ALICE activity report 2023

N. Bianchi, A. Fantoni (Resp.), P. Larionov, V. Muccifora, S. Pisano (Ass.),
F. Ronchetti, E. Spiriti, M. Toppi, O. Vazquez Doce

1 The ALICE experiment

The ALICE collaboration at CERN currently includes 40 countries, 169 institutions, and 1968 members. In turn, INFN participates with 12 groups for a total of about 200 physicists. The INFN-Frascati group is a very active contributor to the scientific output of the collaboration in terms of detector construction, run operation and physics analysis. In fact, the INFN-Frascati group played a key role in the construction and operation of the ALICE electromagnetic calorimeters EMCAL and DCAL, in the upgrade for RUN3 (2021-24) with the construction of 1/4 of the new Inner Tracker System (ITS) Outer Layers (OL), as well as in the operation of the entire ALICE detector (Run and Commissioning Coordination on 2013-15 and 2019-22 and Training Shift Coordinator on 2023-24).

This report briefly summarizes the results obtained by the ALICE-LNF group during the restart of real data taking in 2023 after the LS2 and after the commissioning of the new ALICE Monolithic Active Pixel Sensors (MAPS) ITS which has replaced the old Run 1-2 device (based on hybrid pixel sensors, silicon strips, silicon drifts).

The INFN-Frascati group has been very active also for physics analysis and in particular for the extraction of the π, K, and p spectra from the newest high-energy data set of p-Pb collisions at 8.16 TeV, an essential reference analysis for any light flavor physics as well as for the light antinuclei that can be used to study cosmic-ray interactions and dark-matter annihilation and for the antikaon-deuteron femtoscopic correlations with ALICE.

Since fall 2019 the INFN-Frascati group covered different roles of responsibility, having a leading role to the ALICE Run Coordination (f. Ronchetti), being in charge of the global ALICE commissioning and data taking operations until the end of 2022, being an elected member of the Management Board of the ALICE experiment (A. Fantoni first and F. Ronchetti later), participating by defaults also to the Physics and Technical Boards, being the Training shift Coordinator (S. Pisano) and also the EPN Technical Coordinator (F. ronchetti) for the Run 3.

2 ALICE upgrade for RUN3

One of the major ALICE upgrades in preparation for Run 3 was the design and deployment of a completely new computing model: the O^2 project, which merges online (synchronous) and offline (asynchronous) data processing into a single software framework. In practice, the same code runs online and offline using different selections and parameters. The O^2 computing model also required an upgrade of the experiment's computing farms for the data readout and processing, discussed further in this article.

Due to the increased data volumes, storing all the produced raw data is infeasible, hence the need for efficient online compression and the use of GPUs (Graphic Processing Units) instead of CPUs to speed up processing. GPUs are highly efficient processors, which have a much higher compute throughput than CPUs thanks to the intrinsic parallelism and reduce the cost of the farm and the energy consumption. Without GPUs, about eight times as many servers of the same type and other resources would be required to handle the TPC online processing of Pb–Pb collision data at a 50 kHz interaction rate (equivalent to an instantaneous LHC luminosity of $6x10^27\ cm^{-2}s{-}1$). This data stream is pushed by the ALICE readout farm (FLP, First Level

