:** neutron flux, supercomputing, HPC, exascale, atmospheric radiation

# 1 Introduction

Exascale computing presents several issues, being fault tolerance one of the main ones: while the Mean Time Between Failures (MTBF) of the hardware components (from coolers to memories or random issues) does not grow as fast as the number of resources, the number of cores on a hardware unit experiences continuous growth, and so the probability of one or more tasks being affected by a failure increases [1]. For example, large parallel jobs may fail as frequently as once every 30 minutes on exascale platforms [2]. Also, the higher number of tasks composing a job, the higher will be the computational and economics lost associated with the increasing number in failures. Although these issues pose enough of a risk, additional factors are now coming into play: clusters with lower energy consumption that are designed and fed with a lower voltage, or smaller circuits are more easily upset because they carry smaller charges and are more prone to hardware failures, or supercomputers (partially) built with GPUs cards counting on an amazing number of cores, or much more complex software being executed, etc. All the previous results in a higher failure rate, and so, lower values of the MTBF. Thus, there is a necessity in developing tools and frameworks that reduce the impact of tasks and jobs failure on exascale supercomputers.

Traditionally, general fault-tolerant behaviour has been achieved by redundancy and checkpointing mechanisms. Isolated redundancy is not an ideal approach for HPC as it leads to performance loss, but it has provided nice results in HTC environments (Desktop, Grid, Cloud) or combined with additional methods. Checkpointing techniques have provided good results on a three-fold basis (system-, user-, and application-level) and have demonstrated a wide scenario of solutions on coordinated and uncoordinated actions, roll-back and roll-forward strategies, mono- and multilevel checkpointing, etc.

Even more and beyond the proper interest of resilience, a consequence of the increase of parallelism both on the hardware and applications sides was a series of problems related to task scheduling. The idea was to assign tasks to resources trying to avoid starvation, deadlocks, and performance losses, all while having the cluster as full as possible. This computing efficiency improvement could be achieved by profiting from a proactive (not reactive to failures) checkpointing strategy that could be designed as part of the resource manager scheduler. For example and among other results, the user-level checkpointing library DMTCP was seamlessly integrated into Slurm [3]. By designing several dynamic scheduling algorithms and profiting from a new command (`smigrate`), a more resilient system was provided in which also proactive checkpointing actions could be performed for enhancing the computing and energy efficiency by dynamically migrating tasks previously saved with such a checkpoint with low overhead.

This fact has opened the door to new possibilities such as non-invasive maintenance operations, job preemption, more advanced priority policies, lower energy consumption, etc. Then, further advances must be envisioned once traditional checkpointing and rollback recovery strategies have been accomplished. In this regard, Silent Data Corruption (SDC) errors, or simply, silent errors (SE) have become a cornerstone in the path to exascale computing. Soft errors can be mainly classified into two categories: bit-flipping error (e.g., 1 becomes 0) in RAM; and computation error (e.g. $1 + 1 = 3$) in floating point units. Traditionally, bit-flipping errors have been handled by the Error Correcting Code (ECC) technique, and computation error is dealt with redundancy methods (ECC cannot handle computation error). Unlike aforementioned fail-stop failures, such latent errors cannot be detected immediately, and a mechanism to detect and overcome them must be provided as they are becoming a major drawback as the supercomputer complexity grows. In other words, failures become a normal part of application executions and, among them, SEs are nowadays those with scarce valid solutions properly tested on real environments.

It has been shown that SE are not unusual and must also be accounted for [4]. The cause may be soft efforts in L1 cache, arithmetic errors in the Arithmetic Logic Unit (ALU), (double) bit flips due to cosmic radiation, etc. The problem is that the detection of a latent error is not immediate, because the error is identified only when the corrupted data is activated. One must then account for the detection interval required to detect the error in the error recovery protocol. Indeed, if the last checkpoint saved an already corrupted state, it may not be possible to recover from the error. Hence, the necessity to keep several checkpoints so a valid one could roll back to the last correct state. When dealing with SE, however, faults can propagate to other processes and checkpoints, because processes continue to participate and follow the protocol during the interval that separates the occurrence of the error from its detection.

Summarizing, there is a clear necessity for overcoming SE as they are becoming inevitable with the ever-increasing system scale and execution time, and new technologies that feature increased transistor density and lower voltage. Nevertheless, the question of the source for these SE arises. The answer can be found in the atmospheric cosmic-induced radiation, in which neutrons play a key role. As neutrons are produced during the interaction of cosmic rays with the atmosphere, and since this last experience seasonal changes, the latitude, longitude, and altitude where a data centre hosts an exascale supercomputer as well as the atmospheric seasonal conditions determine the number of the neutrons reaching the infrastructure and, consequently, the predicted MTBF. So, in this work, using the current techniques for calculating the flux of the expected radiation at the ground originated by the cosmic ray flux, the flux of neutrons with energy $E_n \geq 50\,\mathrm{MeV}$ averaged per season in 23 data centres are presented. Among these places, the ones already hosting or expecting to promptly host an exascale supercomputer in China, Europe, Japan, and the United States are included. The geographic distribution of the 23 exascale supercomputing centres is shown in Figure 1 and Table 1.

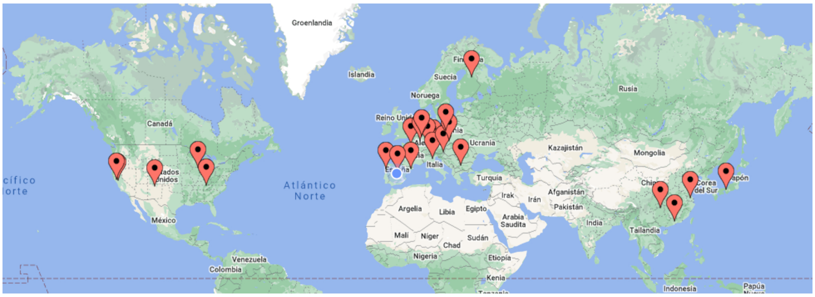


**Fig. 1** Geographic locations of the 23 exascale supercomputing centres that are being built around the World

[paragraph adapted] ~~Roughly speaking, a flux of about $\sim 13\,\mathrm{neutrons\,cm^{-2}\,h^{-1}}$) reaches the ground at sea level.~~ High-energy neutrons, i.e., neutrons with an energy higher than 10 MeV, with a total flux of about $13\,\mathrm{neutrons\,cm^{-2}\,h^{-1}}$ in New York at sea level [5, 6] are expected to cause SE [5], but the flux of neutrons varies with the geographical location [7], altitude [8], atmospheric [9] and geomagnetic and heliospheric conditions [10]. As it will be shown later in this work (see the 8$^{\mathrm{th}}$ column of the table 1 in section 4.1), depending on the location the averaged flux of neutrons for $E_n > 50\,\mathrm{MeV}$ could vary between $(3.7 \pm 0.2)\,\mathrm{cm^{-2}\,h^{-1}}$ in Guangzhou, China, at sea level and $(26.4 \pm 1.1)\,\mathrm{cm^{-2}\,h^{-1}}$ in Los Alamos, USA, at 2,125 m above sea level (asl).

The whole integration of the main source of SE (cosmic radiation) jointly with their prediction process according to the geographical place where such radiation occurs (computing infrastructure location) in a specific season of the year is expected to be useful to the administrators of these supercomputers, given a quantitative measure of the changes in the expected flux of neutrons due to changes in the barometric pressure at the ground level. With all this information, system administrators will be capable of designing and applying different mathematical and software solutions to cope with these SE that will produce more or less overhead. This work is expected to be a decision-making tool for the exascale supercomputers' administrators as they will be able to determine in advance which mitigation methodologies need to be applied for overcoming SE depending on the forecasted neutron flux in a specific period of the year. ~~and both the available data of produced errors in some supercomputers already~~

~~determined and the quantification and qualification of radiation effects on applications' output correlating the number of corrupted elements with their spatial locality.~~

The main result from this exercise will be a higher resilience, better computational efficiency and less energy misuse in exascale supercomputers.

## 2 Related work

Fault tolerance can be defined as the capability of a certain system to overcome hardware, software or communication problems and continue with the execution of applications. This field embraces different sections: the detection of failures, their avoidance if possible, and the recovery from them if not.

~~In order to~~To achieve computational resilience, there are several methodologies for overcoming errors produced in runtime. Technical progress in resilience has been achieved in the last decade, but the problem is not actually solved and the community is still facing the challenge of ensuring that exascale applications complete and generate correct results while running on unstable systems [11]. In this regard, it should be pinpointed that current systems do not have a fully integrated approach to fault tolerance: the different subsystems (hardware, parallel environment software, parallel file system) have their ~~own~~ mechanisms for error detection, notification, recovery, and logging.

The current status can be mostly described in a few articles. In [12], different approaches towards failure detection and prediction are presented. A state-of-the-art description of the approaches to overcome these failures is included in [11], where also a more detailed explanation of checkpoint solutions is presented. An updated status can be found in the compilation of fault detection, fault prediction, and recovery techniques in HPC systems, from electronics to system level, which also analyzes their strengths and limitations and identifies promising paths to meet the reliability levels of exascale systems [13]. These references clearly show that the problem being faced is of real interest in the next generations of supercomputers.

After a failure has been detected (even pre-emptively), checkpoints are a widely used tool devoted to saving the status of the running tasks. A recent survey of checkpointing protocols can be found in the book edited by Hérault and Robert [14]. Strategies also range from coordinated checkpointing (including full and incremental ones) to uncoordinated checkpoint and recovery with message logging, each with different strengths and drawbacks [15]. Checkpoint pursues to reduce the overhead produced by replication methodologies even when the latter is producing valid results still [16].

The coordinated checkpoint technique guarantees consistent global states by enforcing each of the processes to synchronize their checkpoints as it is the most common practical choice due to the simplicity of recovery [17]. The obvious issue is to find a balance between the robustness of iterated checkpoints and the induced overhead. Uncoordinated checkpointing allows different processes to do checkpoints when it is most convenient but is subject to the domino effect, and does not guarantee progress. Although this issue can be avoided with message logging [18], uncoordinated checkpointing does not represent a valid alternative in the majority of current production environments and applications.

Recent advances include multi-level approaches, or the use of SSD or NVRAM as secondary storage [11] as well as the replication for redundant MPI processes [19] and threads [20]. Also, on MPI, it is remarkable the initial FT-MPI introduced to enable MPI based software to recover from process failure [21] and, also, the enlarged capacities via the Checkpoint-on-Failure protocol for forwarding recovery MPI without resulting in a major overhead [22]. Recently, the User Level Failure Mitigation (ULFM) interface provides new opportunities in this field, enabling the implementation of resilient MPI applications, system runtimes, and programming language constructs able to detect and react to failures without aborting their execution [23]. Another development is MANA (MPI-Agnostic Network-Agnostic transparent checkpointing) for MPI [24], which proposes a new solution especially deserved for exascale [25]. The three major approaches to implementing checkpoint systems are application-, user- and system-level (or kernel-level) implementations [26], being the last one always transparent to the user. The most popular approach is the application-level checkpoint [11], where the programmer defines which is the state to be stored in the application by directly injecting the checkpointing routines directly into the code, or by using some automated pre-processors. This approach keeps being of interest as new solutions are proposed, such as the application-based focused recovery (ABFR) [27]. This alternative has however been mostly abandoned in the place of the other two, and up to the authors' knowledge there are currently no significant projects in the area.

With the user-level approach, a library is used to do the checkpointing and the application programs are linked to the library. User-level does not require system privileges to operate either special kernel modules or kernel patches. One of the active projects for transparent user-level checkpoints are DMTCP [28] or BLCR [29], which include support for distributed and multi-threaded applications and do not require modifying either the application executable or the kernel.

Concerning the state-of-the-art of research on SE, (parallel) jobs can be interrupted at any time for checkpointing, for a nominal cost $C$. To deal with fail-stop failures, the execution of divisible-load applications is partitioned into same-size chunks followed by a checkpoint, and there exist well-known formulae by Young & Daly [30] to determine the optimal checkpointing period. To deal with SE, the simplest protocol had been to perform a verification (at a cost $V$) just before taking each checkpoint. If the verification succeeds, then one can safely store the checkpoint and mark it as valid. If the verification fails, then an error has struck since the last checkpoint, which is correct having been verified, and one can safely recover (which takes a time $R$) from that checkpoint to resume the execution of the application. This protocol with verifications zeroes out the risk of fatal errors that would force restarting the execution from scratch, but the key point is to find a pattern that minimizes the expected execution time of the application. Finding the best trade-off between error-free overhead (what is paid due to the resilience method, when there is no failure during execution) and execution time (when errors strike) is not trivial [31].

Later on, it has been published a work for determining the real computational cost in the technique of combining replication and checkpointing [32] for assessing either duplication or triplication, which can be acceptable solutions for specific scenarios (aeronautics, for example, though it also requires manufacturing specific hardware as IBM S/390 in Boeing 777 [33]). Though it does not specifically try to cope with SE, this work is of interest as it provides closed-form formulas that give the optimal checkpointing period and optimal process count as a function of the error rate, checkpoint cost, and platform size. Similar work on predicting an optimal checkpointing period and its relationship with the cluster size has been recently published [34].

In addition to software techniques, SE can be coped with mathematical approaches. The traditional wisdom in computing no longer applies as unorthodox, new algorithmic techniques are emerging linked to the exascale requirements. Aspects related to communicating avoiding algorithms, mixed single-double precision computations or the inclusion of new kinds of randomised algorithms embedded in deterministic portions of the codes are of major concern in the context of faster and more reliable solvers [35].

These new methods are insensitive to the quality of the randomness and produce highly accurate results, besides their simplicity and speed [36, 37]. Hence, there is currently a large interest in conducting further research on them [38, 39]. Specific recent works applied to GMRES [40] or parallel stencil computations [41] also demonstrate the interest in this topic.

