Nuclear Physics Tools – Machine Learning, Artificial Intelligence, and Quantum Computing

Convenors: Valerio Bertone (CEA Paris-Saclay), Jana N. Günther (U. Wuppertal)
NuPECC Liaisons: Hervé Moutarde (CEA Paris-Saclay), Eugenio Nappi (INFN Bari)
WG members:

- Joerg Aichelin,
- Mohammad Al-Turany,
- Sara Collins,
- Tommaso Dorigo,
- Christian S. Fischer,
- Michele Grossi,
- Philipp Hauke,
- Andreas Ipp,
- Marc Labiche,
- Denis Lacroix,
- Johan Messchendorp,
- Stefano Piano,
- Arnau Rios,
- Alessandro Roggero,
- Olga Soloveva,
- Vittorio Somà,
- Liangxiao Wang.

0. Introduction

The tremendous progress in the field of Nuclear Physics has led to the pressing need for appropriate numerical tools aimed at addressing the most relevant experimental, theoretical, and technological challenges, such as those encompassed by the Joint ECFA-NuPECC-APPEC (JENA) initiative. To this end, the advent of algorithms based on Machine-Learning (ML) and Artificial-Intelligence (AI) techniques, and the fast progress in the field of Quantum Computing (QC) have opened an entire new world of possibilities. Nuclear and particle physicists from all around the world have turned their attention to these technologies in quest of more efficient tools to interpret the abundance of experimental data that is currently being delivered at nuclear and high-energy facilities.

However, so far, the exploration of possible applications of ML, AI, and QC to nuclear and high-energy physics in Europe has mostly proceeded incoherently with local or at most national initiatives devoted to it. The purpose of this chapter is to provide a broad and as comprehensive as possible overview of the current status of how these techniques are being employed in nuclear physics, to coordinate this effort at a European level.

The wideness and transversality of this subject make this task complicated to be accomplished. To identify the main guidelines for the final recommendations, this chapter is split into four main sections respectively devoted to:

- machine learning and artificial intelligence,
- quantum computing,
- tools and techniques,
- resources and infrastructure.

This organisation will allow us to focus on the specific aspects of these main branches to eventually issue precise recommendations that have the purpose of optimising European resources.

1. Machine Learning and Artificial Intelligence

Nuclear physics stands at the forefront of our quest to understand the intricacies of subatomic interactions, playing a pivotal role in deciphering the properties of strongly interacting matter. Parallel to this, the meteoric rise of AI has revolutionised our approach to complex problems. ML, as the most prominent subset of AI, is specifically designed to discern patterns in intricate data and encapsulate them with an optimal set of parameters. ML is suitable for rich data and sophisticated explorations, which promises unprecedented opportunities and insights for understanding nuclear matter in today's AI-driven world.

In the realm of nuclear physics, ML applications have been explored in areas such as nuclear experiments, nuclear astrophysics, and various computing-intensive tasks, as shown in Figure 1. Within nuclear physics experiments, ML algorithms have been utilised to process large datasets, aiding in particle identification, improving event reconstructions, and allowing for experiment design and control. In the field of nuclear astrophysics, ML has been applied to analyse signals, which is particularly useful in processing data from noisy space environments. It has also assisted in determining the properties of dense matter, which is crucial for understanding certain celestial events. ML has also been beneficial for computing-intensive challenges. It has been applied in hadronic structure and nuclear collisions [see TWGs 1 and 3], astrophysical simulations [see TWG 4], and notably in Lattice QCD [see TWG 1], a first-principles method, to enhance our understanding of nuclear matter.

Thanks to their flexibility as universal approximators, neural networks (NNs) have been proposed as an effective framework to design accurate ansatzes for the wave function of strongly correlated many-body systems. Currently proposed architectures have already shown impressive results for nuclear structure applications on closed shell nuclei where they can match or even surpass the accuracy achieved by state-of-the-art Quantum Monte Carlo simulations. Similarly, the application of NNs to parameterise the hadronic structure of nucleons and nuclei relevant to the physics of high-energy colliders, such as the Large Hadron Collider (LHC), has already been extensively used and proven superior in many respects to more standard approaches. ML techniques on data simulated by various transport models empower us to facilitate conventional numerical simulations. ML methods also enable us to shed light on events and processes during collisions. This inverse problem can help us realise for the first time real-time event reconstruction or online physics analysis as it has been implemented by, for example, the CBM experiment, the STAR experiment at RHIC and the LHCb experiment at the LHC.