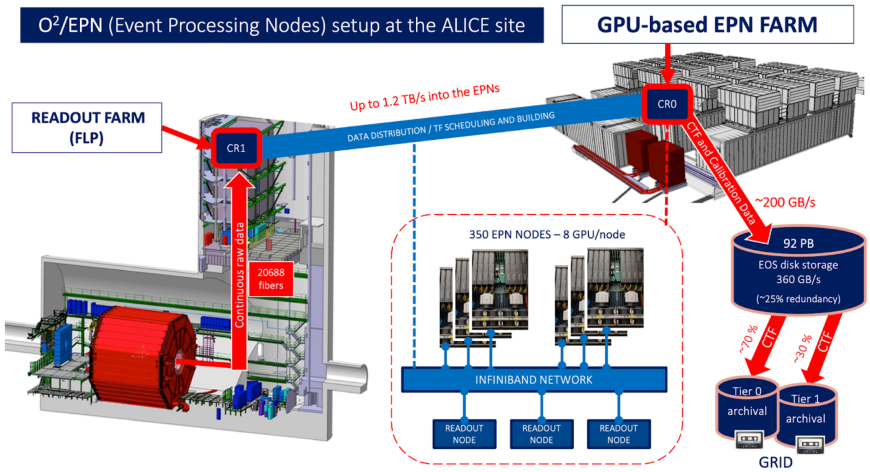

Figure 1: Overview of the ALICE detector dataflow. All detector front-end cards are read out via 20,688 GBT fibers which input to the farm in the CR1 container. The readout nodes are connected via an InfiniBand network to the EPN farm hosted in the surface container (CR0). The data processing occurs on the EPN farm GPUs and the output is then transferred to the CERN distributed storage system, EOS.

Processors) via an InfiniBand (IB) network to the Event Processing Nodes (EPN) using a data distribution software running on both farms (see 1). The EPN farm is hosted in two containers, called CR0, located on the surface close to the ALICE site and is currently composed of 280 servers of the reference configuration, each equipped with eight AMD MI50 GPUs with 32 GB of RAM each, two 32-core AMD Rome CPUs, and 512 GB of memory. The EPN farm is optimized for the fastest possible TPC track reconstruction, which constitutes the bulk of the synchronous processing, and provides most of its computing power in the form of GPU processing following the HLT paradigm. In fact, the initial version of the O2 tracking is an improved version of the old HLT tracking, particularly concerning the track resolution (which now matches the offline resolution in Run 2) and the treatment of data from continuous readout with overlapping collisions in the TPC. The data flow from the front end into the farms and cannot be buffered, so the EPN computing capacity must be sufficient for the highest data rates expected during Run 3.

3 Dataflow and processing

Due to the continuous readout approach, processing does not occur on an "event" triggered by some characteristic pattern in detector signals. Instead, all data is read out and stored during a predefined time slot into a time frame (TF) data structure. The TF length is a customizable parameter and is usually chosen as a multiple of one LHC orbit. Since a full TF is processed at once, it must fit in the GPU's memory. Thus a development effort was put in place for reusing GPU memory in consecutive processing steps. During the proton running in 2022 the TF length was chosen to be 128 LHC orbits, even though it had been verified that the performance for the full reconstruction behaves identically up to 256 LHC orbits. The system was stressed by increasing the proton collision rates beyond those needed to integrate the pp luminosity for physics analysis purposes. Such high-rate tests aimed to reproduce occupancies similar to the expected rates of lead collisions. It was demonstrated that the EPN processing could sustain rates nearly twice the nominal design value of 600 GB/s originally foreseen for Pb–Pb collisions. Using high-rate proton collisions at 2.6 MHz the readout reached 1.24 TB/s which was fully absorbed and processed on the EPNs.

The synchronous processing on the EPN farm has two main objectives: the compression of the raw data into compressed time frames (CTF), which end up on the disk buffers, and second, the extraction of all necessary data for the detector's online calibration, such that successive calibration tasks may then run without the need to access the raw data again.

At the incoming raw data rates it is practically impossible to store the data, even temporarily. Hence, the outgoing data is compressed in real time to a manageable size on the EPN farm. The event building happens during this network transfer. The data distribution suite collects partial TFs from the readout farm and schedules the building of the complete TF such that each EPN node receives and processes a full TF (see 1). Thus, each EPN sees the data from all detectors, but only for the duration of a single time frame each. On the other hand, calibration tasks can run on both the readout nodes and the EPNs. All calibration tasks requiring global information and operating on full TFs, generally run on the EPN farm.