Last but not least, there are some works on radiating computing hardware. More than twenty years ago, it has been demonstrated that neutrons originated in cosmic radiation are the dominant source of soft errors in DRAM devices [42], and cosmic-ray induced soft error rates were measured on 16-Mb DRAM memory chips [43]. Later on, in 2002 and 2003, to prove to the manufacturers that the errors appearing in ASC-Q at Los Alamos National Laboratory were due to cosmic rays, the staff placed one of the servers in a beam of neutrons causing errors to spike [44]. The Jaguar supercomputer logged single-bit ECC errors at a rate of 350 $\text{min}^{-1}$ in 2006 as well as double-bit errors once per day, being the latter detected, but not corrected by ECC technique as previously stated. Also, BlueGene/L at Lawrence Livermore Nat Lab suffered with radioactive lead in the solder to cause bad data in the L1 cache, a problem that ended in slower computations as L1 had to be bypassed.

[new paragraphs about neutron energy error production mechanism and the effective error cross-sections and including some previous references.] The main effects of radiation on semiconductors are the total ionizing dose (TID), the occurrence of Single Event Effects (SEE), and Displacement Damage (DD). For high-energy neutrons, both the elastic and inelastic interactions are possible, and scattering producing a displacement of atoms from their position in the lattice site results in defects altering the electronic properties of the crystal and being one of the main mechanisms of device degradation [45]. The neutron interacts with atoms creating DD and generating secondary charged ionizing particles: a neutron of energy $E_n = 100\,\mathrm{MeV}$ can produce a cascade of secondary particles including secondary neutrons, protons, ions, photons and $\delta$ electrons with energy above 100 eV, extending temporal effects and permanent damage far away from the first interaction site [46]. Detailed simulations show that, while the elastic neutron-$^{28}$Si interaction cross-section decreases from $\sim 1,000\,\mathrm{mb}$ for $E_n \simeq 8\,\mathrm{MeV}$ down to $\simeq 450\,\mathrm{mb}$ at $E_n \simeq 100\,\mathrm{MeV}$ and remains constant up to $E_n \gtrsim 1000\,\mathrm{MeV}$, the corresponding inelastic cross-section curve starts at $\sim 800\,\mathrm{mb}$ for $E_n \simeq 10\,\mathrm{MeV}$, peaking at $\gtrsim 1,000\,\mathrm{mb}$ at $\simeq 80\,\mathrm{MeV}$ and then it stabilizes at $\simeq 200\,\mathrm{mb}$ for $E_n \gtrsim 1\,\mathrm{GeV}$ (see Fig. 3 of [46]), where $100\,\mathrm{mb} = 0.1\,\mathrm{barn} = 10^{-25}\,\mathrm{cm}^2$ means that about 4.2% of the incident neutrons interacts with the $^{28}$Si. Some typical reactions observed involve different mechanisms with energy thresholds between 2.75 and 12.99 MeV, and producing $\alpha$s, such as $^{28}\mathrm{Si}(n,\alpha)^{25}\mathrm{Mg}$ and $^{28}\mathrm{Si}(n,2\alpha)^{21}\mathrm{Ne}$, or neutrons, such as $^{28}\mathrm{Si}(n,n\alpha)^{24}\mathrm{Mg}$, or neutrons and protons, such as $^{28}\mathrm{Si}(n,np)^{27}\mathrm{Al}$ [47]. Similar reactions occur with neutrons and oxygen, increasing the probability of having errors with the incident energy as $SiO_2$ is typically in the proximity to active junction areas [48]. Alia *et al.* [49] exposed commercial SRAM devices to different flux of protons (30-200 MeV) and neutrons (5-300 MeV) and measured the effective $\sigma_{\mathrm{err}}$ for both types of SEE: soft errors, also known as single event upsets (SEU) in the literature, and hard (or catastrophic) errors just as the single event latch-up (SEL). By using fitting their experimental data to Weibull functions they compared the $\sigma_{\mathrm{SEU}}$ for neutrons at different energies with the same magnitude for energetic proton, and observed that the behaviour of $\sigma_{n,\mathrm{SEU}}$ depends both on the neutron energy and on the internal geometry of the device, and that $\sigma_{n,\mathrm{SEU}}$ tends to $\sigma_{p,\mathrm{SEU}}$ of protons at $E_p = 250\,\mathrm{MeV}$ for $E_n \gtrsim 25\,\mathrm{MeV}$ (see Figure 3 of [49]).

As the incident neutron energy gets higher, the number of new reactions in the pathway increases, extending the damage and the probability of having errors from a single reaction. As it will be detailed in section 4, it is possible to characterize the radiation-induced errors in computing devices by defining an effective cross-section, $\sigma_{\mathrm{err}}$, a widely used magnitude to directly evaluate the radiation sensitivity of a particular device [47]. As it is an effective metric, it considers all the possible sources of neutron-induced computing errors, and it is experimentally measured by placing different devices in a neutron beam and calculating the fraction of the observed rate of neutron-induced errors to the injected neutron flux [50]. The Los Alamos Neutron Science Center (LANSCE) irradiation facility is one of the neutron sources typically used to measure the number of fatal soft errors, such as the measurement performed in the ASC-Q supercomputer, one of the world's fastest supercomputers in 2005 [44], and in the Titan supercomputer, which is composed of more than 18,000 Kepler GPUs, has a radiation-induced MTBF in the order of dozens of hours [50].

[end of the new paragraphs.]

Thus, new works on this SE problem produced by radiation have been more recently published focusing on determining the reliability in GPUs [6] and Xeon Phis also applying high-level fault injection [51], where the relative $\sigma_{\mathrm{err}}$ for each device exposed to high-energy neutrons have been obtained. Further steps forward have been the comparison between high-energy and thermal neutrons effects on the error rates on Commercial Off-The-Shelf (COTS) devices [52] by exposing AMD APUs (4 Steamroller CPUs + 1 AMD Raedon R7), Intel XeonPhi processors, Nvidia K20, TitanX and TitanV GPUs and a Zync-7000 FPGA to two beam of neutrons with energies in the range from 1 meV to 1 GeV in the ChipIR and Rotax neutron beam-lines at the ISIS Neutron and Muons Source [52]. They conclude that while high-energy neutrons are the most important source of SE, for some applications in some computing devices thermal-neutrons can account up to 59% of the total MTBF. In a latter work, an experimental evaluation of the ~~probability~~effective cross-section $\sigma_{\mathrm{err}}$ for a high-energy vs thermal neutron to generate an error in the same computing devices is provided as well as an estimation of the thermal neutrons flux modification due to materials heavily present in a supercomputer room [53].

These works also quantify and qualify radiation effects on applications' output correlating the number of corrupted elements with their spatial locality and provide the mean relative error (dataset-wise) to evaluate radiation-induced error magnitude. Might it not be forgotten, as transistors get smaller, the amount of energy it takes to spontaneously flip a bit get smaller too, i.e., as exascale arrives, the number of bit-flip errors caused by radiation increases. Also, previous references about radiating computing hardware are associated to either neutron flux originated in a Lab for quantitatively estimating SE rates or demonstrating how cosmic rays actually affect computations, but what about determining the natural flux that is received in any place in the world? Hence, the evaluation of the contribution of non-thermal neutrons to the error rate of computing devices can be now calculated for the 23 exascale data centres around the World from the work carried out in the previous references and the results provided in this work.

## 3 ~~The physical model~~Atmospheric production of energetic neutrons

Cosmic rays are high-energy particles and atomic nuclei with energies from a few GeVs up to $\gtrsim 10^{20}$ eV [54]. After the pioneering works of Rossi and Auger in the 1930's [55], it is well established that cosmic rays interact with the atmosphere producing cascades of particles via radiative and decay processes, collectively known as Extensive Air Showers (EAS) [56]. Depending on the energy $E_p$ of the primary cosmic ray, an EAS could have up to $\sim 10^{10}$ particles at the moment of its maximum development. The detailed analysis of these phenomena is highly complex, as lot of different processes could be involved as more and more particles are produced. Essentially, the shower starts in the atmosphere at the first interaction point occurring at an atmospheric depth

$X_0$ that depends on primary composition and energy, where it interacts with an atomic nucleus present in the air constituents (see for example [57]). Due to the enormous difference in the energy when compared with the incoming cosmic ray, the target nuclei can be considered at rest. Since the transference at these energies of transverse momentum is small, all the increasing number of secondaries are moving towards the ground in the approximate direction of the primary. However, they can be dispersed, and the small transfer of traverse moment during radiative or decay processes produces a slow drift moving the particles away from the shower axis, and finally remain contained in a curved, thin disk known as the shower front, that moves down to the ground in the direction pointed by the initial momentum of the primary particle. The distribution of secondary particles in the shower front is axially symmetric and the particle density decrease as a power law with the distance $r$ to the shower axis, being well described by the Nishimura-Kamata-Greisen (NKG) lateral distribution function (LDF) [58].

Electromagnetic (EM) showers are initiated by photons or electrons, and most of the processes are mediated by QED interactions. These cascades are mainly ruled by two interaction channels: (i) $e^{\pm}$ Bremsstrahlung, and (ii) pair production of $e^{\pm}$. It is important to notice that both processes are coupled at high energies, as photons produce $e^{\pm}$ pairs by (ii), which in turn produce high-energy $\gamma$s by (i). These processes continue producing EM particles that could initiate new EM sub-cascades and more energy is transferred to the EM channel, which in turn produce new EM secondaries with lower energy. At some point in the cascade evolution during the propagation through the atmosphere, the rate of occurrence of radiative processes begins to decrease as the mean energy as a function of the atmospheric depth $X$, i.e., $\langle E(X)\rangle = E_p/N(X)$, where $N$ is the total number of secondaries in the cascade, drops below the critical energy $E_c$ and the ionization losses start to dominate over the radiative losses. At this point, the cascade reaches its maximum development, with a total number of particles $N_{\max} \propto E_p$ and occurring at an atmospheric depth $X_{\max} \propto \log(E_p)$. The cascade continues collectively moving down to the ground through the atmosphere, and once $X_{\max}$ is surpassed, the total number of particles $N(X)$ starts to monotonically decrease due to: (i) the radiative processes are strongly suppressed for $\langle E(X)\rangle < E_c$; and (ii) the atmospheric absorption raises as the air density increases at lower altitudes.

Instead, a hadron-initiated EAS typically produces new hadrons through fragmentation, and mesons through hadronization of the resulting fragments. Those mesons, typically $\pi^{\pm}$ and $\pi^0$, have different energy losses in the air and, most importantly, their corresponding lifetime and decay products are very different, having at the end a major impact on how these cascades develop. Almost all $\pi^0$, with a lifetime of $\tau_{\pi^0} = 8.4\times10^{-17}\,\mathrm{s}$ [59], decay very close to their production point into two energetic $\gamma$s that initiate new EM showers, transferring more energy into the EM channel. Instead, charged pions can propagate through the atmosphere down to typical altitudes of $4-6\,\mathrm{km}$ due to their longer lifetime $\tau_{\pi^\pm} = 2.6\times10^{-8}\,\mathrm{s}$ [59]. At these altitudes, they start to decay into charged muons $\mu^{\pm}$ generating the muonic component of the cascade. As the shower develops, the energy is continuously transferred to the EM and $\mu$ channels due to the decays of neutral and charged mesons. Close to the ground, $85-90\%$ of $E_p$ is at the EM channel, and the number of particles ratios typically are $10^2 : 1 : 10^{-2}$ for the EM, muon and hadronic channels respectively [60]. This latter is produced by hadronic interactions and so, it remains close to the shower axis as most of the hadrons move in a close direction to the original one, due to the reduced transference of traverse momentum produced by the leading particle effect of hadronic interactions, see e.g. [60–62]. Therefore, the hadronic component is located in a small region located close to the shower axis and is mainly composed of energetic neutrons and protons, with some light nuclei and charged pions, and small traces of other hadrons. Neutrons are mainly produced by spallation processes of protons on $^{14}$N and other nuclei in the atmosphere [63, 64]. As they are the only quasi-stable neutral hadrons present in the cascade[c0] and no ionization or radiative processes affects their propagation in the atmosphere, their evolution is only determined by elastic and quasi-elastic scattering and hadronic interactions. As explained in section 2, the energy distribution of atmospheric neutrons at different places exhibit some similarities and the main variations are related to the location and altitude of the observation site [5, 10, 64, 65]. Energy losses in the atmosphere produce two typical structures in the neutron energy spectrum: first, a single peak in the number of muons is observed at $E_n \simeq 100\,\mathrm{MeV}$, the so-called quasi-elastic peak; and a complex structure observed in the $0.1 \lesssim E_n 10\,\mathrm{MeV}$ caused by many resonances cross-sections depending on the target nuclei. At lower energies, the spectrum follows a typical $E_n^{-1}$ power law distribution with the neutron energy. The exact energy at which these spectral features appear depends on several factors, such as the altitude above sea level, geomagnetic field conditions and Solar activity, and the water vapour content in the air [66]. Due to their energy and the way they propagate through the atmosphere, these neutrons arrive at the ground with a considerable and measurable time delay with respect the primary cascade [67].

To properly simulate the cascade evolution and take into account all the involved physical processes and the propagation and tracking of up to $\sim 10^{10}$ secondary particles is a heavily demanding computing task. To do so, several tools have been developed, but the most extended and validated one is CORSIKA [68], a program for the detailed simulation of extensive air showers initiated by high-energy cosmic ray particles written in FORTRAN and continuously upgraded [69]. However, while it incorporates the possibility to select a specific atmospheric model, the values of the components of the local geomagnetic field and the altitude of the observation level, CORSIKA lacks the possibility to change those values in a dynamic way, or, most importantly, it is not possible to calculate in a direct way the secondary particles at the

---

[c0] It is possible to consider neutron as quasi-stable particles since their lifetime is several orders of magnitude larger than the characteristic time of the cascade evolution.

ground produced by the integrated flux of the primary cosmic rays. These factors are significant for the calculation of the expected background radiation at any particular site around the World and under specific and time evolving atmospheric and geomagnetic conditions.

When calculating the expected flux of secondary particles, the composition of the primary flux, the local atmospheric profile and its variations along the year, or the secular changes and the fast disturbances introduced by the Solar activity in the Earth's magnetic field have to be taken into account as they affect the number of primaries impinging the Earth's atmosphere, the evolution of the EAS in the air and the consequent flux of secondary particles at the ground.

To accomplish these tasks in a semi-autonomous way, the Latin American Giant Observatory (LAGO) [70] developed ARTI [71], a toolkit designed to effortlessly calculate and analyze the total background flux of secondaries and the corresponding detector signals produced by the atmospheric response to the primary flux of galactic cosmic rays (GCR). ARTI is publicly available at the LAGO GitHub repository [72].