The application of ML techniques for real-time event reconstruction will be opportune to efficiently process raw, low-level, data from the large variety of detectors. The success of the in-situ data reconstruction and selection will depend crucially on the stability of individual sensor elements and accelerator operations. The usage of ML algorithms and AI in operating nuclear physics experiments and accelerators, enabling fast and vast data processing and analysis, is taking a central role, reducing the traditional dichotomy between offline and online computation.

Deep learning (DL), a subset of ML, has proven to be successful in representing complex processes. Its data-driven nature aligns well with nuclear physics, especially given the massive amount of data from experiments, observations, and simulations. Generative models, a prominent subset of DL, are designed to generate data samples resembling the input data. These models have found significant applications across various domains of nuclear physics. In Lattice Quantum Chromodynamics (QCD), flow-based models have now been introduced to improve the efficiency of Monte Carlo algorithms. The incorporation of gauge symmetry, a fundamental principle governing the strong force, has been successfully realised in lattice-gauge-equivariant convolutional NNs, which can also be applied to real-time glasma simulations, describing the earliest stages of heavy-ion collisions. Generative models have also been instrumental in nuclear experiments, aiding in data augmentation, noise reduction, and event simulation, relevant to optimizing particle accelerator operations. Currently, as the intersection of ML and nuclear physics deepens, tools like Generative Pre-trained Transformer (GPT) models are set to play a pivotal role in both experimental and theoretical advancements.

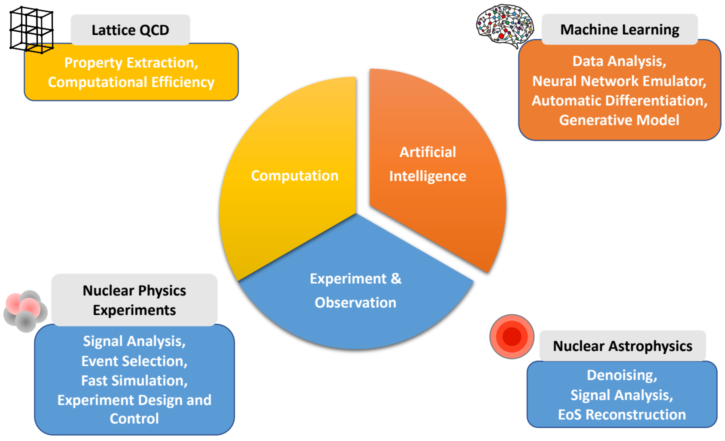

Figure 1: Chart of ML application in nuclear physics.

ML offers numerous benefits but also faces challenges like interpretability which is crucial in nuclear physics. However, physics-informed and physics-driven ML methods are providing promising solutions. Techniques like automatic differentiation are being employed to refine experimental designs, ensuring that the functional designs align with physical principles. It has also been used in tackling inverse problems when one wants to extract spectral functions from correlators calculated in Lattice QCD, build dense nuclear matter equation of states from neutron star observations, and obtain QCD matter properties from heavy-ion collision experiments. Additionally, constraint-based learning ensures that deep models follow physical laws, improving their reliability in, *e.g.*, simulating nuclear reactions or predicting material attributes.

Recommendations

- **Transform ML prototypes into applications for production.** Formulate a strategic approach to advance from current short-term, proof-of-concept ML projects in nuclear physics towards practical applications usable in production. This process begins with the identification of specific pilot projects that warrant conversion into long-term projects. Collaborative support from computer science research centres can aid the transformation process. In transitioning from pilot ML projects to production in nuclear physics, leveraging established ML frameworks and codebase management platforms is pivotal.

- **Fostering data sharing in nuclear physics.** Promote data-sharing initiatives in nuclear physics by establishing a user-friendly hub for open databases, akin to existing platforms in related scientific fields such as the *International Lattice Data Grid* (ILDG) [https://hpc.desy.de/ildg/], the *Gravitational Wave Open Science Center (GWOSC)* [https://gwosc.org/], and the *CERN Open Data Portal* [https://opendata.cern.ch/], or the *European Open Science Cloud* (EOSC) [https://eosc-portal.eu/] and *ESCAPE* [https://projectescape.eu/].
- **Strengthen computational resources.** Allocate funding for enhanced GPU clusters within established HPC centres across Europe. With the increasing demand for computational power in emerging fields like large language models, which require extensive computational resources, we anticipate a similar need for substantial computational capabilities in scientific applications. At the same time, it is crucial to support GPU access at various scales, ranging from centralised GPU centres designed for large-scale projects to localised GPU resources for the rapid development needs of individual research groups.
- **Training of scientific foundation models.** Recognizing the potential of foundation models in scientific research, we recommend dedicated investments in training and fine-tuning models tailored for scientific purposes, such as GPT models specialised for nuclear science.