4 The first high-luminosity Pb-Pb run

The ALICE commissioning during 2022 revealed that the TPC was producing roughly 30% more clusters then anticipated. Already at the end of 2022 30 additional EPN worker nodes where purchased to try to keep the required "safety compute margin" of 20% with respect to the nominal rates expected at 50 kHz. The need to process the extra data size coming from the TPC has required a further enhance of the farm computing capacity in preparation for the first high-rate heavy-ion run scheduled for October 2023. Hence, 70 new computing nodes equipped with more advanced GPUs, where installed, resulting in a 30% increase in overall computing power. The upgrade has requested overseeing node installation, configuration, network load optimization towards the CERN IT EOS system using IB to Ethernet gateways, and conducting stability and performance tests. Additional network connectivity from the FLP farm to the EPN far was put in place and one of spare IT containers was used to host part of the new 70 EPN nodes. The EPN farm is strucured in building blocks of 3 adjacent rack where all the worker node are connected to a TOR (Top Of the Rack) IB switch. The extra nodes were installed in the spare slots of the already existing building block and a new block was created in one of the spare IT containers. The whole operation required the opening of the trenches located at the base of the CR0 infrastructure to install the new trunk fibres and the placing patch panels to connect the new building block to the core of the IB network. The installation of the new nodes which required also the reshuffling of the physical positions of 70 of the "old" servers in order to obtain 10 balanced building blocks each with 28 "old" servers and 7 "new servers". Such reshuffling was critical to achieve the correct network transfer balancing from the internal EPN core network (InfiniBand) to the CERN IT EOS (Ethernet) via the 4 IB-to-Ethernet gateways at the ALICE P2. The new servers were configured via Ansible and their inclusion in the production environment was done after running FSTs (Full System Test, a stand alone reconstruction task) in order to validate the GPU processing and the overall hardware stability.

The resulting farm was featuring 350 nodes which had different GPU cards: 280 nodes with previous generation GPUs (MI50) and 70 nodes with more powerful models (MI100). As a consequence and optimization of the synchronous reconstruction workflows to leverage the upgraded farm was carried out. The two main action items were the identification of the suitable software stack that performed correctly on both kind of GPUs (AMD base installation software, compilers and kernel modules) since the vendor standard stack don't work out of the box with the ALICE software.

The ability of the core framework to handle etherogeneous topologies was absolutely critical for the Pb-Pb run in october in order to use the new nodes. Extensive work was done also to automatise the deployment of all the observability tools (EPN InfoLogger and the Telgraf-Influxdb-Grafana stack) which are used to monitor the status of the EPN farm and propagate this information to the other online systems and in the ALICE control room. This allowed a rapid re-deployment of these services during the Pb-Pb run in case of failures or when the infrastructure nodes were overloaded. The day-by-day operations for the Pb-Pb run was closely followed up either being on call or by supporting the rest of the EPN team and the Run Coordination team. A continous effort was done in configuring the EPN farm to run new patched versions of the processing software quickly between fills or during any slot with non physics beam. This bursts of on-the-fly software updates were critical and strongly needed since the LHC accelerator had operational issues and there where several weeks of stop of the pp beam operations right before staring the October Pb-Pb run. As a consequence, most of the validations for the processing using hi-rate p-p data could not be done and our software had to be tuned directly during the Pb-Pb data taking. However the overall behaviour or the EPN far was very stable and peaks of 47 kHz Pb-Pb hadronic rates with 770 GB/s in input to the farm were absorbed without problems and data to the CERN IT EOS could be transferred at 200 GB/s sustained, hence demonstrating the the EPN system can cope with the expected data rates for Run 3.

5 Contribution to the ITS QA

It is worth to remind that the Frascati group provided 1/4 of the total Outer Barrel staves, building and assembling 29 staves between the end of 2018 and end of 2019. After the installation, the Frascati group has contributed to the commissioning of the detector and during the data taking in 2023 (Run3) has contributed in the Offline Data Quality Assurance (QA) of the whole detector, being responsible for the Cluster (cluster size and occupancy) and for the Tracks (vertex parameters and track angular distributions).

In 2023 14.5 pb^{-1} of luminosity has been delivered to ALICE for pp data taking at 13.6 TeV with production for physics at 500 kHz interaction rate. A fraction of fills has been dedicated to Heavy Ions (HI) preparation for the TPC firmware validation and high rate scan at 4 MHz. Pb-Pb data taking at 5.36 TeV has been possible with most of the fills leveled at $\sim$25 kHz, but some data also collected at $\sim$45 kHz. The QA team has analyzed the new whole periods 3 times per week, coordinated via the JIRA ticket system.

In order to define a bad run from the cluster point of view, the following criteria have been applied:

- at least 1 layer with >25% empty staves (cluster occupancy is 0 cluster/pixel/ nChip);
- the run has >10% empty lanes overall;
- the average cluster size is out of limits by 3-7 pixels

The detector occupancy has been studied also with the cluster task and it has been found that cluster size is independent of the Interaction Rate (IR) and that the decrease of the cluster size by the end of the fill can be due to the beam-gas interactions.