LAGO operates a network of water Cherenkov detectors (WCD) at different sites in Latin America, spanning over different altitudes and geomagnetic rigidity cutoffs [73]. The geographic distribution of the LAGO sites, combined with the new electronics for control, atmospheric sensing, and data acquisition, allows the realisation of diverse astrophysics studies at a continental scale [74]. By using ARTI, LAGO is capable to obtain a better characterization of its distributed detection network and determining the sensitivity to the different phenomena studied, such as the measurement of space weather phenomena [75] or the observation of high-energy transients [76].

ARTI is a computational tool that integrates CORSIKA, Magneto-Cosmic and Geant4 with its own designed control and data analysis codes, allowing the calculation of the expected integrated flux of atmospheric radiation in any geographic location under realistic and time-evolving atmospheric and geomagnetic conditions [77]. The expected flux at the ground calculated by ARTI has been contrasted and verified with measurements performed at different astroparticles observatories, as most of them take advantage of the atmospheric muon background for the detector calibration [74, 78–81]. ARTI also has been extensively used for different applications, such as the characterization of new high altitude sites for the observation of steady gamma sources or astrophysical transients, such as the sudden occurrence of a gamma ray burst [76]; or to study the impact of space weather phenomena from ground level by using water Cherenkov detectors [74, 82, 83]; or to calculate the most statistically significant flux of high-energy muons at underground laboratories [83, 84]; to help in the assessment of active volcanoes risks in Latin America [85–88]; and even to contribute to the detection of improvised explosive devise at warfare fields in Colombia [89]. In particular, we have used ARTI to estimate the expected response of water Cherenkov detectors, commonly used for astroparticles observation, to the atmospheric neutron flux and its relation with the observation of space weather phenomena [90], and for the design of new safeguard neutron detectors for the identification of traffic of fissile materials [91, 92], which involves in both cases the calculation of the expected flux of atmospheric neutrons and the corresponding detector responses [90, 92].

Added to the intrinsic complexity of tracking all the relevant interactions of up to billions of particles with the atmosphere just for a single EAS, the atmospheric radiation at the ground level is originated by the convolution of the cascade developments of billions of cosmic rays that simultaneously impinge the Earth's atmosphere. Therefore, to obtain a statistically significant distribution of secondary particles at the ground, the time integration should be long enough to avoid statistical fluctuations[71, 74]. For example, a typical calculation of the expected number of secondaries per square metre per day for a high-latitude site involves the computation of $\sim 10^9$ EAS. For this reason, ARTI is prepared for running at high-performance computing (HPC) clusters operating with the SLURM workload manager, and in Docker containers running at virtualized cloud-based environments such as the European Open Scientific Cloud (EOSC) and capable to store and access the produced data catalogues at federated cloud storage servers [83, 93].

In this work previous calculations of the expected flux or particles at the ground level are extended, with special emphasis on the neutron flux, as one of the possible sources of silent and non-silent errors as described in previous sections. [modified paragraph for account on the energy jargon usage described in the point-by-point accompanying letter.] For doing this, we selected the minimum possible available value of the kinetic energy cuts for hadrons in CORSIKA, i.e., $E_{h_{\min}} = 5 \times 10^{-2}$ GeV, and so, for the case of neutrons, they have not tracked anymore once they reach this energy limit of $E_{n_{\min}} = 50$ MeV, that corresponds to a total energy of 989.6 MeV.

As can be inferred from the development of the showers described above, the atmosphere has a crucial role in the final distribution of particles at the ground. Any atmospheric model describes the atmosphere's main parameters (such as the atmospheric density profile) at a given time and position. So, to account for the atmospheric impacts on the cascades developments, ARTI can use four different types of atmospheric models: i) the broad MODTRAN atmospheric model [94], that assigns a general profile for different areas of the World depending on latitude and season (tropical, subtropical summer and winter, arctic or antarctic summer and winter) [94]; ii) local atmospheric profiles based on the Linsley's layers model [95] for predefined sites; iii) extract real-time atmospheric profiles from the Global Data Assimilation[c1] System (GDAS) [96] using the Linsley's model; and iv) calculate and use the typically monthly-averaged atmospheric profiles for a given location [9, 83, 93]. As we will show in the next section, by using these functionalities we can model the expected seasonal variation in the flux of secondary particles at the ground level for each one of the 23 exascale data centres shown in Figure 1.

---

[c1] Data assimilation is the adjustment of the parameters of any specific atmospheric model to the real state of the atmosphere as measured by meteorological observations

Given all the relevant primaries are charged particles and nuclei, another important factor that should be taken into account is the secular variation of the Earth's magnetic field (EMF) and its fast disturbances. These effects could be significant for the case of high latitude sites, such as the CSC Kajaani data centre in Finland. As it is described in [71, 77], ARTI incorporates specific modules to calculate the status of the EMF by using the different EMF models taken into account both the secular variation of the EMF and its disturbances.

In the next section, we show the expected flux of atmospheric radiation at the ground and its corresponding seasonal variations for the 23 exascale supercomputing centres.

# 4 Results and Discussions

## 4.1 Barometric effects in the flux of high-energy neutrons

The first step in the calculation of the expected flux at the ground is to obtain the magnetic field components $B_x$ (north component) and $B_z$ (vertical component) from the current version of the International Geomagnetic Field Reference (IGRF) model (IGRF13-2019) [97]. To reduce the impact produced by Solar activity, all the calculations were performed using the configuration of the EMF for December, 20th, 2021, as no disturbances in the magnetosphere were observed for this day.

Once the EMF components are defined, the next step is to obtain the atmospheric profiles we shall use at each of the 23 sites. For this calculation we use the monthly atmospheric profile for 2020 at each site, which was averaged from two local daily profiles extracted from the GDAS database and averaged following the ARTI methodology [9], obtaining $23 \times 12 = 276$ atmospheric profiles. A sample of the obtained density profiles and their seasonal variations can be seen in the left panel of Figure 2, where the seasonal density profiles of Los Alamos, are shown as a function of the altitude above sea level. Density profiles follow the expected seasonal variations, with denser air at the ground level in winter and a decrease in the density in the summer's warm air. In the right panel of the same Figure, expected variations along the year are shown for each atmospheric layer between ground level and 8 km asl. These variations are characterised by the minimum, maximum and one sigma deviation from the mean observed during 2020. We also included the variations observed at the High-performance Computing Center Stuttgart (HLRS, 453 m asl), the Centre de Calcul Recherche et Technology (CCRT, 94 m asl) and the Minho Advanced Computing Centre (MACC, 207 m asl) for comparative analysis. The observed differences in the density profiles along the year are small, at the level of a few per cent, but they are critical when observing the atmospheric radiation at the ground level, as the atmospheric depth at a given altitude $h_i$, defined as the integral of the atmospheric density profile within the atmospheric layer of thickness $\delta h_i$, $X(h_i) = \int_{\delta h_i} \rho(h')\, \mathrm{d}h'$, has a direct impact on the particles production, interactions and absorption at each particular layer (especially for altitudes below $\sim 15\,\text{km}$ asl), and therefore, on the final secondary particle distribution at the ground.

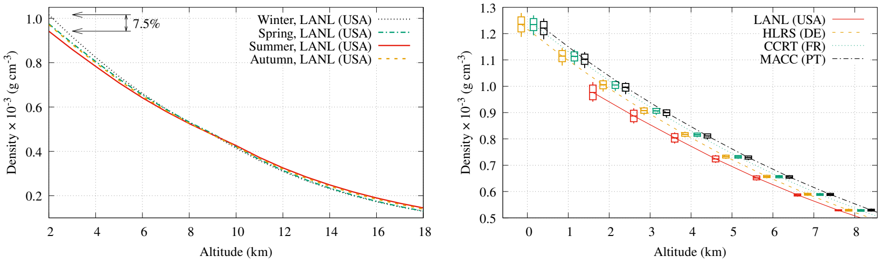


**Fig. 2** Left: The atmospheric density profiles for LANL are shown for the Winter (dotted black line), Spring (dash-dotted green line), Summer (solid red line) and Autumn (dashed yellow line) of 2020. These profiles were extracted from the GDAS database and averaged for each month. Differences of up to 7.5% can be observed in the density at the ground level in the LANL site, at an altitude of $2,125\,\text{m}$ asl. Atmospheric profiles used extends up to an altitude of $\sim 110\,\text{km}$, corresponding to the limit of the Earth's atmosphere according to Linsley's atmospheric model [95].
Right: Density variations observed at different altitudes along 2020 at LANL (solid red line), HLRS (dashed yellow line), CCRT (dotted green line) and MACC (dash-dotted black line). For each altitude between 0 and 8 km asl, candlesticks show minimums, maximums and 1-sigma deviation from the mean of the density at each atmospheric layer. See Table 1 for a summary of the characteristics of each site. Altitudes were slightly shifted for the sake of clarity

Given the stochastic nature of the development of the EAS, a large sample of showers is needed to observe these effects on the expected flux at the ground in a statistically significant manner. So the third step in the calculation is to integrate the primary spectrum $j$ to determine the total number of primary cosmic rays $N(A,Z) = \int j \; \mathrm{d}\Omega \; \mathrm{d}t \; \mathrm{d}S \; \mathrm{d}E_p$ of each relevant nucleus (identified by its atomic mass $A$ and number $Z$), which needs to be injected for a given integration time $t$, observation area $S$, solid angle interval $\Omega$, and primary energy $E_p$ range.

The cosmic ray energy spectrum ranges from GeV and up to more than $100\,\text{EeV}$ and can be very well approximated by a simple monotonically decreasing power law, i.e.,

$$\Phi(E_p, A, Z) \simeq \Phi_0(E_0, A, Z) \times (E_p/E_0)^{\alpha(E_p, A, Z)}, \qquad (1)$$

where $\Phi(E_p)$ is the expected flux of the considered primary nucleus $(A, Z)$, $\Phi_0$ is the reference flux at a certain energy $E_0$ for this particular nucleus, and $\alpha$ is the spectral index that depends on the primary energy and, while it can slightly vary from nucleus to nucleus, it can be well approximated by $\alpha \approx -3$ for the whole spectrum. Thus, we can use this property of the primary flux to limit the upper energy limit when calculating the total number of primaries for each species that need to be injected. Even more, at the PeV scale, the spectral index becomes steeper in the so-called *knee* of the cosmic ray spectrum, i.e., $\alpha \approx -3.3$ at $E_p = 4.5\,\mathrm{PeV}$ [98]. At the lowest energies, primaries are much more abundant but secondary particle production is limited and most of them are absorbed by the atmosphere before reaching ground level. For all these reasons, we limit the primary energy range for the calculation of the expected background at the ground to $E_{\min} < E_p < 10^6\,\mathrm{GeV}$, where $E_{\min} = m(Z,A)c^2 + 0.1\,\mathrm{GeV}$, being $m(A,Z)$ the mass of the injected primary [71].

The second important parameter to be considered is the total integration time $t$. While lower times reduce the total number of primaries needed to be simulated, the risk of the calculation being dominated by a statistical fluctuation increases as $t$ decreases. So, in the end, a compromise has to be taken between the saving of computing resources and the statistical significance of the calculations. While typical values for $t$ in astrophysics studies are up to a few hours [77, 82], in this case, we want to evaluate the atmospheric impact on the flux of secondary particles, and so we considered a total integration time $t$ of 1.5 days, i.e., $t = 129,600\,\mathrm{s}$ for each month at $S = 1\,\mathrm{m}^2$ in each one of the 23 sites to reduce statistical fluctuations.

Finally, since at these energies the primary flux is isotropic, we considered all the primaries following a uniform distribution in solid angle for the complete sky hemisphere around each site, i.e., $-\pi \leq \varphi \leq \pi$ and $0 \leq \theta \leq \pi/2$ for the local azimuth and zenith angle respectively.

Once the integration intervals are defined, the expected primary flux is integrated for all the relevant cosmic nuclei, obtaining $N \simeq 1.6 \times 10^9$ primaries from protons to irons ($1 \leq Z \leq 26$) for each month at each site, resulting in $\simeq 4.3 \times 10^{11}$ simulated showers in $12 \times 23 = 276$ individual runs. Calculations and analysis were done using the ARTI framework v1r9 [72], including CORSIKA v7.7402 [68] for the EAS simulations, and QGSJET-II-04 [99] and GEISHA-2002 libraries for accounting for the high- and low-energy interactions respectively. The total flux of secondaries, $\Xi_{\mathrm{All}}$, ranges from $\sim 700$ to $\sim 2,000$ particles per square metre per second, depending mainly on the EMF conditions, affecting the low energy sector of the primary flux [77]; and the atmospheric profile, having a direct influence on particle production and absorption. All the computations were performed on the ACME (equipped with Intel Gold 6138 processors) and TURGALIUM (Intel Gold 6254) clusters, demanding $\sim 450$ kCPU·hours and occupying a storage space of 1 TB for the final binary compressed files.