2. Quantum Computing

Since the recent realization of the first quantum computer demonstrators, the field of quantum QC has seen remarkable growth worldwide. In the physics community, much of the interest is motivated by the exciting potential of these technologies to simulate strongly correlated many-body systems. There are two main platform approaches to this purpose, digital universal quantum computers and analogue special-purpose devices, called quantum simulators.

Despite the formidable challenges in designing algorithms suitable to harness the capabilities of these early prototypes, an increasing number of applications have been proposed and tested. This includes efforts from the nuclear theory community to study the structure and dynamics of many-body systems of fermions, such as high- and low-energy nuclear theory, gauge and field theories, and neutrino physics (see Figure 2).

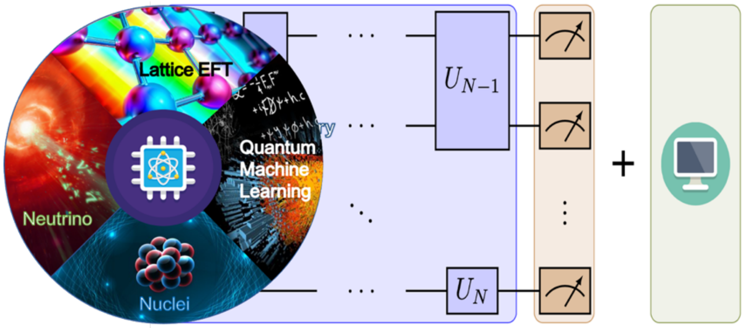

Figure 2: Schematic illustration of some of the fields where quantum computing is now being explored (left). Examples include physics of quarks, neutrinos and nuclei, as well as a pictorial view of quantum machine learning. All these problems are being now considered as pilot applications that could be treated on digital quantum computers using quantum circuits (right).

Another important effort is now being made on data mining through the development of advanced Quantum ML (QML) techniques, which blend two unique disciplines: quantum theory and ML. This novel approach provides the possibility of defining quantum and classical algorithms on classic or quantum data. Supervised and unsupervised learning together with reinforcement learning are still valid definitions within the QML framework, supporting the interdisciplinary collaboration of ML and QC experts. QML is thus becoming an excellent candidate to deal with the incoming computational tools outbreak caused by increasingly complex data.

Quantum information processing has conjointly seen rapid advances on the level of devices as well as our theoretical understanding of quantum many-body systems. These hand us paradigmatically novel computational methods, *e.g.*, in the form of tensor networks, and move fundamental questions such as the role of entanglement in quantum many-body systems into focus. The development of quantum computers and related aspects offer new perspectives and might rapidly revolutionise the way complex problems are treated together with the potential of recent advancements in quantum sensors which open the possibility of looking directly into quantum data.

Small-scale quantum computers are now widely available, often as computational devices accessible as cloud services. Complex quantum computers with hundreds of qubits are being commissioned or built by many academic institutions and industrial partners worldwide. We are currently in a transition period where the size and the fidelity of quantum resources are not sufficient for fault tolerance and decoherence is still an issue. The technological challenges associated with these issues are being addressed by the quantum computing community on the hardware and algorithm level, with substantial progress every year. Despite these difficulties, scientists in different fields have started to seriously consider (noisy) quantum computers as a disruptive technology that opens new horizons.

Internationally, the potential for game-changing applications of quantum technologies and quantum information theory has already been recognised. For example, in the USA it has given rise to a strategic effort towards quantum computing for nuclear theory through an influx of funding and a series of strategic white papers collecting community input and commissioned by government bodies. In Europe, several countries have decided to boost their national strategy on quantum technologies from which one can, in the coming years, kickstart real quantum device applications. Indeed, several relevant initiatives have been put in place, *e.g.*, by CEA, CERN (CERN Quantum Technology Initiative), CNRS, and INFN. Some of these initiatives have the goal of establishing joint research, setting up the supporting computing infrastructure, and providing dedicated mechanisms for the exchange of knowledge and innovation.

Moreover, several bottom-up initiatives have been made to promote quantum computing in the NuPECC community. These include European collaborative research grants (*e.g.,* Horizon 2020/Europe or Quantera), white papers, as well as the organization of workshops or doctoral training programs. Thanks also to such initiatives, quantum computing is currently gaining interest in the NuPECC community, with constantly emerging new groups all across Europe working on various subjects, (see Figure 2). These are encouraging first steps, but a wider consensus with dedicated infrastructure tailored to nuclear physics needs has not yet been achieved. While on a methodological and technical level, Europe is on par with the best international players, a concerted strategy that bundles and streamlines individual efforts is lacking.