From the Tracks point of view a run is defined good if the following quality criteria for offline QA have been satisfied:

- no anomalies in angular track distribution;
- the Z vertex shape ranging between -1.5 and 1.5 cm;
- the average nClusters per track ranging between 5 to 6.

ITS demonstrated stable performance to both periods, so the Quality Check (QC) is an effective tool for the run quality assessment and filtering bad data, showing no significant differences in average efficiency wrt p-p data taking.

6 Physics contribution

6.1 Analysis on light-flavor

During 2023, the study of the formation mechanism of light nuclei and anti-nuclei, in particular deuterium and ^{3}He, to which the LNF group contributed with the extraction of the proton spectra, has been finally published in *Physics Letters B* (https://doi.org/10.1016/j.physletb.2023.137795). The nain resul is reported in Fig. 2

Two main models are presently used to describe the mechanism that nucleons undergo for forming a two or three-nucleon nucleus: the Statistical Hadronization Model and the Coalescence model. These two approaches are sensitive to different characteristics of the formation process: while the former, indeed, is mainly sensitive to the mass and the spin degeneracy of the formed nucleus, the second is also sensitive to the size of the latter, since the coalescence process implies the proximity of two nucleons (that should also be close in the phase space) and is then sensitive to the source size. A possible approach for disentangling between the two models consist in testing them for nuclei of different sizes, and, in this perspective, ALICE produced the most accurate sets of data extracting the relevant observables for different light nuclei (as deuterium and helium) in different colliding systems and at different energies. In this perspective, one of the most interesting observable is the ratio of the deuteron and helium yields to the proton ones, and the evolution of this ratio as a function of the final charged particle multiplicity. The measurement provided by the ALICE Collaboration is then expected to play a crucial role for testing the phenomenological power of the different models.

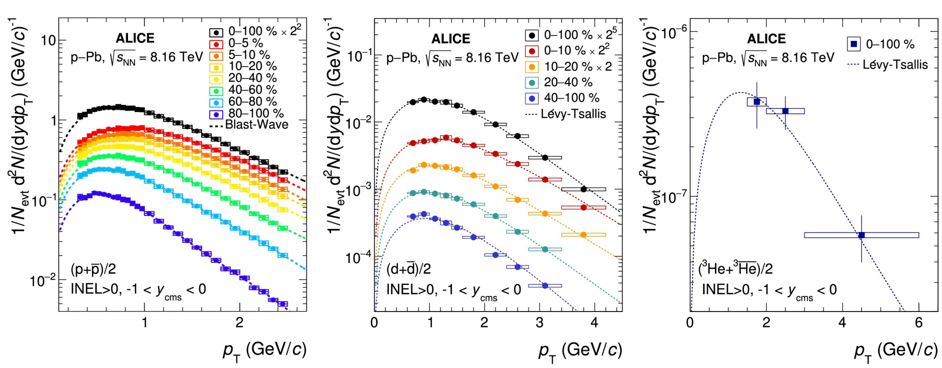

Figure 2: Spectra for proton (left), deuteron (middle) and helium (right) as a function of the transverse momentum p_T

6.2 Analysis on light antinuclei

The observation of antinuclei such as $^{3}\bar{\mathrm{He}}$ is one of the most promising signatures of Dark Matter (DM) annihilation of weakly interacting massive particles. The kinetic-energy distribution of antinuclei produced in DM annihilation peaks at low kinetic energies (E_{kin} per nucleon $\leq$0.1 GeV/A) for most assumptions of DM mass. In contrast, for antinuclei originating from cosmic-ray interactions the spectrum peaks at much larger E_{kin} per nucleon ($\simeq$10 GeV/A). Thus, the low-energy region is almost free of background for DM searches. Since no $^{3}\bar{\mathrm{He}}$ beams are available, the antimatter production at the LHC and the excellent identification and momentum determination for ^{3}He in ALICE has been used as an equivalent setup. In this study, the ALICE detector itself has served as a target for the inelastic processes.