Typically, secondary particles are grouped into three main groups: the electromagnetic component, composed of $\gamma$s and $e^{\pm}$, the hadronic component composed of neutrons, protons, nuclei and other baryons and mesons, and the muon $\mu^{\pm}$ component. In Figure 3, the secondary momentum $p_s$ spectra are shown for these different components for the Minho Advanced Computing Centre (MACC) in Portugal, at an altitude of 200 m asl in February 2020. Several important features of the cascade development can be inferred from this Figure. At low $p_s$ values the flux is dominated by the electromagnetic (EM) component. As explained in the previous section, as the shower evolves in the atmosphere, more and more energy is transferred to the EM component via

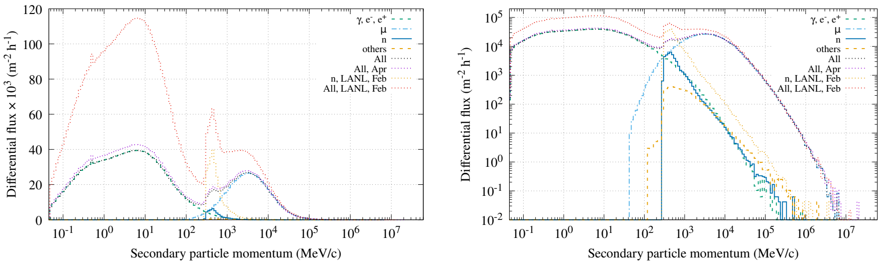


**Fig. 3** Linear-Log (left) and Log-Log (right) distribution of the momentum of the secondary particles $p_s$ expected in February 2020 at the ground level in MACC (207 m asl). The main components of the showers, i.e., the electromagnetic component (dot-dashed green line), the muons $\mu^{\pm}$ (dot-long-dashed light blue line) and the neutrons (solid blue line) and other hadrons (dot-dashed yellow line), are identifiable by their own characteristics as described in the text. Total flux for February 2020 (dotted black line) and April 2020 are also show to evidence the seasonal effects. However, the major impact is produced by the altitude above sea level, as can be seen by comparison with the neutron (dotted yellow line) and total (dotted red line) fluxes expected in February 2020 at the ground level in LANL (2,125 m asl)

particle decay and radiative processes. However, EM particles are coupled to each other through different radiative processes and, thus, EM becomes the most important component of the shower development. In the left panel of the Figure 3, a significant increase in the photon flux at the $510 - 520\,\mathrm{keV}$ energy bin is seen, corresponding to the production of $E_\gamma = 511\,\mathrm{keV}$ photons via pair annihilation $e^+e^- \to \gamma\gamma$ processes in the atmosphere.

The high-energy flux, shown in the right panel of Figure 3, is dominated by muons, charged leptons that carry the same interaction charges as $e^{\pm}$ but they are $\sim 200$ times as massive. Thus, energy losses are relatively small compared with their typical energies: $\mathrm{d}E/\mathrm{d}X$ is in the range of $2-6\,\mathrm{MeV\,cm^2\,g^{-1}}$, i.e.,$5-15\,\mathrm{MeV\,cm}$ in silicon, for muons in the $10^0 - 10^3$ GeV energy range [100]. Muons at the TeV scale, as those observed in Figure 3, possess enough energy to traverse hundreds and up to thousands of metres of rock and could be the main source for signals in muography studies [101] or background noise at underground laboratories [102]. For the same reason, it is almost impossible to shield critical devices from muons, where they could induce SET and SEU soft errors by ionization for both types of muons, plus nuclear capture only for low-energy negative muons ($\sim 50\%$ of the total muon flux). Recent works started to analyse the impact of atmospheric muons producing soft errors in different types of devices [103, 104].

Finally, at intermediate values of $p_s$, the non-thermal flux of atmospheric neutrons produces an important contribution to the total flux, especially at high-altitude sites. The impact of the altitude and local atmospheric conditions can also be seen in the same Figure, where we also included the total flux of secondaries at the MACC site but for April 2020, and at Los Alamos National Laboratory (LANL, US, 2125 m asl) for February 2020. Except for the flux of high-energy muons, which are essentially not affected by atmospheric absorption, the altitude effect is, by far, the dominant one when comparing the flux between different sites. An increase of up to 3 times in the flux of secondary particles can be observed between the MACC and LANL sites. It is also noticeable a lower but still statistically significant change in the flux originated from the change in the atmospheric profile at MACC between February and April 2020.

A denser atmosphere shall produce more absorption during the final stages of the development of the EAS, and so, a lower number of secondary particles at the ground will be observed, producing the well known anti-correlation between the atmospheric pressure and the rate of particles at the ground level [105]. The atmospheric effect can be easily observed when studying the atmospheric pressure $P(h_0)$ at the ground level[c0] and the relative temporal variations in the expected flux of secondary type $j$, i.e.,

$$\zeta_j = \frac{\Delta\Xi_j}{\overline{\Xi_j}} = \frac{\Xi_j(t)}{\overline{\Xi_j}} - 1, \tag{2}$$

where $\Xi_j(t)$ is the instantaneous flux at time $t$ and $\overline{\Xi_j}$ is the reference flux. In Figure 4, the values for $\zeta_j$ for high-energy neutrons, muons and total number of secondaries are shown together with the atmospheric pressure at the ground for the supercomputing centres of the National Energy Research Scientific Computing Center (NERSC, USA, $h_0 \simeq 210\,\mathrm{m}$ asl) and the National Supercomputing Center in Wuxi (NSCW, China, $h_0 \simeq 10\,\mathrm{m}$ asl).

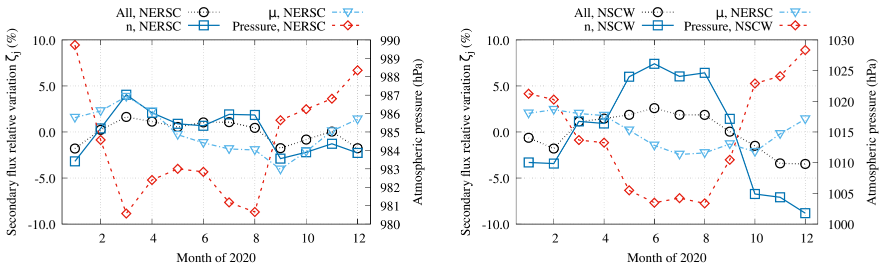


**Fig. 4** Expected relative flux variations $\zeta_j$ for neutrons (blue solid line, empty squares), muons (light blue dot-dashed line, empty triangles) and all the secondaries (black dotted line, empty circles); and the local atmospheric pressure at the ground (red dashed line, empty rhombus, right axis), are shown for each month of 2020 at the data centres of the National Energy Research Scientific Computing Center (NERSC, USA, $h_0 \simeq 210\,\mathrm{m}$ asl) and National Supercomputing Center in Wuxi (NSCW, sea level, right). As described in the text, except for muons, the anti-correlation is remarkable at all the studied sites, especially for the neutron flux

Depending on the secondary type, the atmospheric dependence could be more or less important. For example, in the right panel of Figure 4 the flux of electromagnetic particles is reduced due to the air absorption in the denser

[c0] Atmospheric pressure at a certain altitude $P(h)$ can be obtained from the atmospheric profiles by simply integrating the density profile, i.e., $P(h) = \int_{\infty}^{h} g\rho(h')\mathrm{d}h'$, where $g$ is the acceleration due to gravity.

layers of the low atmosphere and, thus, the barometric modulation in Wuxi for the total flux is not as large as for neutrons. Instead for muons, atmospheric absorption effect can be considered negligible, as can be appreciated in both panels of Figure 4, where even a correlation can be observed during part of the year at some sites. This can be explained by recalling that muons are mainly produced after charged pions decay, and so, local changes in density profiles at the muon production atmospheric depth are more relevant than the integral effect, that is related with the absorption.

On the other hand, the atmosphere has a greater impact on neutron production, propagation, moderation and absorption, as can be also seen in Figure 5, where the average, deviation and extrema in the expected number of neutrons at the ground per squared metre and hour are shown as a function of their energy for the complete year of 2020 at four sites: Los Alamos National Laboratory (LANL, 2125 m asl), High-performance Computing Center Stuttgart (HLRS, 453 m asl), Centre de Calcul Recherche et Technology (CCRT, 94 m asl) and Minho Advanced Computing Centre (MACC, 207 m asl). While the altitude effect is still dominant, the seasonal atmospheric variations have a noticeable effect on the flux of these high-energy neutrons ( [replaced previous value using the new convention (kinetic energy instead of total energy)]$E_n > 50\,\mathrm{MeV}$), even at higher energies. A detailed view of the $60-110\,\mathrm{GeV}$ neutron energy range is included, where a slightly significant deviation from the averaged power law is observed at $E_n \simeq 75)\,\mathrm{GeV}$ for all the sites. This deviation is originated from the convolution of the decreasing energy at the production level with the increase in the neutron-nucleon cross-section at the 100 GeV scale [106].

In the right panel of the same Figure, it is detailed the flux and its variations in the range $50 \leqslant E_n/\mathrm{MeV} < 450$ [(adapted to the new convention)], where the neutron flux increases by a factor of $1-2$ as the impact of the seasonal effects are enlarged. At LANL, for example, the expected neutron flux in the 100 MeV [(adapted to the new convention)] energy bin could vary by $+15\%$, from $3.5 \times 10^4$ up to $4.0 \times 10^4$ neutrons per hour per squared metre, due only to the seasonal effect.

To get a quantitative measure of the impact of the temporal variations of the atmosphere, in Figure 6 the relative variation in the flux $\zeta_j$ for different types of secondaries $j$ is shown as a function of the variation of the local atmospheric pressure at all the low-altitude ($h < 1,000\,\mathrm{m}$ asl) data centres. The barometric effect has a different impact on each type of component of the showers due to their different development in the atmosphere. This is visible in this Figure from the large differences in the observed slopes for each type of particle. The biggest impact is for neutrons and other hadrons, evidencing global variations of up to $+40\%$ for a $-4\%$ decrease in the atmospheric pressure, with the flux $\Xi_n$ ranging from $42,500$ up to $68,500$ neutrons per squared metre per hour.

It is important to notice that, besides the obvious influence of the temperature on the air density, it also impacts the single shower distribution of particles at the ground due to local changes in the lateral development of the

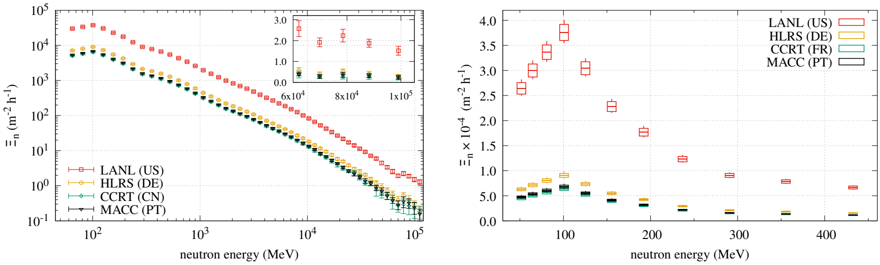


**Fig. 5** Energy distribution of the expected flux of neutrons $\Xi_n$ and its variations along 2020 at four sites: LANL (red squares), HLRS (yellow circles), CCRT (green rhombuses) and MACC (black triangles). At left, the $\Xi_n$ in the energy range $50 < E_n < 10^5$ MeV [(adapted to the new convention)] is shown as long as the 1-sigma observed variation along the year. A slight increment in the flux is observed at $E_n \simeq 80$ GeV (inset), consistent with neutron-nucleon cross-section increase at this energy range. The significant peak in $\Xi_n$, observed at $E_n \simeq 100$ MeV [(adapted to the new convention)], is detailed in the right panel, where the mean, $1\sigma$ deviations and the extrema in $\Xi_n$ for each energy bin are also shown. It can be noticed that the flux within each energy bin is not symmetric to the mean. [figure adapted to the new convention)]

cascade [107]. However, we are not interested in studying single EAS but looking for the global effect over the development of whole primary flux in the air producing the atmospheric radiation at the ground. So, given the GCR flux isotropy and uniformity at the relevant energy ranges for this study, and the stochastic (Poissonian) and self-similarity [71] nature of the atmospheric radiation production, the only effect that needs to be considered is related to the integral variation of the air density profile, i.e., the atmospheric pressure at the ground level.

Thereby, it is possible to take advantage of these effects to anticipate the expected flux of neutrons in different energy ranges at each data centre facility just by simply using the local atmospheric pressure at the ground as a tracer for the expected number of neutrons.

Local variations at each site are not as large as those shown in Figure 6, where all the observed seasonal variations with the global mean of the barometric pressure and the flux for each type of particle are shown together for the 22 low-altitude ($h < 1,000$) data centres. The slight deviation from the straight line is evidence of the exponential dependence of the flux of any secondary particle $j$, $\Xi_j(t)$, with the local barometric pressure $p(t)$ at time $t$. Nevertheless, the observed variations in the barometric pressure at every single site are significantly smaller than the global ones, and thus, they can be modelled by:

$$\zeta_j = \beta_j \Delta P, \tag{3}$$

where $\beta^j$ is the barometric coefficient for secondary $j$ and $\Delta P = P(t) - \overline{P}$ is the variation of the atmospheric pressure to the local reference $\overline{P}$. As this can be also done for different energy ranges, in this work we considered three different ones: the complete simulated energy range, [adapted to the new convention used ]$E_n \geqslant 50)$ MeV; $(50 \leqslant E_n \leqslant 1,000)$ MeV; and $(E_n > 1,000)$ MeV; respectively

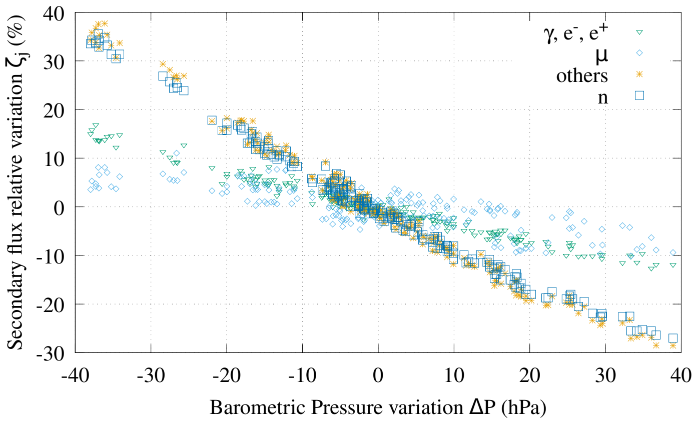


**Fig. 6** Effect of the changes in the barometric pressure in the expected flux of electromagnetic radiation (green triangles), muons (light blue rhombuses) and neutrons (blue squares) and other hadrons (yellow stars) at the ground level for low altitude data centres ($h < 1,000$ m asl). Large variations are observed in the neutron flux for slightly small changes in the pressure. Due to their different atmospheric development, each type of particle evidence a very different response to changes in the barometric pressure, as it is evidenced in the slopes of these curves. For muons, on the other hand, local variations are most influenced by changes in the atmospheric profile at muon production layers than the barometric pressure. The exponential atmospheric dependence of the flux is visible in the slight deviation from a straight line.

labelled as $i = 0$, $i = 1$ and $i = 2$. For the sake of clarity and given we are mainly focused on the neutron flux, we can obviate the subscript $n$ and so, equation (3) could be written as:

$$\zeta_i = \beta_i \Delta P, \tag{4}$$

where now the subscript $i$ refers to the corresponding neutron energy range $i = 0, 1, 2$ described above. The obtained results for all the sites are compiled in the table 2. It is important to notice that slight differences could be observed in both the total flux and the barometric coefficients at sites with similar altitudes due to differences of the atmospheric profiles and their impacts on the neutron flux.