Recommendations

A sustained effort is necessary to advance QC towards realistic applications. Among the requirements for future applications of QC, we highlight the need to:

i. Test and validate the applicability of state-of-the-art algorithms and methods developed in the general theory of quantum computing to fields of direct interest to NuPECC.
ii. Define good benchmarks to leverage cutting-edge QML solutions to problems of interest for NuPECC.
iii. Design new algorithms and new quantum ansatzes tailored for instance to complex many-body systems of relevance for nuclear physics, neutrino physics, Lattice field theory, and/or Quantum Machine Learning.
iv. Perform applications on existing quantum devices with an appropriate effort dedicated to contrasting the presence of noise using a variety of error mitigation/corrections techniques.

In this respect, a major objective to be addressed soon is the development of new quantum algorithms and error-mitigation techniques that can efficiently exploit the capabilities of current (noisy) and future quantum computers. In parallel to these efforts, a set of recommendations meant to foster collaborations, strengthen expertise, and address emerging challenges within the NuPECC quantum computing community to accelerate advancements in this rapidly evolving field are given below.

- **Establish a transnational European network on quantum activities.** Create a collaborative network at the European level, focused on QC and quantum information (QI), in alignment with the interests of NuPECC. Foster cooperation and knowledge exchange among researchers from different institutions and countries. Bridge QC theory, QI processing, and machine learning, enabling cross-community collaboration, research, and development in future technologies and knowledge-sharing.
- **Organise workshops, schools, and training programs.** Regularly host workshops, schools, and hands-on events to facilitate the transition towards quantum computing [see TWG 10]. Consider partnering with academic institutions to strengthen their role in QC and provide comprehensive training opportunities. Address the emerging topic of QC applications by providing specific training for young and experienced scientists/engineers in quantum mechanics, quantum information, and quantum algorithms. Prioritise the training of researchers in quantum computing. Beyond workshops, explore measures such as student exchanges and joint fellowships to build a strong interdisciplinary network.
- **Facilitate access to quantum platforms.** Ensure access to state-of-the-art quantum platforms by bridging the gap between academic institutions and private companies. Consider forming agreements with national High-Performance Computing (HPC) centres like CINECA, Jülich, GENCI, etc., to enhance accessibility. These include agreements for the use of classical simulators (*e.g.*, based on tensor network codes) to benchmark quantum hardware. Another key aspect of accessing certain quantum platforms is to promote collaboration between theorists and experimentalists developing current and future quantum devices to collaboratively design application-specific quantum-simulation algorithms, capitalizing on their combined expertise.
- **Develop strategies for quantum-classical interfaces.** Formulate a clear strategy for interfacing quantum and classical machines. As most algorithms will be hybrid quantum/classical, effective integration is crucial for future developments. This includes

addressing data mining challenges given the large flow of data generated by quantum computers with increasing qubits by developing specific techniques to extract information from the data through classical post-processing. Also, develop common open-source libraries for a frictionless integration of classical and quantum software frameworks.

3. Numerical Tools and Techniques

A comprehensive and exhaustive treatment of the numerical tools and techniques that have been and are being employed in nuclear and high-energy physics is an enormous task that is impossible to exhaustively fit in this report. However, it remains relevant to discuss some prominent cases that will eventually help us formulate a set of general recommendations concerning the development of new numerical technologies aimed at advancing the field of nuclear and high-energy physics. To be as representative as possible, we will distinguish between computing techniques applied to nuclear and high-energy physics from a theory viewpoint on the one hand, and from an experimental viewpoint on the other. On the theory side, we will consider transport approaches in heavy-ion collisions, the study of the hadronic structure, and computations on the lattice. On the experimental side, we will discuss the development of Monte Carlo generators aimed at developing detector systems.

a. Heavy-ion collisions and hadronic structure

Presently, we are observing a distinctive change of paradigm:

- Experiments, especially those at RHIC and the LHC, have moved on from an exploratory phase to a high-precision phase.
- The theoretical approaches evolved from simple models to describe single observables to more complex frameworks aiming at the simultaneous description of all observables.

In several contexts ranging from transport approaches, hydrodynamics, and thermalisation in heavy-ion collisions [see TWG 2] to the study of the nucleon structure [see TWG 1], these developments have led to advancements of unprecedented complexity. Developing them further requires software specialists, a new structure of the community (by regrouping and joining efforts) and new computational approaches. This aspect is often not addressed in funding schemes. In addition, the underlying theories, like lattice QCD or perturbative QCD, do not usually provide predictions in the whole kinematic range which is studied in heavy-ion collisions. Therefore, forceful extrapolations and approximations are often necessary. To make progress in the understanding of the physics questions at hand, an assessment of the different approaches is necessary.