The disappearance probability of $^3\overline{\mathrm{He}}$ when it encounters matter particles and annihilates or disintegrates within the ALICE detector at the Large Hadron Collider has been determined. The inelastic interaction cross section has been extracted (Fig 3) and then used as input to calculations of the transparency of our Galaxy to the propagation of $^3\overline{\mathrm{He}}$ stemming from dark-matter annihilation and cosmic-ray interactions within the interstellar medium.

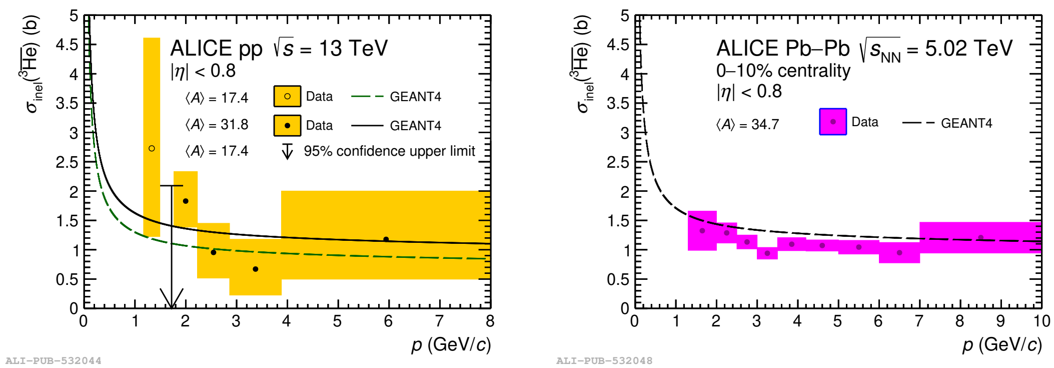

Figure 3: *Results for $\sigma_{inel}(^3\overline{\mathrm{He}})$ as a function of $^3\overline{\mathrm{He}}$ momentum. Results obtained from pp collisions at $\sqrt{s}$=13 TeV (up); results from the 10% most central Pb–Pb collisions at $\sqrt{s_{NN}}$=5.02 TeV (down). The curves represent the GEANT4 cross sections corresponding to the effective material probed by the different analyses. The arrow on the upper plot shows the 95% confidence limit on $\sigma_{inel}(^3\overline{\mathrm{He}})$ for $< A >$=17.4. The different values of average mass number $< A >$ correspond to the three different effective targets (see the paper for details). All the indicated uncertainties represent standard deviations.*

For a specific DM profile, a transparency of about 50% has been estimated, varying with increasing $^3\overline{\mathrm{He}}$ momentum from 25% to 90% for cosmic-ray sources. The results indicate that $^3\overline{\mathrm{He}}$ nuclei can travel long distances in the Galaxy, and can be used to study cosmic-ray interactions and dark-matter annihilation. The Frascati group has been involved in this analysis and one of the member (P. Larionov) has been one of the main analyser and author of the paper, published on Nature Physics 19 (2023) 1, 61-71.

The measured $\sigma_{inel}(^3\overline{\mathrm{He}})$ and the developed methodology can be employed to carry out the propagation of $^3\overline{\mathrm{He}}$ using any DM or cosmic-ray interaction modelling as a source. Since a large separation between signal and background is retained for low kinetic energies,our results clearly underline that the search for $^3\overline{\mathrm{He}}$ in space remains a very promising channel for the discovery of DM. These studies will be extended to $^4\overline{\mathrm{He}}$ and to the lower momentum region in the near future with much larger data sets that will be collected in the coming few years.

7 Three body dynamics via femtoscopy

The INFN Frascati group has been involved in two recent papers published by ALICE regarding the access to the three-body dynamics via femtoscopy studies in pp collisions. Two different approaches are used, on the one hand it is demonstrated that, given the small distances at which particles are emitted in pp collisions, hadron-deuteron correlations are sensitive to the three-body dynamics of the involved particles, considering the internal structure of the deuteron. On the other hand, a more direct approach is used in the analysis of the three-body $K^{+/-}$–p–p correlation function.

A member of the INFN Frascati group (O. Vazquez Doce) takes part of the Paper Committee of both articles, and has been the main analyzer of the anti-Kaon–deuteron correlations (see Analysis Note in https://alice-notes.web.cern.ch/node/1259https://alice-notes.web.cern.ch/node/1259).

