HPC community access: Lattice QCD Community

How is your community structured? Who represents the community?

The lattice QCD community in Europe is comprised of several collaborations, as well as smaller groups. Together they are part of an international research effort, and many of the larger collaborations have participating researchers from outside Europe. Major collaborations are ALPHA, BMW, CLS, ETMC, HotQCD, QCDSF, and UKQCD, each with its own simulation and analysis codes. The community has a long history in working together in coordinated, EU funded research programs. The first such program was the European Community Network "Computational Particle Physics" (1993-1997). Currently members of the lattice QCD community coordinate the European Network for Particle Physics, Lattice Field Theory and Extreme Computing (EurPLEx), and participate in a number of EU projects, including STRONG-2020 funded under the European Integrating Activity for Advanced Community initiative.

Representatives of all the aforementioned collaborations and selected groups have been contacted and a consensus has been reached to establish a five-member working group to respond to the request for input by PRACE, consisting of C. Alexandrou (Cyprus), L. Del Debbio (UK), F. Karsch (Germany), S. Ryan (Ireland) and H. Wittig (Germany).

At the beginning of 2020, a meeting is planned to which all European lattice QCD researchers will be invited in order to put in place a commonly accepted structure and representation of the community.

What is the scientific case? (with sharp goals)

To describe the structure and interactions of hadrons and light nuclei from QCD: How hadrons emerge from QCD and how they interact are fundamental questions driving a rich experimental research program at the LHC, CERN, Jefferson Lab and Brookhaven National Laboratory in the US, the GSI and MAMI/MESA in Germany, J-PARC in Japan and BESII in China. The Electron Ion Collider planned (EIC) in the US will be the next facility for the QCD frontier to map the dynamics of quarks and gluons, and lattice QCD simulations are central to that program. The goals are: i) precision computation of the mass, spin decomposition, size and 3-D structure of the proton, ii) computation of the properties of resonances and exotics, and hadron-hadron interactions, and iv) determination of the structure of light nuclei.

To understand the fundamental symmetries determining the structure of the universe and search for new interactions Beyond the Standard Model (BSM): Discovering the basic forces of nature and understanding their consequences are at the forefront of physics research. Among the fundamental questions dictating high energy and precision experiments at CERN, Fermilab, Mainz, LBNF, etc is the puzzle why there is more matter than antimatter, what the nature of dark matter is, and whether there is physics beyond the Standard Model. The goals are: i) precision computation of the muon anomalous magnetic moment, axial nucleon couplings and form factors for neutrino-nucleon scattering, neutron electric dipole moment and nucleon σ-terms, ii) weak decays of charm and bottom in the meson and baryon sectors, mixing in the B- and D- meson sector and CP violating form factors that provide constrains on the CKM matrix and BSM physics, iii) determination of the parameters of BSM models, such as composite Higgs, Dark Matter, new global symmetries, different gauge groups and matter representations.

To understand the phases of strongly interacting matter and the properties of the quark-gluon plasma (QGP): This is one of the key missions of the nuclear physics program in Europe and internationally. It aligns with the experimental research programs at the LHC at CERN, (ALICE, ATLAS, CMS, LHCb) and the Relativistic Heavy Ion Collider (RHIC) at BNL in the USA as well as FAIR at GSI, Darmstadt, which will start operation in the middle of the next decade with a large fraction of its research program dedicated to the exploration of the QGP. Goals are: i) Establishing accurately the critical parameters for a transition from ordinary hadronic matter to the QGP, ii) determining the dependence of these parameters on the net baryon-number, (iii) establishing whether or not the theoretically conjectured and experimentally searched for critical point exists, (iv) calculating the modifications of hadron properties in hot and dense matter, and (v) quantifying the transport properties of the medium itself.

What resources can you use now? And what would you need to reach your goals?

Lattice QCD groups are known to be early adopters of new computer technologies. In the first 5 years of PRACE Tier-0 calls lattice QCD groups have used on the average 20% of the resources. In later calls this percentage dropped as the participation of other communities started to pick up. Since PRACE started to provide access to Tier-0 systems in 2009, 52 lattice QCD projects have been awarded grants amounting to 2185.21 million core-hours and 397 million GPU-hours. We expect that lattice QCD groups will similarly be early adopters of exascale infrastructure and a number of groups are already preparing highly optimized codes. Thus, they form an important community to have on board.