This is a worldwide problem, which has already been successfully addressed in the USA with the creation of structures like JETSCAPE [https://jetscape.org/], MUSES [https://muses.physics.illinois.edu/] or HEFTY [https://hefty.tamu.edu/]. These topical structures typically have a budget of 5 million US dollars for 4 years. In these well-funded structures, several research groups (typically between 10 and 15) collaborate on a given topic. The different approaches to a given topic have been centralised and manpower made available to make the programs comparable, to benchmark the results, and to subdivide them into different modules. Structures of that type should be created in Europe as well.

The comparison of high-precision data to theoretical predictions of these advanced models needed to answer the essential physics questions at stake and that aims at an accuracy of 10%, must rely on an identical treatment of experimental data and numerical results from transport approaches. This needs a forum where the numerical data is made available and where the software tools for such comparisons are developed and maintained.

The study of the hadronic structure in the context of high-energy collisions has also witnessed an impressive development catalysed by experimental, theoretical, and technological advances. The quantity and quality of experimental results delivered by past and current facilities have pushed the need for a detailed quantitative understanding of the internal structure of hadrons to an unprecedented level. The community has reacted by producing more accurate theoretical calculations and by developing numerical tools aimed at streamlining the interpretation of experimental results. In many cases, ML techniques have also come to the rescue by enhancing our ability to manage complex data structures.

These successful efforts have been widely recognised and supported. At the European level, the STRONG-2020 project [http://www.strong-2020.eu/scientific-frontiers/high-energy-frontier.html] has invested significant resources in the virtual-access packages NLOAccess and 3DPartons devoted to the development and maintenance of open-source numerical frameworks for high-precision phenomenology at high-energy colliders. Similarly, the Centre for Nuclear Femtography [https://www.femtocenter.org] in the US has been recently instituted with similar purposes. Another stepping stone towards a detailed understanding of the hadronic structure is the recent approval of the Electron-Ion Collider (EIC) to be built in the US within the next decade. This initiative has boosted even further the demand for precision phenomenology at high-energy colliders that now more than ever needs a more incisive organisation and support.

Recommendations

- The creation of a new, European-wide structure which hosts, maintains, and develops open-source numerical tools for the physical interpretation of present and future heavy-ion and collider data. These structures should: bring together the community, allow for a detailed comparison of different approaches, and provide support in software engineering for future development. US efforts like JETSCAPE could serve as a model.
- The current efforts to provide a solid basis for systematic comparisons of the experimental results to theoretical predictions should be intensified. This requires a European-wide platform where theoretical predictions are stored and where open-source analysis tools (like RIVET) are developed to facilitate detailed comparisons with data.
- International collaborations should be maintained and extended.
- Access to substantial computational resources and funding for the development of optimised algorithms and codes should be guaranteed.
- Further support and centralisation of the initiatives aimed at providing open-source numerical frameworks devoted to the study of the hadron structure should be supported.

Highlight: Functional methods

The functional approach to QCD via Dyson-Schwinger equations (DSEs), functional renormalisation group equations (FRGs) and Bethe-Salpeter equations (BSEs) works at the level of non-perturbative correlation functions. In principle, these objects encode all information on the physics content of QCD. This includes the conventional and exotic hadron spectrum of ground and excited states, form factors, decays and electromagnetic as well as hadronic processes like Compton scattering and pion-nucleon interaction. Approximation schemes at very different levels of sophistication have been developed which approach a versatile tool: on the one hand, simple approximations allow us to make contact with quark-model calculations; on the other hand, highly sophisticated and numerically demanding schemes allow for a direct comparison with lattice QCD. A very recent example is the self-consistent and parameter-free calculation of the glueball spectrum shown in Figure 4 which demonstrates the excellent agreement of both approaches. Functional methods are furthermore fully covariant and do not rely on heavy quark expansions. The approach can be used in the light- and heavy-quark region allowing for fruitful interactions with chiral perturbation theory as well as with heavy-quark effective theory. Systematic comparisons with other approaches have a high potential for the identification of the physical mechanisms behind the observable phenomena.