From these values and using (4), it is possible to estimate the expected flux of high-energy neutrons and its variations at each site just by measuring the local atmospheric pressure, since $\beta_i$ corresponds to the relative decrease (increase) in the neutron flux for an 1 hPa increase (decrease) in the local barometric pressure. For example, from the second and the fourth column of Table 2, the reference atmospheric pressure and the global barometric coefficient for the site of Los Alamos (LANL) are $\overline{P} = 777$ hPa and $\beta_0 = -9.2 \times 10^{-3}$ hPa$^{-1}$ respectively. Therefore, on a typical sunny day at LANL, when the barometric pressure should be higher than usual, say, $P(t) = 779$ hPa, a reduction of $\beta_0(P(t) - \overline{P}) = -9.2 \times 10^{-3}$ hPa$^{-1} \times 2$ hPa =

**Table 1** Altitude, latitude and longitude of the 23 new exascale facilities, including the total, averaged flux $\Xi_{All}$ in cm$^{-2}$ hour$^{-1}$, of all the secondary particles, muons and neutrons with $E_n \geqslant 50\,\mathrm{MeV}$ [adapted for the new convention)] expected at the ground level in each site.

| **Supercomputing centre** | Country | Alt. | Lat. | Lon. | Code | $\Xi_{All} \times 10^{-6}$ | $\Xi_n \times 10^{-4}$ | $\Xi_\mu \times 10^{-5}$ |
|---|---|---|---|---|---|---|---|---|
| Los Alamos National Laboratory | USA | 2,125 | 35.85 | -106.29 | LANL | $(6.70 \pm 0.19)$ | $(26.4 \pm 1.1)$ | $(9.7 \pm 0.3)$ |
| National University of Defense Technology | China | 750 | 27.93 | 107.71 | NUDT | $(3.43 \pm 0.08)$ | $(7.7 \pm 0.4)$ | $(7.4 \pm 0.1)$ |
| Centro de Investigaciones Energéticas, Medioambientales y Tecnológicas | Spain | 700 | 40.50 | -3.67 | MAD | $(3.47 \pm 0.06)$ | $(7.7 \pm 0.3)$ | $(7.7 \pm 0.1)$ |
| Sofia Tech Park | Bulgaria | 565 | 42.67 | 23.37 | SOFIA | $(3.32 \pm 0.05)$ | $(6.9 \pm 0.2)$ | $(7.6 \pm 0.1)$ |
| Leibniz Supercomputing Centre | Germany | 471 | 48.26 | 11.67 | LRZ | $(3.21 \pm 0.07)$ | $(6.4 \pm 0.3)$ | $(7.5 \pm 0.1)$ |
| High-performance Computing Center Stuttgart | Germany | 453 | 48.74 | 9.10 | HLRS | $(3.19 \pm 0.07)$ | $(6.3 \pm 0.3)$ | $(7.5 \pm 0.1)$ |
| Institute of Information Science | Slovenia | 280 | 46.56 | 15.65 | IZUM | $(2.97 \pm 0.05)$ | $(5.3 \pm 0.2)$ | $(7.3 \pm 0.1)$ |
| uxConnec's Data Center DC2 | Luxembourg | 275 | 49.79 | 6.09 | DC2 | $(2.98 \pm 0.07)$ | $(5.3 \pm 0.2)$ | $(7.3 \pm 0.1)$ |
| IT4 Innovations National Supercomputing Center | Czechia | 261 | 49.84 | 18.16 | IT4 | $(2.96 \pm 0.07)$ | $(5.2 \pm 0.2)$ | $(7.3 \pm 0.1)$ |
| Oak Ridge National Laboratory | USA | 250 | 35.93 | -84.31 | ORNL | $(2.85 \pm 0.05)$ | $(4.8 \pm 0.1)$ | $(7.0 \pm 0.1)$ |
| Argonne National Laboratory | USA | 214 | 41.72 | -87.98 | ANL | $(2.87 \pm 0.04)$ | $(4.9 \pm 0.2)$ | $(7.1 \pm 0.2)$ |
| National Energy Research Scientific Computing Center | USA | 210 | 37.88 | -122.25 | NERSC | $(2.85 \pm 0.03)$ | $(4.8 \pm 0.1)$ | $(7.0 \pm 0.2)$ |
| Minho Advanced Computing Centre | Portugal | 207 | 41.56 | -8.40 | MACC | $(2.85 \pm 0.05)$ | $(4.8 \pm 0.2)$ | $(7.1 \pm 0.1)$ |
| Lawrence Livermore National Laboratory | USA | 188 | 37.69 | -121.70 | LLNL | $(2.84 \pm 0.03)$ | $(4.7 \pm 0.1)$ | $(7.0 \pm 0.1)$ |
| Datacenter CSC Kajaani | Finland | 128 | 64.23 | 27.70 | CSCF | $(2.94 \pm 0.08)$ | $(5.1 \pm 0.3)$ | $(7.4 \pm 0.2)$ |
| Barcelona Supercomputing Center | Spain | 100 | 41.39 | 2.12 | BSC | $(2.76 \pm 0.05)$ | $(4.3 \pm 0.2)$ | $(7.0 \pm 0.1)$ |
| Jülich Supercomputing Centre | Germany | 100 | 50.92 | 6.54 | JSC | $(2.79 \pm 0.05)$ | $(4.5 \pm 0.2)$ | $(7.1 \pm 0.1)$ |
| Poznan Supercomputing and Networking Center | Poland | 100 | 52.41 | 16.92 | PSNC | $(2.80 \pm 0.05)$ | $(4.5 \pm 0.2)$ | $(7.1 \pm 0.1)$ |
| Centre de Calcul Recherche et Technologie | France | 94 | 48.60 | 2.20 | CCRT | $(2.78 \pm 0.05)$ | $(4.4 \pm 0.2)$ | $(7.1 \pm 0.1)$ |
| Bologna Technopole | Italy | 40 | 44.52 | 11.36 | BOLT | $(2.72 \pm 0.06)$ | $(4.2 \pm 0.2)$ | $(7.0 \pm 0.1)$ |
| National Supercomputer Center in Guangzhou | China | 10 | 23.07 | 113.39 | NSCG | $(2.58 \pm 0.05)$ | $(3.7 \pm 0.2)$ | $(6.4 \pm 0.1)$ |
| National Supercomputing Center in Wuxi | China | 10 | 31.57 | 120.30 | NSCW | $(2.60 \pm 0.05)$ | $(3.8 \pm 0.2)$ | $(6.6 \pm 0.1)$ |
| RIKEN Center for Computational Science | Japan | 10 | 34.68 | 135.22 | RCCS | $(2.63 \pm 0.04)$ | $(3.9 \pm 0.1)$ | $(6.7 \pm 0.2)$ |

**Table 2** Reference pressure $\overline{P}$ and neutron flux $\overline{\Xi_i}$, and barometric coefficients $\beta_i$ (in $\text{hPa}^{-1}$) for the $i$-th energy range at the 23 exascale facilities. With these values, it is possible to calculate using equation (4) the local variations in the flux of neutrons just from the local barometric pressure (use $\overline{\Xi_0}$ and $\beta_0$ for the neutron flux with $E_n \geqslant 50\,\text{MeV}$ [(adapted to the new convention)]). Pressure is given in hPa and fluxes are given in $\text{m}^{-2}\,\text{hour}^{-1}$.

| **Site** | Alt. | $\overline{P}$ | $\overline{\Xi_0}$ | $\beta_0$ | $\overline{\Xi_1}$ | $\beta_1$ | $\overline{\Xi_2}$ | $\beta_2$ |
|---|---|---|---|---|---|---|---|---|
| LANL | 2,125 | 777 | 26.4 | -9.2 | 25.7 | -9.2 | 7.0 | -9.7 |
| NUDT | 750 | 927 | 7.7 | -6.9 | 7.5 | -6.9 | 1.9 | -7.1 |
| MAD | 700 | 927 | 7.7 | -7.7 | 7.5 | -7.7 | 1.9 | -7.5 |
| SOFIA | 565 | 940 | 6.9 | -6.8 | 6.8 | -6.8 | 1.7 | -7.0 |
| LRZ | 471 | 950 | 6.4 | -7.8 | 6.2 | -7.8 | 1.6 | -7.9 |
| HLRS | 453 | 952 | 6.3 | -8.1 | 6.1 | -8.1 | 1.6 | -8.3 |
| IZUM | 280 | 974 | 5.3 | -6.8 | 5.2 | -6.8 | 1.3 | -7.0 |
| DC2 | 275 | 973 | 5.3 | -7.9 | 5.2 | -7.9 | 1.3 | -8.2 |
| IT4 | 261 | 975 | 5.2 | -7.6 | 5.1 | -7.6 | 1.3 | -7.8 |
| ORNL | 250 | 984 | 4.8 | -7.9 | 4.7 | -7.9 | 1.2 | -8.2 |
| ANL | 214 | 983 | 4.9 | -8.8 | 4.8 | -8.8 | 1.2 | -9.6 |
| NERSC | 210 | 984 | 4.8 | -7.0 | 4.7 | -7.0 | 1.2 | -6.4 |
| MACC | 207 | 986 | 4.8 | -7.6 | 4.6 | -7.6 | 1.2 | -8.0 |
| LLNL | 188 | 987 | 4.7 | -7.1 | 4.6 | -7.1 | 1.1 | -7.7 |
| CSCF | 128 | 978 | 5.1 | -8.4 | 5.0 | -8.4 | 1.3 | -8.7 |
| BSC | 100 | 997 | 4.3 | -7.7 | 4.2 | -7.7 | 1.1 | -7.9 |
| JSC | 100 | 993 | 4.5 | -7.7 | 4.4 | -7.8 | 1.1 | -7.7 |
| PSNC | 100 | 993 | 4.5 | -7.3 | 4.4 | -7.3 | 1.1 | -7.7 |
| CCRT | 94 | 995 | 4.4 | -8.2 | 4.3 | -8.2 | 1.1 | -8.4 |
| BOLT | 40 | 1002 | 4.2 | -7.2 | 4.1 | -7.2 | 1.0 | -7.2 |
| NSCG | 10 | 1015 | 3.7 | -6.9 | 3.6 | -6.9 | 0.9 | -7.2 |
| NSCW | 10 | 1014 | 3.8 | -6.4 | 3.7 | -6.4 | 0.9 | -6.5 |
| RCCS | 10 | 1010 | 3.9 | -6.7 | 3.8 | -6.7 | 0.9 | -6.7 |
| | | | $\times 10^4$ | $\times 10^{-3}$ | $\times 10^4$ | $\times 10^{-3}$ | $\times 10^4$ | $\times 10^{-3}$ |

$-1.84 \times 10^{-2} \simeq -2\%$ in the $E_n \gtrsim 50\,\text{MeV}$ [(adapted to the new convention)] neutron flux shall be expected. Thunderstorms, on the other hand, are preceded by a drop in the atmospheric pressure of several hPa in a few hours, with typical drop rates of at least $-1\,\text{hPa}\,\text{h}^{-1}$. So, at sea level, the barometric pressure could be as low as $1{,}002\,\text{hPa}$, or even less, during a thunderstorm. Thus, for example, during the preclude of a thunderstorm at the RIKEN Center for Computational Science (RCCS) in Kobe, Japan, where the average atmospheric pressure is $\overline{P} = 1{,}010\,\text{hPa}$, an increase of $\sim 6\%$ in the flux of neutrons with energies above $50\,\text{MeV}$ [(adapted to the new convention)] could be expected[c4], and the situation could be even worst when considering the effective moderation of neutrons produced by rain. As a consequence, an increase (decrease) in the flux of high-energy neutrons will result in a similar increase (decrease) in the probability of errors produced in the supercomputer.

For muons, local changes in the profile at muon production depth are the dominant effect. Expected average muon flux at each data centre for $E_\mu \geqslant 15\,\text{MeV}$ [(adapted to the new convention)] are also included in Table 1.

Space weather phenomena, such as the disturbances of the magnetosphere produced by the passage of an interplanetary coronal mass ejection (iCME)

[c4]Since, according to Table 2 for RCSS: $\beta_0(P(t) - \overline{P}) = -6.7 \times 10^{-3}\ \text{hPa}^{-1} \times (-8)\ \text{hPa} \simeq 6\%$.

by Earth [108], also impacts the flux of high-energy neutrons and for this reason, atmospheric neutrons have been used since decades ago to monitor Solar activity [109]. These phenomena are observed as decreases in the total flux of atmospheric neutrons, where reductions of up to 35% could be expected for $E_n \simeq 100\,\mathrm{MeV}$ [(adapted to the new convention)] neutrons during severe geomagnetic storms [77], and some astroparticle observatories, such as LAGO, are focused on enhancing their neutron detection capabilities [90, 91].

These scenarios are important when anticipating possible errors associated with the flux of high-energy neutrons at supercomputer centres, as it will be discussed in subsection 4.2.

## 4.2 High-energy neutrons modulations and soft error rates at supercomputers

A typical magnitude used to describe the device performance in terms of its sensitivity to radiation is the FIT (failures-in-time) rate, i.e., the number of observed failures of a certain (or any) kind in $10^9$ (one billion) hours of device operation, and so, the total FIT is just the sum of each kind of failure: $\mathrm{FIT} = \sum_k^N \mathrm{FIT}_k$. From this definition, the MTBF measured in hours is just the reciprocal of FIT times $10^9$:

$$\mathrm{MTBF} = \frac{10^9}{\mathrm{FIT}}. \tag{5}$$

It is possible to obtain the FIT rate from the effective cross-section $\sigma_{\mathrm{err}}$, as it is just an effective measure of the probability that a neutron triggers a certain type of error in a device, and it is typically expressed in units of area ($\mathrm{cm}^2$) [47]. Thus, in general,

$$\mathrm{FIT}_{\mathrm{err}} = 10^5\ \Xi\ \sigma_{\mathrm{err}}, \tag{6}$$

when the flux $\Xi$ is expressed in units of $\mathrm{m}^{-2}\,\mathrm{h}^{-1}$. Then, by combining this result with equations (2) and (4) for neutrons:

$$\mathrm{FIT}_{\mathrm{err}}(t) = 10^5\ \sigma_{\mathrm{err}}\ \overline{\Xi_i}\left[1 + \beta_i \left(P(t) - \overline{P}\right)\right], \tag{7}$$

in the $i$-th neutron energy range, for pressure expressed in hPa and $\sigma$ in $\mathrm{cm}^2$.