7.1 Hadron-deuteron correlations

By studying the hadron-deuteron correlation function in pp collisions, a new experimental method to study three-body nuclear systems has been established. The study of three- and many-body dynamics has been a long-standing goal in nuclear physics, particularly for understanding the structure of light nuclei and describing neutron-rich and dense nuclear matter. ALICE has published the measurement of correlations in the momentum space of deuteron-hadron pairs. The K^+–d and p–d correlation function are measured in high-multiplicity proton-proton collisions at $\sqrt{s} = 13$ TeV and are shown on Fig. 4

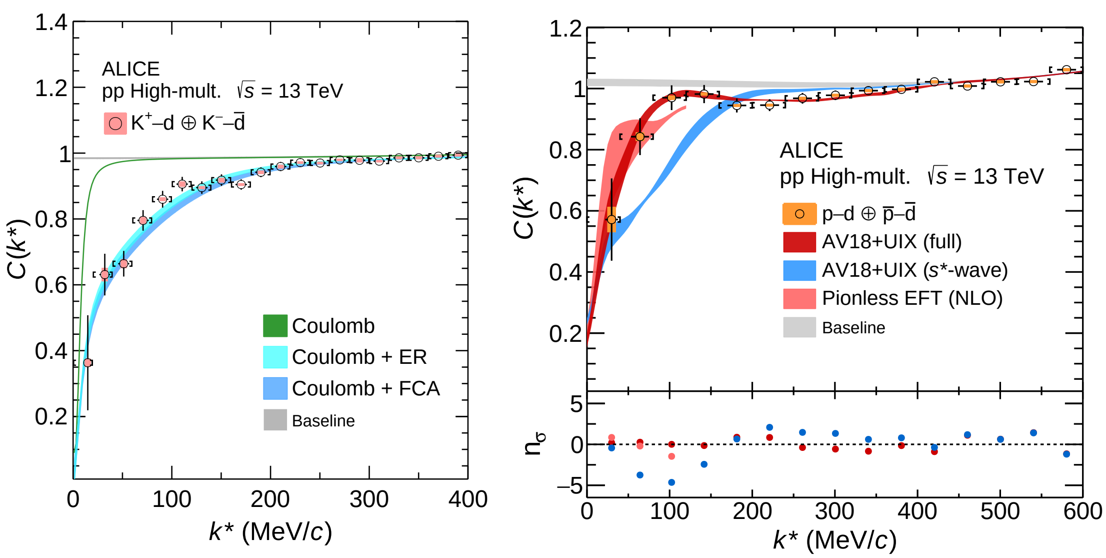

Figure 4: Measured K^+–d (left panel) and p–d (right panel) correlation functions as a function of the relative momentum k*. The blue curves in the left panel represent the theoretical expectation under the two-body approach. The red curve in the right panel represents the theoretical expectation using a full-fledged three-body calculation. See text for details.

The correlation functions are compared with effective two-body calculations anchored to: i) scattering parameters describing the K^+–d and p–d interactions obtained via scattering experiments; ii) to the characterization of the size of the particle source, of around 1 fm, determined using hadron-hadron correlations in pp collisions. In the case of the K^+–d correlation, an excellent description of the measurement is obtained, as can be seen by the blue bands in the left panel of Fig.4, demonstrating that the two assumptions made (small distances and two-body interactions) hold for this system.

However such calculations fail to describe the p–d system. In this case, it is clear that one cannot assume just two-body interactions between the proton and the deuteron, and the internal structure of the deuteron should be considered. The discrepancy between the theoretical expectation and the ALICE data can only be resolved by performing a full three-body calculation that accounts for the underlying three-nucleon dynamics. Detailed calculations based on the deuteron wave function, the formation probability of the deuteron under the coalesence model, and the underlying two-body (via Argonne V18 potential) and three-body interactions (via Urbana IX potential), are used. The right panel in Fig. 4 shows the level of agreement between the full-fledged three-body calculations (red curve) and the ALICE measurement.

The analysis demonstrates that nucleons are the explicit degrees of freedom also in the correlations among light nuclei produced at short distances in hadronic collisions and opens the possibility of investigating the effect of genuine many-body nuclear interactions at the LHC in the future, including the strangeness and heavy flavour sectors.