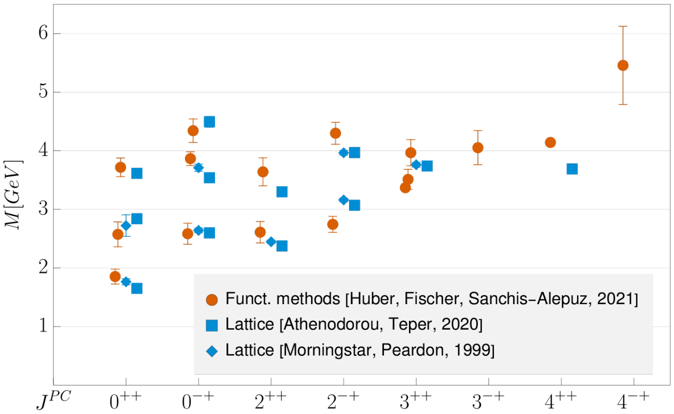

Figure 2: Glueball spectrum of pure Yang-Mills theory determined on the lattice and with functional methods.

After two decades of advances in terms of both high-quality truncations and applications to three (baryon, hybrid), four (tetraquarks) and multi-body problems including scattering processes, many current projects in the functional community have reached a level where the bottleneck is computational power.

b. Lattice QCD

Numerical simulations of Quantum Chromodynamics (QCD) formulated on a discrete space-time lattice have become an essential tool in strong-interaction physics. Lattice QCD provides a rigorous framework for defining and solving QCD non-perturbatively that is systematically improvable, thus providing stringent tests of QCD [see TWG 1]. Lattice QCD aims to provide quantitative information on nuclear and hadronic properties in terms of the fundamental constituents of matter and their interactions. Physical properties are extracted from correlation functions which are evaluated through high-dimensional integrals over quark and gluon fields (with the number of degrees of freedom depending on the size of the space-time grid). The integrals over the gluon fields are estimated via Markov-chain Monte Carlo such that predictions for physical quantities have an associated statistical error. Representative sets (ensembles) of the gluon fields are generated for lattices with a given lattice spacing and finite volume and, due to the computational expense, sometimes for unphysically large light-quark masses. The simulations are repeated for different lattice spacings, volumes, and quark masses and physically relevant results are obtained by extrapolation to the continuum and infinite volume, and physical quark-mass limits. Controlling the systematics associated with these limits requires, depending on the observable, pushing the simulations towards ever smaller lattice spacings and larger volumes. Current applications include: exploring the phase diagram of strongly interacting matter, and studies of hadron spectroscopy and structure as well as of the properties of small nuclei. Since the LRP in 2017, there has been tremendous progress in lattice QCD. This has resulted from the development of new theoretical approaches, simulation algorithms and solvers as well as the possibility to access large pan-European supercomputing resources enabled by PRACE and EuroHPC Joint Undertaking (JU) and supplemented by national computing facilities.

Due to the complexity of the calculations and the huge computational resources required, the lattice community is largely organised into European and worldwide collaborations, each with its code bases some of which are freely available. Examples are Chroma, the Columbia physics system (CPS) and MIMD Lattice Computation (MILC), which are application packages which utilise the USQCD software suite and the GRID, OpenQCD, and twisted-mass LQCD (tmLQCD) software packages [see TWG 1 for a list of collaborations]. Such community efforts need to continue to be supported in the future.

Frontier research in lattice QCD necessitates high-performance computing capabilities and continuous access to large computational resources. Therefore, for European lattice QCD groups to maintain their competitiveness and achieve their scientific objectives, access to substantial computational resources on GPUs and CPUs (of the order of billions of core hours on conventional CPUs) is an absolute requirement. In particular, in terms of the hardware, the (mostly) double precision floating point performance must be matched with adequate memory and internode communication bandwidth. Furthermore, to effectively harness the available computing power, it is crucial to develop highly optimised algorithms and codes. This means that funding experts to develop appropriate algorithms that can efficiently run on new computer architectures is particularly important, as is the education and training of early-career researchers. Lattice QCD groups have a reputation for being early adopters of cutting-edge computer technologies. We anticipate that these groups will continue to embrace new technologies and become users of exascale infrastructure.

With the advent of exascale facilities, more long-term storage for the gauge ensembles needs to be made available in the future. The need for coordinated data handling has led to the formation of the International Lattice Data Grid (ILDG), with several regional sub-grids, including the European LDG. Support for coordinated data management and storage needs to be intensified in the coming years.

Emerging techniques based on lattice QCD have recently broken through in the field of hadronic structure. Indeed, they offer the long-sought possibility to extract information on partonic distributions defined on the light cone through lattice computations. This challenge has now been taken on by many groups in Europe and all around the world leading to significant advances. This effort needs to be supported and organised at a European level to optimise the large numerical resources required to push it forward and make it effective.