Oliveira *et al.* [52, 53] irradiate different types of commercial off-the-shelf (COTS) devices by exposing them to neutron beams in energy scales from thermal to 1 GeV, obtaining the device sensitivity to neutrons measured through the identification of unrecoverable errors (DUE) or SDC in APUs (CPUs+GPUs integrated in the same device), FPGAs and DDR memories. Unfortunately, they only present cross-sections "relative to the lowest one measure for each vendor to prevent the leakage of business-sensitive data" [53]. However, is it possible to see that, for all the tested devices, thermal neutron cross-sections are far for being negligible[53], but in most cases they are still considerable smaller than the corresponding effective cross-section of high-energy neutrons (the observed differences are up to one order of magnitude for APUs). Similar conclusions can be obtained from Figure 6 of [52], where it is possible to observe that, in presence of the nominal atmospheric flux of high-energy neutrons ($E_n > 10\,\text{MeV}$), the FIT rates are totally dominated by them.

As mentioned in section 2, Tiwari *et al.* [50], analyzed the error logs of two GPU supercomputing facilities: the Titan supercomputer at the Oak Ridge National Laboratory (ORNL), consisting of 18,688 K20X GPUs; and of the Moonlight GPGPU cluster at Los Alamos National Laboratory (LANL), consisting of 616 M2090 GPGPUs. By exposing K20X GPU to the ISIS and LANSCE white neutron sources, that emulate the atmospheric neutron flux in the $10 < E_n < 750\,\text{MeV}$ energy range [110], they were able to obtain the SDC and program crashes effective cross-sections $\sigma_{\text{err}}$, that are compiled in the table 2 of [50] and can be averaged obtaining $\sigma_{\text{SDC}} = (4.8 \pm 0.4) \times 10^{-7}\,\text{cm}^2$ and $\sigma_{\text{crash}} = (2.7 \pm 0.2) \times 10^{-7}\,\text{cm}^2$ respectively. While the energy ranges of the neutron sources used for the irradiation of the K20s devices are lower than the complete energy range simulated in this work, it is possible to assume that the neutron-error cross-sections in the energy range $E_n > 1,000\,\text{MeV}$ [(adapted to the new convention)] should not be far from the reported values. Moreover, at these high energies, the flux is considerably lower than in the $50 \leq E_n/\text{MeV} \leq 1,000$ energy range, and so the error rates will be dominated by the flux within this range. [(adapted to the new convention)] Therefore, following equation (7) and using the tabulated values for $\overline{P}$, $\overline{\Xi_1}$ and $\beta_1$ for the ORNL site, the expected $\text{FIT}_{\text{SDC}}$ rate when the atmospheric pressure drops by, say, $-5\,\text{hPa}$ respect to the barometric reference pressure, should be[c2] of $\text{FIT}_{\text{SDC}} \sim 2,300$, and so, from equation (5), the corresponding MTBF for the whole Titan supercomputer should be of $\simeq 23$ hours, i.e., about 1 silent error per day due to the expected flux of neutrons with $50 < E_n < 1,000\,\text{MeV}$ [(adapted to the new convention)] when the atmospheric pressure drops by $-5\,\text{hPa}$.

Once the expected flux of neutrons was determined for each site, calculation of effective flux at computing devices, including CPUs, GPUs, APUs, storage and memories, have to take into account the geometry and materials of computing racks, buildings and other infrastructures in the surroundings, even, on the supercomputing cooling system, especially those using water or any other aqueous solutions as coolants. All these components will have a profound impact in the flux of high-energy neutrons, producing thermal and epi-thermal neutrons having different cross-sections with the materials used for making the different types of devices available in any data centre.

[c2] $FIT_{\text{SDC}} = (10^5)(4.7 \times 10^4)(4.8 \times 10^{-7})[1 + (-7.9 \times 10^{-3})(979 - 984)] = 2,345 \simeq 2,300$ failures in $10^9$ device·hours of operation.

As a final remark, given the linearity of equations (4) and (6), it is easy to see that the relative variation of the FIT rates,

$$\psi_{\text{err}} = \frac{\text{FIT}_{\text{err}}(t)}{\overline{\text{FIT}}_{\text{err}}} - 1,$$

where $\overline{\text{FIT}_{\text{err}}}$ is the reference $\text{FIT}_{\text{err}}$ rate at the site, is equal to the relative variation $\zeta$ of the high-energy neutron flux, i.e, $\psi = \zeta$, and so:

$$\psi_{\text{err}} = \beta \Delta P, \tag{8}$$

that is, the FIT rate associated with the flux of high-energy neutrons at each site should evidence a small anti-correlation ($\beta > -1\%$) with the local changes in the barometric pressure that increases with the altitude of the supercomputing centre.

# 5 Conclusions

In this work we presented the calculation of the expected flux of atmospheric neutrons and their seasonal variations at each one of the 23 future sites of the next generation of exascale supercomputing facilities. This was done by simulating the interaction of the measured galactic cosmic rays flux and including real atmospheric conditions at each site using the state-of-the-art techniques and codes heavily used, tested and validated in the astroparticle physics community.

By using real atmospheric profiles, extracted from the GDAS database and averaged to obtain the atmosphere conditions for each month of 2020, the expected flux of high-energy neutrons with $E_n \gtrsim 50\,\text{MeV}$ [(adapted to the new convention)] and its seasonal variations at each exascale supercomputing centre were obtained and parametrised. The dependence on the total flux of particles and neutron flux with the atmospheric pressure was observed and the barometric pressure coefficient for neutrons at different energy ranges were obtained and they are summarised in Table 2. The reported barometric coefficients, $\beta_i$ corresponds to the relative change in the expected flux in different energy ranges when the atmospheric pressure changes by $\pm 1\,\text{hPa}$. The provided information makes it possible to easily estimate the expected flux of neutrons under different atmospheric conditions (equation (4)) and to evaluate the corresponding FIT rates of silent errors due to high-energy neutrons (equation (7)) and its relative seasonal variations (equation (8)). This can be done by using the instantaneous barometric pressure that can be easily measured at each facility, being a simple and direct way to anticipate potential silent and non-silent errors that could appear during critical calculations that could be performed soon at the next generation of exascale supercomputing facilities.

To avoid the intrinsic limitation of CORSIKA for low energy neutrons, we are currently developing a special module in ARTI, based on FLUKA [111], to extend current calculations down to the meV neutron energy scale. Extensions of the atmospheric flux simulations using real atmospheres presented here but including other effects such as the rain, that could double the thermal neutron flux at the ground as water droplets acts as neutrons moderators, and the corresponding Geant4 [112] simulations of neutron moderation in infrastructures are being considered and will be published as a follow-up of the analysis presented here.

## Declarations

**Ethics approval.** Not applicable.

**Consent to participate.** All authors agreed to participate.

**Consent for publication.** Not applicable.

**Availability of data and materials.** The datasets generated and analysed during the current study are available in the Zenodo repository, 10.5281/zenodo.6721615. The ARTI code is available in the LAGO GitHub repository: github.com/lagoproject/arti.

**Competing interests.** Not applicable.

**Funding.** This work has been partially funded by the co-funded Spanish Ministry of Science and Innovation project CODEC-OSE (RTI2018-096006-B-I00) with European Regional Development Fund (ERDF) funds, by the co-funded European Union Horizon 2020 research and innovation Programme project EOSC-SYNERGY (grant agreement No 857647), and by the co-funded Comunidad de Madrid project CABAHLA-CM (S2018/TCS-4423).

**Authors' contributions.** All authors contributed equally to this work and reviewed the manuscript.

**Acknowledgments.** This work has been partially funded by the co-funded Spanish Ministry of Science and Innovation project CODEC-OSE (RTI2018-096006-B-I00) with European Regional Development Fund (ERDF) funds, by the co-funded European Union Horizon 2020 research and innovation Programme project EOSC-SYNERGY (grant agreement No 857647), and by the co-funded Comunidad de Madrid project CABAHLA-CM (S2018/TCS-4423). Also, this work was partially supported by the computing facilities (Turgalium) of Extremadura Research Centre for Advanced Technologies (CETA-CIEMAT), funded by the ERDF too.

The authors are grateful to Antonio Juan Rubio-Montero and Angelines Alberto-Morillas from CIEMAT, Alfonso Pardo-Diaz from CETA/CIEMAT and Iván Sidelnik from CNEA for their continuous support and fruitful discussions.

HA thanks Rafael Mayo-García for his warm welcome and continuous support during his stay at CIEMAT in Madrid, Spain.

# References

[1] Dongarra, J., Beckman, P., Aerts, P., Cappello, F., Lippert, T., Matsuoka, S., Messina, P., Moore, T., Stevens, R., Trefethen, A., *et al.*: The international exascale software project: a call to cooperative action by the global high-performance community. The International Journal of High Performance Computing Applications **23**(4), 309–322 (2009)

[2] Snir, M., Wisniewski, R.W., Abraham, J.A., Adve, S.V., Bagchi, S., Balaji, P., Belak, J., Bose, P., Cappello, F., Carlson, B., *et al.*: Addressing failures in exascale computing. The International Journal of High Performance Computing Applications **28**(2), 129–173 (2014)

[3] Rodríguez-Pascual, M., Cao, J., Moríñigo, J.A., Cooperman, G., Mayo-García, R.: Job migration in hpc clusters by means of checkpoint/restart. The Journal of Supercomputing **75**(10), 6517–6541 (2019)

[4] Yeom, J.-S., Nikolopoulos, D.S.: Strider: Runtime support for optimizing strided data accesses on multi-cores with explicitly managed memories. In: SC'10: Proceedings of the 2010 ACM/IEEE International Conference for High Performance Computing, Networking, Storage and Analysis, pp. 1–11 (2010). IEEE

[5] Gordon, M.S., Goldhagen, P., Rodbell, K.P., Zabel, T.H., Tang, H.H.K., Clem, J.M., Bailey, P.: Measurement of the flux and energy spectrum of cosmic-ray induced neutrons on the ground, vol. 51, pp. 3427–3434 (2004). https://doi.org/10.1109/TNS.2004.839134

[6] Oliveira, D.A.G.D., Pilla, L.L., Hanzich, M., Fratin, V., Fernandes, F., Lunardi, C., Cela, J.M., Navaux, P.O.A., Carro, L., Rech, P.: Radiation-induced error criticality in modern hpc parallel accelerators, pp. 577–588 (2017). https://doi.org/10.1109/HPCA.2017.41

[7] Rojdev, K., Koontz, S., Reddell, B., Atwell, W., Boeder, P.: Comparison and validation of fluka and hzetrn as tools for investigating the secondary neutron production in large space vehicles. (2015). https://doi.org/10.2514/6.2015-4601

[8] Dilillo, L., Wrobel, F., Galliere, J.-M., Saigne, F.: Neutron detection through an sram-based test bench, pp. 64–69 (2009). https://doi.org/10.1109/IWASI.2009.5184769. http://ieeexplore.ieee.org/document/5184769/

[9] Grisales-Casadiegos, J., Sarmiento-Cano, C., Núñez, L.A.: Impact of global data assimilation system atmospheric models on astroparticle showers. Canadian Journal of Physics **100**(3), 152–157 (2022)

[10] Infantino, A., Blackmore, E.W., Brugger, M., Alía, R.G., Stukel, M., Trinczek, M.: Fluka monte carlo assessment of the terrestrial muon flux at low energies and comparison against experimental measurements. Nuclear Instruments and Methods in Physics Research, Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **838**, 109–116 (2016). https://doi.org/10.1016/j.nima.2016.09.012

[11] Cappello, F., Al, G., Gropp, W., Kale, S., Kramer, B., Snir, M.: Toward exascale resilience: 2014 update. Supercomputing Frontiers and Innovations: an International Journal **1**(1), 5–28 (2014)

[12] Gainaru, A., Cappello, F., Snir, M., Kramer, W.: Failure prediction for hpc systems and applications: Current situation and open issues. The International journal of high performance computing applications **27**(3), 273–282 (2013)

[13] Canal, R., Hernandez, C., Tornero, R., Cilardo, A., Massari, G., Reghenzani, F., Fornaciari, W., Zapater, M., Atienza, D., Oleksiak, A., *et al.*: Predictive reliability and fault management in exascale systems: State of the art and perspectives. ACM Computing Surveys (CSUR) **53**(5), 1–32 (2020)

[14] Fault-tolerance Techniques for High-performance Computing. Springer, Switzerland (2015). https://doi.org/10.1007/978-3-319-20943-2

[15] Artho, C., Suzaki, K., Hagiya, M., Leungwattanakit, W., Potter, R., Platon, E., Tanabe, Y., Weitl, F., Yamamoto, M.: Using checkpointing and virtualization for fault injection. International Journal of Networking and Computing **5**(2), 347–372 (2015)

[16] Samfass, P., Weinzierl, T., Reinarz, A., Bader, M.: Doubt and redundancy kill soft errors—towards detection and correction of silent data corruption in task-based numerical software. In: 2021 IEEE/ACM 11th Workshop on Fault Tolerance for HPC at eXtreme Scale (FTXS), pp. 1–10 (2021). IEEE

[17] Rodríguez, G., Martín, M.J., González, P., Tourino, J., Doallo, R.: Cppc: a compiler-assisted tool for portable checkpointing of message-passing applications. Concurrency and Computation: Practice and Experience **22**(6), 749–766 (2010)

[18] Bosilca, G., Bouteiller, A., Brunet, E., Cappello, F., Dongarra, J., Guermouche, A., Herault, T., Robert, Y., Vivien, F., Zaidouni, D.: Unified model for assessing checkpointing protocols at extreme-scale. Concurrency and Computation: Practice and Experience **26**(17), 2772–2791 (2014)