7.2 Three-body correlation function

While the Kaon–nucleon and in particular the anti-Kaon–nucleon interactions have been extensively studied, details of the three-body KNN and $\overline{\mathrm{K}}$NN dynamics are still not well understood, mainly due to the overlap with multi-nucleon interactions in nuclei. An alternative method to probe the dynamics of three-body systems with kaons is to study the final state interaction within triplet of particles emitted in pp collisions, which are free from effects due to the presence of bound nucleons. ALICE has reported in Eur. Phys. Journal A **59** (2023) no.12, 298 the first femtoscopic study of p–p–K^+ and p–p–K^- correlations measured in high-multiplicity pp collisions at $\sqrt{s} = 13$ TeV.

Three-body correlation functions are presented for the triplets as a function of the Q_3, the hyper-momentum that represents the relative momentum of the particles in the triplet. Additionally, by utilising the expansion formalism of the three-particle cumulants, a three-body cumulant is calculated, where the lower order correlations, deriving for the low momentum region from the two-body interactions, are removed from the original correlation function. In order to extract such lower order correlations, a simultaneous measurement of the correlation function of all possible particle pairs within the triplet of interest is performed, and used to constrain a projector method that allows to obtain an expected distribution as a function of the relative momentum of the three particles as expectation for the effect of lower order correlations. The resulting cumulants, where such exepectation from the two-body interactions only are removed, are shown Fig. 5. The cumulants do not present deviations from the null value for the whole range of Q_3, indicating than effects that go beyond the two-body interaction are absent.

As a conclusion, the analysis shows that the measured correlation functions can be interpreted in terms of pairwise interactions in the triplets, indicating that the dynamics of such systems is dominated by the two-body interactions without significant contributions from three-body effects or bound states. Such conclusions constitute important inputs for the calculation of the equation of state of neutron stars and the formation of kaonic nuclei.

8 ALICE scientific output

The ALICE Collaboration has published 65 papers in 2023 and to date 473 papers submitted to international referred physics journals, of which 430 already published (Fig. **??**) since the birth of the ALICE experiment in 2009.

The full list of ALICE publications for the year 2023 can be found online at the link: https://alice-publications.web.cern.ch/statistics/2023 while the full list is available at the link: https://alice-publications.web.cern.ch/publications

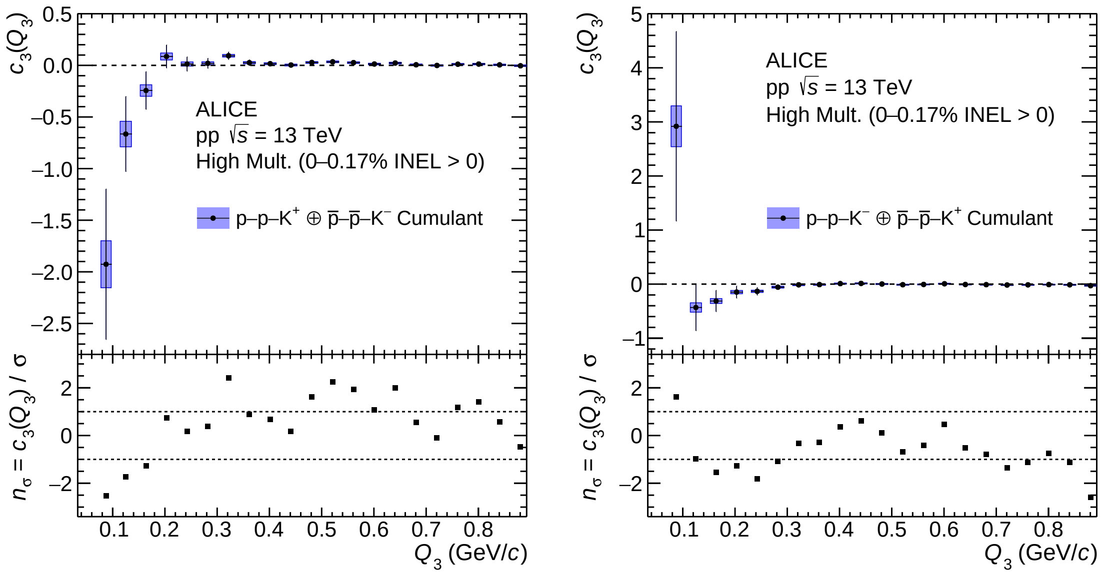

Figure 5: Cumulants for the p–p–K$^+$ (left panel) and p–p–K$^-$ (right panel) primary triplets. The n_σ deviations from zero in each bin are shown in the bottom panels.