Recommendations

Considering the above, we recommend the following:

- International collaborations and freely available software packages for lattice QCD should be maintained and extended.
- access to substantial computational resources fitted to the need of lattice computations (double precision floating point performance, adequate memory and internode communication bandwidth) is essential.
- it is important to fund the development of optimised algorithms and codes.
- Call for more long-term storage solutions for gauge ensembles, particularly with the advent of exascale facilities.
- Coordinated data management provided by the ILDG and the European LDG should be supported.
- Support for and coordination of the many European groups devoted to the study of hadronic structure on the lattice.

c. Low-energy nuclear structure

Nuclear structure theory aims to describe the properties of nuclei and nuclear matter by solving the many-nucleon Schrödinger equations either in an effective or in an *ab initio* fashion [see TWGs 3 and 4]. The latter strategy has witnessed considerable progress in the last decade, with calculations extending beyond light nuclei to more and more complex systems. Current frontiers include pushing ab initio calculations to heavy and/or deformed isotopes, developing accurate and systematically improvable nuclear interactions and many-body techniques, as well as bridging nuclear structure and reactions.

These open issues involve specific computational challenges. The modelling of nuclear forces requires sophisticated fitting procedures in multi-dimensional parameter spaces. On top of it, meticulous Bayesian analyses are currently deployed to gain insight into the associated statistical and systematic uncertainties. Many-body techniques typically involve large-scale diagonalisations or iterative solutions of tensor networks, where the dimensions of the matrices or tensors increase with the required accuracy. Applications to nuclear reactions further augment the size and complexity of the objects to be manipulated.

This workflow entails CPU- and memory-greedy algorithms that are required to fully exploit the potentialities of state-of-the-art high-performance clusters. We envisage that in the coming years, many of the existing codes will be ported to exascale machines. Furthermore, in the longer term, the possibilities offered by machine learning and quantum computing will be increasingly explored by the community.

In a context where algorithmic, computational, but also conceptual boundaries are rapidly progressing, collaborative initiatives will be key for systematic and sound advances in the field. In this respect, crucial aspects include the public availability of numerical codes, and the involvement of computer scientists and applied mathematicians, possibly leading to additional, dedicated manpower to support nuclear physicists.

Recommendations

- Promote cross-fertilisation from/to other fields utilising many-body methods and applied mathematics.

- Enhance and systematise collaborations with computer scientists to exploit new computer architectures.
- Support virtual-access facilities such as Theo4Epx of EURO-LABS.
- Encourage the sharing and publication of numerical codes.

d. Computing Techniques in Experiments

An extremely valuable numerical tool in experimental nuclear and high-energy physics is Monte Carlo (MC) event generators. For example, MC simulations of radiation transport are used for the development and optimisation of detector systems and to aid in the interpretation of experimental results. MC generators rely on the generation of pseudo-random numbers which is typically a CPU-bound task. In this respect, a centralised computing server with very high CPU power would be a valuable resource for such applications. Since MC simulations are intrinsically parallelisable, multi-core systems can be exploited. However, this sometimes requires significant work from the user to implement, so such a system would ideally be chosen for per-core CPU performance.

Adding a powerful multi-core CPU and larger memory to the data storage and backup systems mentioned above would allow users to log in remotely and execute data analysis codes directly on the server. Such a system would reduce the need for time-consuming transfers of large data samples between multiple locations and allow sorting to be completed and/or iterated more quickly. Such systems have already been used with great benefits in high-energy physics and could be equally beneficial in nuclear physics.

A centralised data storage and a backup facility would help the nuclear physics community secure large data sets collected at great expense from laboratories all around the world. In addition, it would help scientists meet their responsibilities under open-access data-sharing agreements. Such a system would need a very large array of redundant hard disks for storage but would only have basic requirements for CPU and computing power.

Recommendations

Given the above, we recommend what following:

- Deployment of simulation software such as GEANT4, MNCP, and FLUKA on fast computer systems to enable the generation and analysis of large numbers of events and the simulation of complex experimental setups.
- Adaptation of physics models or development of dedicated interfaces to extend the range of physics applications that event generators can currently deal with. Complex theoretical models can be CPU or memory-bound depending on the specific case. A centralised computing resource with high total and per-core CPU performance together with very large RAM is valuable for the nuclear physics community.

4. Resources and Infrastructure

To address the challenges posed by the end of Moore's law, the industry has embraced multi-core processors (Microprocessor Trend Data [https://github.com/karlrupp/microprocessor-trend-data]). Modern CPUs commonly have homogeneous multiple cores, enabling parallel processing of tasks. While this approach provides better performance, it has also highlighted the need for software optimization to fully exploit the potential of multiple cores. More specialised hardware components, the coprocessor units, are needed to perform specific tasks more efficiently than the general-purpose CPU cores.