[19] Ferreira, K., Stearley, J., Laros III, J.H., Oldfield, R., Pedretti, K., Brightwell, R., Riesen, R., Bridges, P.G., Arnold, D.: Evaluating the viability of process replication reliability for exascale systems. In: Proceedings of 2011 International Conference for High Performance Computing, Networking, Storage and Analysis, pp. 1–12 (2011)

[20] Yu, J., Jian, D., Wu, Z., Liu, H.: Thread-level redundancy fault tolerant cmp based on relaxed input replication. In: 2011 6th International Conference on Computer Sciences and Convergence Information Technology (ICCIT), pp. 544–549 (2011). IEEE

[21] Fagg, G.E., Dongarra, J.J.: Ft-mpi: Fault tolerant mpi, supporting dynamic applications in a dynamic world. In: European Parallel Virtual Machine/message Passing Interface Users' Group Meeting, pp. 346–353 (2000). Springer

[22] Bland, W., Du, P., Bouteiller, A., Herault, T., Bosilca, G., Dongarra, J.J.: Extending the scope of the checkpoint-on-failure protocol for forward recovery in standard mpi. Concurrency and computation: Practice and experience **25**(17), 2381–2393 (2013)

[23] Losada, N., González, P., Martín, M.J., Bosilca, G., Bouteiller, A., Teranishi, K.: Fault tolerance of mpi applications in exascale systems: The ulfm solution. Future Generation Computer Systems **106**, 467–481 (2020)

[24] Garg, R., Price, G., Cooperman, G.: Mana for mpi: Mpi-agnostic network-agnostic transparent checkpointing. In: Proceedings of the 28th International Symposium on High-Performance Parallel and Distributed Computing, pp. 49–60 (2019)

[25] Xu, Y., Zhao, Z., Garg, R., Khetawat, H., Hartman-Baker, R., Cooperman, G.: Mana-2.0: A future-proof design for transparent checkpointing of mpi at scale. arXiv preprint arXiv:2112.05858 (2021)

[26] Egwutuoha, I.P., Levy, D., Selic, B., Chen, S.: A survey of fault tolerance mechanisms and checkpoint/restart implementations for high performance computing systems. The Journal of Supercomputing **65**(3), 1302–1326 (2013)

[27] Fang, A., Chien, A.A.: Abfr: convenient management of latent error resilience using application knowledge. In: Proceedings of the 27th International Symposium on High-Performance Parallel and Distributed Computing, pp. 27–39 (2018)

[28] Ansel, J., Arya, K., Cooperman, G.: Dmtcp: Transparent checkpointing for cluster computations and the desktop. In: 2009 IEEE International Symposium on Parallel & Distributed Processing, pp. 1–12 (2009). IEEE

[29] Hargrove, P.H., Duell, J.C.: Berkeley lab checkpoint/restart (blcr) for linux clusters. In: Journal of Physics: Conference Series, vol. 46, p. 067 (2006). IOP Publishing

[30] Daly, J.T.: A higher order estimate of the optimum checkpoint interval for restart dumps. Future generation computer systems **22**(3), 303–312 (2006)

[31] Aupy, G., Benoit, A., Cavelan, A., Fasi, M., Robert, Y., Sun, H., Uçar, B.: Coping with silent errors in hpc applications. In: Emergent Computation, pp. 269–292. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-46376-6_11

[32] Benoit, A., Cavelan, A., Cappello, F., Raghavan, P., Robert, Y., Sun, H.: Coping with silent and fail-stop errors at scale by combining replication and checkpointing. Journal of Parallel and Distributed Computing **122**, 209–225 (2018)

[33] Yeh, Y.C.: Triple-triple redundant 777 primary flight computer. In: 1996 IEEE Aerospace Applications Conference. Proceedings, vol. 1, pp. 293–307 (1996). IEEE

[34] Moríñigo, J.A., Rodríguez-Pascual, M., Mayo-García, R.: On the modelling of optimal coordinated checkpoint period in supercomputers. The Journal of Supercomputing **75**(2), 930–954 (2019)

[35] Avron, H., Maymounkov, P., Toledo, S.: Blendenpik: Supercharging lapack's least-squares solver. SIAM Journal on Scientific Computing **32**(3), 1217–1236 (2010)

[36] Baboulin, M., Li, X.S., Rouet, F.-H.: Using random butterfly transformations to avoid pivoting in sparse direct methods. In: International Conference on High Performance Computing for Computational Science, pp. 135–144 (2014). Springer

[37] Howell, G.W., Baboulin, M.: Iterative solution of sparse linear least squares using lu factorization. In: Proceedings of the International Conference on High Performance Computing in Asia-Pacific Region, pp. 47–53 (2018)

[38] Dongarra, J., Tomov, S., Luszczek, P., Kurzak, J., Gates, M., Yamazaki, I., Anzt, H., Haidar, A., Abdelfattah, A.: With extreme computing, the rules have changed. Computing in Science & Engineering **19**(3), 52–62 (2017)

[39] Chetverushkin, B.N., Yakobovskiy, M.V., Kornilina, M.A., Semenova, A.V.: Numerical algorithms for hpc systems and fault tolerance. In:

International Conference on Parallel Computational Technologies, pp. 34–44 (2019). Springer

[40] Moríñigo, J.A., Bustos, A., Mayo-García, R.: Error resilience of three gmres implementations under fault injection. The Journal of Supercomputing, 1–28 (2021)

[41] Cavelan, A., Ciorba, F.M.: Algorithm-based fault tolerance for parallel stencil computations. In: 2019 IEEE International Conference on Cluster Computing (CLUSTER), pp. 1–11 (2019). IEEE

[42] Normand, E.: Single event upset at ground level. IEEE Transactions on Nuclear Science **43**, 2742–2750 (1996). https://doi.org/10.1109/23.556861

[43] Ziegler, J.F., Nelson, M.E., Shell, J.D., Peterson, R.J., Gelderloos, C.J., Muhlfeld, H.P., Montrose, C.J.: Cosmic ray soft error rates of 16-mb dram memory chips. IEEE Journal of Solid-state circuits **33**(2), 246–252 (1998)

[44] Michalak, S., Harris, K., Hengartner, N., Takala, B., Wender, S.: Using the lansce irradiation facility to predict the number of fatal soft errors in one of the world's fastest supercomputers. Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms **241**(1-4), 414–418 (2005)

[45] Presti, D.L., Medina, N.H., Guazzelli, M.A., Moralles, M., Aguiar, V.A.P., Oliveira, J.R.B., Added, N., Macchione, E.L.A., de Tarso D. Siqueira, P., Zahn, G., Genezini, F., Bonanno, D., Gallo, G., Russo, S., Sgouros, O., Muoio, A., Pandola, L., Cappuzzello, F.: Neutron radiation effects on an electronic system on module. Review of Scientific Instruments **91**, 083301 (2020). https://doi.org/10.1063/5.0010968

[46] Han, J., Guo, G.: Characteristics of energy deposition from 1-1000 mev proton and neutron induced nuclear reactions in silicon. AIP Advances **7**, 115220 (2017). https://doi.org/10.1063/1.4995529

[47] Baumann, R.C.: Radiation-induced soft errors in advanced semiconductor technologies. IEEE Transactions on Device and Materials Reliability **5**, 305–316 (2005). https://doi.org/10.1109/TDMR.2005.853449

[48] Baumann, R.C.: Soft errors in advanced semiconductor devices - part i: the three radiation sources. IEEE Transactions on Device and Materials Reliability **1**, 17–22 (2001). https://doi.org/10.1109/7298.946456

[49] Alía, R.G., Bonaldo, S., Brugger, M., Danzeca, S., Ferrari, A., Frost, C., Infantino, A., Iwamoto, Y., Mekki, J., Theis, C., Thornton, A.:

Single event effect cross section calibration and application to quasi-monoenergetic and spallation facilities. EPJ Nuclear Sciences and Technologies **4**, 1 (2018). https://doi.org/10.1051/epjn/2017031

[50] Tiwari, D., Gupta, S., Rogers, J., Maxwell, D., Rech, P., Vazhkudai, S., Oliveira, D., Londo, D., DeBardeleben, N., Navaux, P., Carro, L., Bland, A.: Understanding gpu errors on large-scale hpc systems and the implications for system design and operation, pp. 331–342 (2015). https://doi.org/10.1109/HPCA.2015.7056044

[51] Oliveira, D., Pilla, L., DeBardeleben, N., Blanchard, S., Quinn, H., Koren, I., Navaux, P., Rech, P.: Experimental and analytical study of xeon phi reliability, pp. 1–12 (2017). https://doi.org/10.1145/3126908.3126960. https://dl.acm.org/doi/10.1145/3126908.3126960

[52] Oliveira, D., dos Santos, F.F., Davila, G.P., Cazzaniga, C., Frost, C., Baumann, R.C., Rech, P.: High-energy versus thermal neutron contribution to processor and memory error rates. IEEE Transactions on Nuclear Science **67**, 1161–1168 (2020). https://doi.org/10.1109/TNS.2020.2970535

[53] Oliveira, D., Blanchard, S., DeBardeleben, N., dos Santos, F.F., Dávila, G.P., Navaux, P., Favalli, A., Schappert, O., Wender, S., Cazzaniga, C., Frost, C., Rech, P.: Thermal neutrons: a possible threat for supercomputer reliability. Journal of Supercomputing **77**, 1612–1634 (2021). https://doi.org/10.1007/s11227-020-03324-9

[54] Blümer, J., Engel, R., Hörandel, J.R.: Cosmic rays from the knee to the highest energies. Progress in Particle and Nuclear Physics **63**(2), 293–338 (2009). https://doi.org/10.1016/j.ppnp.2009.05.002

[55] Kampert, K.-H., Watson, A.A.: Extensive air showers and ultra high-energy cosmic rays: a historical review. The European Physical Journal H **37**(3), 359–412 (2012)

[56] Grieder, P.K.F.: Exentsive Air Showers and High Energy Phenomena, (2010). https://doi.org/10.1007/978-3-540-76941-5

[57] Abreu, P., Aglietta, M., Ahn, E., Albuquerque, I.F.d.M., Allard, D., Allekotte, I., Allen, J., Allison, P., Almeda, A., Castillo, J.A., *et al.*: Measurement of the proton-air cross section at $\sqrt{s} = 57$ Tev with the pierre auger observatory. Physical review letters **109**(6), 062002 (2012)

[58] Greisen, K.: Cosmic ray showers. Annual Review of Nuclear Science **10**(1), 63–108 (1960). https://doi.org/10.1146/annurev.ns.10.120160.000431

[59] Zyla, P.A., *et al.*: Review of Particle Physics. PTEP **2020**(8), 083–01 (2020). https://doi.org/10.1093/ptep/ptaa104. and 2021 update

[60] Matthews, J.: A heitler model of extensive air showers. Astroparticle Physics **22**(5-6), 387–397 (2005)

[61] Roberts, C.D., Bhagwat, M., Höll, A., Wright, S.: Aspects of hadron physics. The European Physical Journal Special Topics **140**(1), 53–116 (2007)

[62] Capdevielle, J.: The influence of baryon resonances and vector mesons on cosmic ray cascades. Journal of Physics G: Nuclear and Particle Physics **18**(2), 43 (1992)

[63] Silberberg, R., Tsao, C.: Spallation processes and nuclear interaction products of cosmic rays. Physics Reports **191**(6), 351–408 (1990)

[64] Goldhagen, P.: Cosmic-ray neutrons on the ground and in the atmosphere. MRS bulletin **28**(2), 131–135 (2003)

[65] Vuković, B., Poje, M., Varga, M., Radolić, V., Miklavčić, I., Faj, D., Stanić, D., Planinić, J.: Measurements of neutron radiation in aircraft. Applied Radiation and Isotopes **68**, 2398–2402 (2010). https://doi.org/10.1016/j.apradiso.2010.06.017

[66] Clem, J.M., Angelis, G.D., Goldhagen, P., Wilson, J.W.: New calculations of the atmospheric cosmic radiation field–results for neutron spectra. Radiation Protection Dosimetry **110**, 423–428 (2004). https://doi.org/10.1093/rpd/nch175

[67] Erlykin, A.: The neutron 'thunder' accompanying large extensive air showers. Nuclear Physics B-Proceedings Supplements **175**, 330–333 (2008). https://doi.org/10.1016/j.nuclphysbps.2007.11.024

[68] Heck, D., Knapp, J., Capdevielle, J.N., Schatz, G., Thouw, T.: CORSIKA: A Monte Carlo Code to Simulate Extensive Air Showers

[69] Engel, R., Heck, D., Huege, T., Pierog, T., Reininghaus, M., Riehn, F., Ulrich, R., Unger, M., Veberič, D.: Towards a next generation of corsika: A framework for the simulation of particle cascades in astroparticle physics. Computing and Software for Big Science **3**, 2 (2019). https://doi.org/10.1007/s41781-018-0013-0

[70] Sidelnik, I., Asorey, H., Collaboration, L., *et al.*: Lago: The latin american giant observatory. Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **876**, 173–175 (2017)

[71] Sarmiento-Cano, C., Suárez-Durán, M., Calderón-Ardila, R., Vásquez-Ramírez, A., Jaimes-Motta, A., Núñez, L.A., Dasso, S., Sidelnik, I., Asorey, H.: The arti framework: cosmic rays atmospheric background simulations. The European Physical Journal C **82**, 1019 (2022). https://doi.org/10.1140/epjc/s10052-022-10883-z

[72] Asorey, H., Sarmiento-Cano, C., Suárez-Durán, M., Rubio-Montero, A.J.: The ARTI Framework. https://doi.org/10.5281/zenodo.7316555. https://github.com/lagoproject/arti

[73] Sidelnik, I.: The sites of the latin american giant observatory, p. 665 (2016). https://doi.org/10.22323/1.236.0665. https://pos.sissa.it/236/665

[74] Asorey, H.: Lago: the latin american giant observatory, p. 247 (2016). https://doi.org/10.22323/1.236.0247. https://pos.sissa.it/236/247