9 ALICE LNF talks

S. Pisano "The ALICE overview", 153rd LHCC meeting, CERN March 2023
O. Vazquez Doce "Precise tests of the hadron-hadron strong interaction via femtoscopy", INFN CSN3 meeting Feb. 2023
O. Vazquez Doce "Precise tests of the hadron-hadron strong interaction via femtoscopy", IFIC Valencia Physics Seminar, February 2023 O. Vazquez Doce "Precise tests of the hadron-hadron strong interaction via femtoscopy (in small systems!)", Theory Seminar at Instituto de Física de Partículas y del Cosmos (Universidad Complutense Madrid) June 2023
O. Vazquez Doce "Updated experimental insight into the KN interaction" EMMI Workshop Bound states and particle interactions in the 21st century, July 2023, Trieste (Italy)
O. Vazquez Doce "Accessing the antiKN interaction via femtoscopy and antiKaonic atoms" Mini-Workshop "Kaonic atoms: present status and future plans", Laboratori Nazionali di Frascati (Italy) July 2023
O. Vazquez Doce "Three-body dynamics at short range via deuteron-hadron correlations by ALICE" Poster presentation, Quark Matter Conference Houston (USA), September 2023
O. Vazquez Doce "Accessing hadron-hadron strong interaction and three-body dynamics via femtoscopy" TNPI2023 - XIX Conference on Theoretical Nuclear Physics in Italy TNPI2023, Cortona (Italy) October 2023
O. Vazquez Doce "Study of the three-body dynamics at short range via femtoscopy by ALICE at the LHC", WPCF 2023 - XVI Workshop on Particle Correlations and Femtoscopy IV Resonance Workshop (WPCF 2023), Catania (Italy) November 2023.

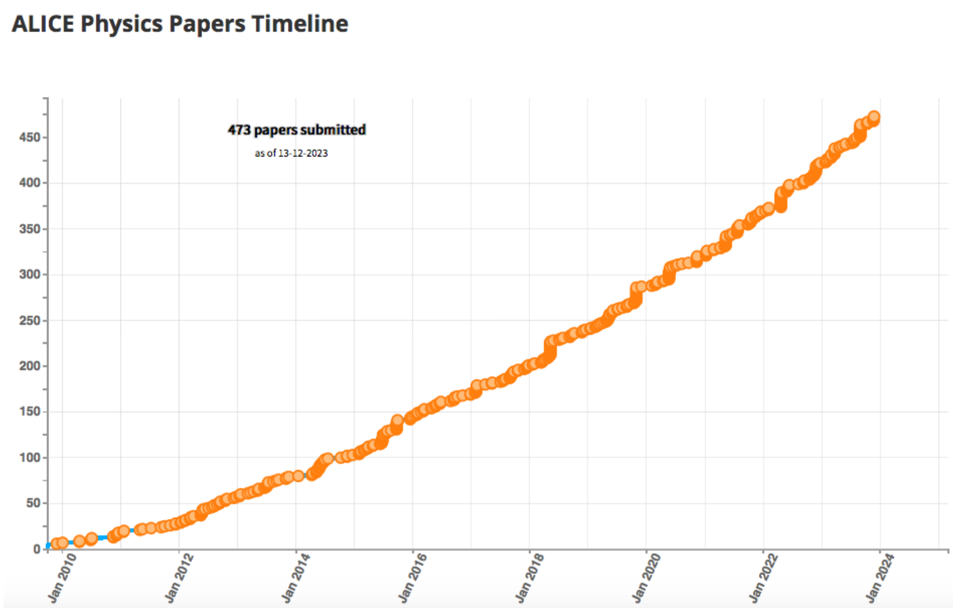

Figure 6: *Timeline of the total number of ALICE papers ("submitted" is to be intended as published+submitted) since the fist LHC beam at 900 GeV on November 23, 2009).*