Historically, CPUs, GPUs, and FPGAs have been separate entities within a computing system. As computational demands have grown, the need for efficient and flexible architectures that can handle diverse workloads has become increasingly important. This has led to the development of chiplets, also called tiles, which represent a novel approach to CPU design (Universal Chiplet Interconnect Express (UCIe) [https://www.uciexpress.org/]). Instead of monolithic, single-die processors, chiplets break down the CPU into smaller, specialised components that can be interconnected on a single package. This approach allows for improved modularity, scalability, and cost-effective manufacturing where different types of chiplets, such as CPU, I/O, and GPU ones, are already available or will soon be available on the market.

Exploiting the parallel processing capability with such a heterogeneous variety becomes crucial for simulating more and more complex nuclear models, handling the vast amounts of data generated by experiments and training the deep neural networks that are increasingly used in nuclear physics. On the other hand, developing programs to fully exploit the parallelism and heterogeneity poses several challenges, where programming models that abstract the underlying hardware complexity and enable efficient use of resources and portability across different architectures are required. Fortunately, there are several models for writing portable applications on the market such as Alpaka, Kokkos, oneAPI and SYCL.

The advent of HPC exa-scale infrastructure and QC offers European researchers an unprecedented processing capacity. This complements and integrates the capabilities provided by research High-Throughput Computing (HTC) and Cloud facilities at both national and European levels. However, exploiting HPC centres comes with several challenges, such as the diversity in access and usage policies, and the heterogeneous and different computing architectures. To address these challenges and create a cohesive data processing system, the SPECTRUM project, funded by the EU, aims to integrate different European computing resources. This includes on-premises data centres, HPC clusters, and Quantum nodes. The ultimate long-term objective is to establish a European exabyte-scale research data federation and a seamless compute continuum.

As the computational demands within various nuclear-physics communities continue to grow, the performance, accessibility, and interoperability of IT infrastructures have become indispensable elements for ensuring the success of research endeavours. One of the central challenges to address is the *heterogeneity* in the usage of computational tools, driven by the evolving technological demands and the diverse needs of research communities.

Future experiments are designed to handle data in situ, without the need for hardware triggers, at high rates, while leveraging sophisticated high-level event selections. The effective operation of these experiments hinges on the successful deployment of often intricate data processing algorithms that exhibit a high level of scalability. These algorithms must run seamlessly on a variety of architectures, including FPGA, GPU, ARM, and conventional x86 processors. This imposes rigorous demands on the IT infrastructure supporting these experiments, as well as on the quality, reproducibility, and portability of the software, frequently developed and maintained by young researchers collaborating in the scientific field.

Similarly, forthcoming theoretical activities will become even more demanding, necessitating precision calculations for complex many-body systems to achieve significant breakthroughs in research. These activities also require the use of a diverse array of IT infrastructures, spanning from high-performance computing centres utilising conventional x86 processors in tandem with GPU clusters to the potential integration with future quantum computers.

The high data rates received (expected) for the next generation of particle-physics experiments (*e.g.* new experiments at FAIR/GSI and the upgrade of CERN experiments) call for dedicated attention concerning the design of the computing infrastructure needed online and offline [see TWG 6]. The traditional separation between DAQ, online and offline vanishes and one has to consider all these stages in data processing together; *i.e.* the traditional DAQ/Trigger designs are not able to handle the amount and kind of data which will come from the detection systems, and therefore more data processing is needed online. This change in the role of DAQ/Trigger systems will introduce more software-based components into the DAQ/Trigger systems than ever before. Most of the traditional hardware-based triggers are going to be (being) replaced with systems based on commodity hardware processors and co-processors running standard software algorithms. Handling such systems includes not only the design of efficient and scalable algorithms, but also the development of key software building blocks for ultrafast data processing on large-scale heterogeneous computing infrastructures, synchronization of multiple data streams, transport services, container orchestration, and efficient binding to storage and network.

Recommendations

- Technology evolution demands significant investment in software frameworks, that need to support parallelism on multi- or many-core CPU and the execution of algorithms on heterogeneous platforms.
- To harness the synergies among research projects undertaken by diverse collaborations and across various IT infrastructures, it is highly beneficial to promote the implementation of a federated computing model. Emphasis should be placed on the accessibility of resources (data, software, and computing) and facilitating interoperability among IT activities within the nuclear-physics community and closely aligned fields, including particle and astroparticle physics.
- Developing interfaces located between the system software layer (i.e. storage, RDMA networks, container runtime, etc.) and the core software developed by the domain scientist implementing the core of the scientific analysis requires collaboration between IT experts and domain scientist, the community should encourage and actively implement such collaborations.