To achieve several major scientific breakthroughs will require community access and exascale resources. Three examples are:

- The contribution of gluons to the nucleon mass, and the calculation of the low moments of the spin-averaged and spin-dependent gluon distributions, very relevant for the EIC.
- The neutron electric dipole moment important for understanding the presence of CP violation whose origin must lie beyond the Standard Model, and more generally constrains models of physics BSM.
- Establishing the structure of the QCD phase diagram and reaching quantitative conclusions on the location of a critical point in the phase diagram at high baryon number density is of utmost importance for the international experimental heavy ion research programs.

What is the portfolio or type of projects (Excellence vs Production)?

Lattice QCD research heavily relies on access to leadership computers and for this reason it has been at the frontier of HPC. An integral part of lattice QCD research has been the development of architecture specific, highly optimized algorithms and codes to best exploit the available compute resources. Frontier research in lattice QCD typically needs speed and throughput and continuous access to large computational resources.

A lattice QCD calculation is comprised of three stages: i) generation of gauge field configurations; ii) generation of quark propagators; and iii) construction of correlation functions and other complex propagator combinations needed for the computation of physical observables. The first two stages are computationally very demanding and require continuous access to large computational and storage resources. The generation of gauge fields is sequential and thus relies on access to a large number of cores or GPUs demanding good strong scaling of algorithms and codes. The computation of propagators is done on each gauge field configuration independently and it thus relies on throughput.

The community access modes that would benefit the lattice QCD community are thus twofold:

A. **Excellence driven research**: This follows the traditional PRACE project allocation approach, which is based on scientific excellence. This high-risk – high-gain access mode has an excellent track record across the lattice QCD community. It works very well for the computation of targeted observables utilizing novel ideas.

B. **Long term production-type research:** The production of gauge field configurations and the computation of a large number of quark propagators require long-term continuous access. Such access is needed in order to reach the precision required for impacting the interpretation of experimental results and for probing new physics beyond the standard model. Large production runs on pre-exascale and exascale machines will generate large amounts of data in the hundreds of TB. Managing and analyzing these data is a challenge that will require workflow technologies similar to those used by e.g. experimental groups at CERN.

A community access mode allowing continuous access to large computational resources, in addition to the excellence-driven allocations based on the PRACE model (access mode A), would be essential for ensuring that the European lattice QCD community remains competitive.

The mode access A has been well defined and successfully delivered by PRACE. It should be maintained allowing each group to submit projects.

Community access mode B will provide additional resources for long-term objectives of the community. We stress here that the community would allow the various collaborations using different approaches and codes to pursue the scientific program defined in the white paper that the community is to prepare. The earmarked resources should be allocated among the various collaborations and groups within the lattice QCD community according to a peer review process following similar guidelines as for access mode A.

The mechanism for distributing resources to the various communities needs a well-defined and transparent process that takes into account research excellence. Lattice QCD pursues fundamental research with socio-economic impact characteristic similar to that of other pure research fields, i.e. high impact that typically is made obvious indirectly or later in time. The research results provided by the community, however, often have immediate impact on international experimental programs.

The lattice QCD community strongly believes that the guiding principle of any community access model must be to promote excellent science and to optimally utilize the computational infrastructure.

What is your level of readiness for exascale? And what do you need to be ready for exascale?

The lattice QCD community has an excellent track record of being ready to use new generations of supercomputers. However, this has been achieved in the past in a co-design approach at both hardware and software levels through early access to and even involvement in the design of new hardware. In particular, it is evident that the exascale hardware will be novel, involving also novel software features. This will require the targeted development and optimization of new software.

In the US this has been realized early on and lattice QCD groups there have already been involved for more than two years in the development of code for the first US exascale computers, with access to inside knowledge about this future hardware. If the European lattice QCD community is to be ready for European exascale computers with efficient code that can immediately exploit the available resources, it needs to have access to prototype hardware or hardware simulators. In addition, training and creating a career path for a new generation of domain-specific scientists who can drive the new algorithmic and code development is necessary to ensure the best exploitation of exascale facilities.

What is the potential of Machine Learning and Deep Neural Networks for the applications?

Although machine learning (ML) has only recently been applied in lattice QCD it is emerging as a potentially powerful tool in lattice QCD simulations, production of correlators and their analysis. ML approaches, such as deep neural networks have been successfully applied to the determination of order parameters and the study of phase transitions, including the deconfinement transition in SU(2) Yang-Mills theory at finite temperature. The development of further application of such approaches is ongoing research work. This includes the finding of optimized parameters that speed up the simulations and the generation of gauge configurations, the computation of quark propagators and correlators and the extraction of observables in an inverse problem analysis.