[75] Durán, M.S., Asorey, H., Dasso, S., Nunez, L.A., Pérez, Y.F., Sarmiento, C.: The lago space weather program: Directional geomagnetic effects, background fluence calculations and multi-spectral data anal, p. 142 (2016). https://doi.org/10.22323/1.236.0142. https://pos.sissa.it/236/142

[76] Sarmiento-Cano, C., Asorey, H., Sacahui, J., Otiniano, L., Sidelnik, I.: The latin american giant observatory (lago) capabilities for detecting gamma ray bursts. In: Proceedings of 37th International Cosmic Ray Conference, vol. PoS(ICRC2021), pp. 1–4 (2021). https://doi.org/10.22323/1.395.0929

[77] Asorey, H., Núñez, L.A., Suárez-Durán, M.: Preliminary results from the latin american giant observatory space weather simulation chain. Space Weather **16**, 461–475 (2018). https://doi.org/10.1002/2017SW001774

[78] collaboration, P.A., *et al.*: The pierre auger observatory and its upgrade. Science Reviews-from the end of the world **1**(4), 8–33 (2020)

[79] Galindo, A., Moreno, E., Carrasco, E., Torres, I., Carramiñana, A., Bonilla, M., Salazar, H., Conde, R., Alvarez, W., Alvarez, C., *et al.*: Calibration of a large water-cherenkov detector at the sierra negra site of lago. Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **861**, 28–37 (2017)

[80] Peña-Rodríguez, J., Núñez, L.A., Asorey, H.: Characterization of the muography background using the muon telescope (mute). In: Proceedings of 40th International Conference on High Energy Physics, vol. PoS(ICHEP2020), pp. 1–4 (2021). https://doi.org/10.22323/1.390.0984

[81] Aab, A., Abreu, P., Aglietta, M., Albury, J.M., Allekotte, I., Almela, A., Castillo, J.A., Alvarez-Muñiz, J., Batista, R.A., Anastasi, G.A., *et al.*: Studies on the response of a water-cherenkov detector of the pierre auger observatory to atmospheric muons using an rpc hodoscope. Journal of Instrumentation **15**(09), 09002 (2020)

[82] Sarmiento-Cano, C., Suárez-Durán, M., Vásquez Ramírez, A., Jaimes-Motta, A., Calderón-Ardila, R., Peña-Rodríguez, J.: Modeling the lago's detectors response to secondary particles at ground level from the antarctic to mexico. In: Proceedings of 36th International Cosmic Ray Conference, vol. PoS(ICRC2019), pp. 1–4 (2019). https://doi.org/10.22323/1.358.0412

[83] Rubio-Montero, A.J., Pagán-Muñoz, R., Mayo-García, R., Pardo-Diaz, A., Sidelnik, I., Asorey, H.: The eosc-synergy cloud services implementation for the latin american giant observatory (lago), p. 261 (2021). https://doi.org/10.22323/1.395.0261. https://pos.sissa.it/395/261

[84] Bertolli, C.P., Sarmiento-Cano, C., Asorey, H.: Estimación del flujo de muones en el laboratorio subterráneo andes. In: ANALES AFA, vol. 32, pp. 106–111 (2022)

[85] Peña-Rodríguez, J., Vesga-Ramírez, A., Vásquez-Ramírez, A., Suárez-Durán, M., de León-Barrios, R., Sierra-Porta, D., Calderón-Ardila, R., Pisco-Guavabe, J., Asorey, H., Sanabria-Gómez, J., et al.: Muography in colombia: simulation framework, instrumentation and data analysis. Journal for Advanced Instrumentation in Science **2022**(in press) (2022)

[86] Taboada, A., Sarmiento-Cano, C., Sedoski, A., Asorey, H.: Meiga, a dedicated framework used for muography applications. Journal for Advanced Instrumentation in Science **2022**(1) (2022). https://doi.org/10.31526/jais.2022.266

[87] Vásquez-Ramírez, A., Suárez-Durán, M., Jaimes-Motta, A., Calderón-Ardila, R., Peña-Rodríguez, J., Sánchez-Villafrades, J., Sanabria-Gómez, J., Asorey, H., Núñez, L.: Simulated response of mute, a hybrid muon telescope. Journal of Instrumentation **15**(08), 08004 (2020)

[88] Vesga-Ramírez, A., Sanabria-Gómez, J., Sierra-Porta, D., Arana-Salinas, L., Asorey, H., Kudryavtsev, V., Calderón-Ardila, R., Núñez, L.: Simulated annealing for volcano muography. Journal of South American Earth Sciences **109**, 103248 (2021)

[89] Vásquez-Ramírez, A., Ariza-Gómez, M., Carrillo-Moreno, M., Baldovino-Medrano, V., Asorey, H., Núñez, L.: Improvised explosive devices and cosmic rays. In: Proceedings of 37th International Cosmic Ray Conference, vol. PoS(ICRC2021), pp. 1–4 (2021).

https://doi.org/10.22323/1.395.0480

[90] Sidelnik, I., Asorey, H., Guarín, N., Durán, M.S., Berisso, M.G., Lipovetzky, J., Blostein, J.J.: Simulation of 500 mev neutrons by using nacl doped water cherenkov detector. Advances in Space Research **65**(9), 2216–2222 (2020)

[91] Sidelnik, I., Asorey, H., Guarin, N., Durán, M.S., Lipovetzky, J., Arnaldi, L.H., Pérez, M., Haro, M.S., Berisso, M.G., Bessia, F.A., *et al.*: Enhancing neutron detection capabilities of a water cherenkov detector. Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **955**, 163172 (2020)

[92] Sidelnik, I., Asorey, H., Guarin, N., Durán, M.S., Bessia, F.A., Arnaldi, L.H., Berisso, M.G., Lipovetzky, J., Pérez, M., Haro, M.S., *et al.*: Neutron detection capabilities of water cherenkov detectors. Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **952**, 161962 (2020)

[93] Rubio-Montero, A.J., Pagan-Munoz, R., Mayo-Garcia, R., Pardo-Diaz, A., Sidelnik, I., Asorey, H.: A novel cloud-based framework for standardized simulations in the latin american giant observatory (lago), vol. 2021-December, pp. 1–12 (2021). https://doi.org/10.1109/WSC52266.2021.9715360. https://ieeexplore.ieee.org/document/9715360/

[94] Kneizys, F.X., et al.: The modtran 2/3 report and lowtran 7 model. Technical report, Phillips Laboratory, Hanscom AFB, MA (USA) (1996)

[95] National Aerospace Administration (NASA), N.O., (NOAA), A.A., Force, U.A.: Us standard atmosphere 1976. NOAA technical report NOAA-S/T-76-1562, National Oceanic and Atmospheric Administration (1976)

[96] NOAA Air Resources Laboratory (ARL): Global Data Assimilation System (GDAS1) Archive Information. http://ready.arl.noaa.gov/gdas1.php

[97] Alken, P., Thébault, E., Beggan, C.D., Amit, H., Aubert, J., Baerenzung, J., Bondar, T.N., Brown, W.J., Califf, S., Chambodut, A., Chulliat, A., Cox, G.A., Finlay, C.C., Fournier, A., Gillet, N., Grayver, A., Hammer, M.D., Holschneider, M., Huder, L., Hulot, G., Jager, T., Kloss, C., Korte, M., Kuang, W., Kuvshinov, A., Langlais, B., Léger, J.-M., Lesur, V., Livermore, P.W., Lowes, F.J., Macmillan, S., Magnes, W., Mandea, M., Marsal, S., Matzka, J., Metman, M.C., Minami, T., Morschhauser, A., Mound, J.E., Nair, M., Nakano, S., Olsen, N., Pavón-Carrasco, F.J., Petrov, V.G., Ropp, G., Rother, M., Sabaka, T.J., Sanchez, S.,

Saturnino, D., Schnepf, N.R., Shen, X., Stolle, C., Tangborn, A., Tøffner-Clausen, L., Toh, H., Torta, J.M., Varner, J., Vervelidou, F., Vigneron, P., Wardinski, I., Wicht, J., Woods, A., Yang, Y., Zeren, Z., Zhou, B.: International geomagnetic reference field: the thirteenth generation. Earth, Planets and Space **73**, 49 (2021). https://doi.org/10.1186/s40623-020-01288-x

[98] Hoerandel, J.R.: On the knee in the energy spectrum of cosmic rays. Astroparticle Physics **19**(2), 193–220 (2003)

[99] Ostapchenko, S.: Monte carlo treatment of hadronic interactions in enhanced pomeron scheme: Qgsjet-ii model. Physical Review D **83**, 014018 (2011). https://doi.org/10.1103/PhysRevD.83.014018

[100] Groom, D.E., Mokhov, N.V., Striganov, S.I.: Muon stopping power and range tables 10 mev–100 tev. Atomic Data and Nuclear Data Tables **78**(2), 183–356 (2001)

[101] Bonechi, L., D'Alessandro, R., Giammanco, A.: Atmospheric muons as an imaging tool. Reviews in Physics **5**, 100038 (2020)

[102] y C. Sarmiento-Cano y H. Asorey, C.P.B.: Estimation of the muon flux expected at the andes underground laboratory. Anales de la Asociación de Física Argentina (AFA) **32**(4), 106–111 (2022)

[103] Liao, W., Hashimoto, M., Manabe, S., Watanabe, Y., Abe, S.-I., Nakano, K., Takeshita, H., Tampo, M., Takeshita, S., Miyake, Y.: Negative and positive muon-induced seu cross sections in 28-nm and 65-nm planar bulk cmos srams. In: 2019 IEEE International Reliability Physics Symposium (IRPS), pp. 1–5 (2019). IEEE

[104] Hashimoto, M., Liao, W.: Soft error and its countermeasures in terrestrial environment, vol. 2020-January, pp. 617–622 (2020). https://doi.org/10.1109/ASP-DAC47756.2020.9045161

[105] Dasso, S., Asorey, H., for the Pierre Auger Collaboration: The scaler mode in the pierre auger observatory to study heliospheric modulation of cosmic rays. Advances in space research **49**(11), 1563–1569 (2012)

[106] Ayre, C.A., Gustafson, H.R., Jones, L.W., Longo, M.J., Ramana Murthy, P.V.: Neutron-nucleus total cross section, 30-270 GeV. In: International Cosmic Ray Conference. International Cosmic Ray Conference, vol. 7, p. 2268 (1975)

[107] Abraham, J., Abreu, P., Aglietta, M., Aguirre, C., Ahn, E., Allard, D., Allekotte, I., Allen, J., Allison, P., Alvarez-Muniz, J., *et al.*: Atmospheric effects on extensive air showers observed with the surface detector of the

pierre auger observatory. Astroparticle Physics **32**(2), 89–99 (2009)

[108] Masías-Meza, J.J., Dasso, S., Démoulin, P., Rodriguez, L., Janvier, M.: Superposed epoch study of icme sub-structures near earth and their effects on galactic cosmic rays. Astronomy & Astrophysics **592**, 118 (2016)

[109] Mishev, A., Usoskin, I.: Current status and possible extension of the global neutron monitor network. Journal of Space Weather and Space Climate **10**, 17 (2020)

[110] Violante, M., Sterpone, L., Manuzzato, A., Gerardin, S., Rech, P., Bagatin, M., Paccagnella, A., Andreani, C., Gorini, G., Pietropaolo, A., *et al.*: A new hardware/software platform and a new 1/e neutron source for soft error studies: Testing fpgas at the isis facility. IEEE Transactions on Nuclear Science **54**(4), 1184–1189 (2007)

[111] Böhlen, T., Cerutti, F., Chin, M., Fassò, A., Ferrari, A., Ortega, P.G., Mairani, A., Sala, P.R., Smirnov, G., Vlachoudis, V.: The fluka code: developments and challenges for high energy and medical applications. Nuclear data sheets **120**, 211–214 (2014)

[112] Agostinelli, S., Allison, J., Amako, K., Apostolakis, J., Araujo, H., Arce, P., Asai, M., Axen, D., Banerjee, S., Barrand, G., Behner, F., Bellagamba, L., Boudreau, J., Broglia, L., Brunengo, A., Burkhardt, H., Chauvie, S., Chuma, J., Chytracek, R., Cooperman, G., Cosmo, G., Degtyarenko, P., Dell'Acqua, A., Depaola, G., Dietrich, D., Enami, R., Feliciello, A., Ferguson, C., Fesefeldt, H., Folger, G., Foppiano, F., Forti, A., Garelli, S., Giani, S., Giannitrapani, R., Gibin, D., Cadenas, J.J.G., González, I., Abril, G.G., Greeniaus, G., Greiner, W., Grichine, V., Grossheim, A., Guatelli, S., Gumplinger, P., Hamatsu, R., Hashimoto, K., Hasui, H., Heikkinen, A., Howard, A., Ivanchenko, V., Johnson, A., Jones, F.W., Kallenbach, J., Kanaya, N., Kawabata, M., Kawabata, Y., Kawaguti, M., Kelner, S., Kent, P., Kimura, A., Kodama, T., Kokoulin, R., Kossov, M., Kurashige, H., Lamanna, E., Lampén, T., Lara, V., Lefebure, V., Lei, F., Liendl, M., Lockman, W., Longo, F., Magni, S., Maire, M., Medernach, E., Minamimoto, K., de Freitas, P.M., Morita, Y., Murakami, K., Nagamatu, M., Nartallo, R., Nieminen, P., Nishimura, T., Ohtsubo, K., Okamura, M., O'Neale, S., Oohata, Y., Paech, K., Perl, J., Pfeiffer, A., Pia, M.G., Ranjard, F., Rybin, A., Sadilov, S., Salvo, E.D., Santin, G., Sasaki, T., Savvas, N., Sawada, Y., Scherer, S., Sei, S., Sirotenko, V., Smith, D., Starkov, N., Stoecker, H., Sulkimo, J., Takahata, M., Tanaka, S., Tcherniaev, E., Tehrani, E.S., Tropeano, M., Truscott, P., Uno, H., Urban, L., Urban, P., Verderi, M., Walkden, A., Wander, W., Weber, H., Wellisch, J.P., Wenaus, T., Williams, D.C., Wright, D., Yamada, T., Yoshida, H., Zschiesche,

D.: Geant4—a simulation toolkit. Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **506**, 250–303 (2003). https://doi.org/10.1016/S0168-9002(03)01368-8