Project ID **bep00085** — Initial Application

QCD+QED Simulations with C* Boundary Conditions

Application for CPU-time on SKL Systems at HLRN

RC$\star$OR collaboration

Prof. Dr. **Agostino Patella** (principal investigator)
Institut für Physik, Humboldt Universität, Newtonstr. 15, 14289 Berlin
DESY, Platanenallee 6, 15738 Zeuthen

Dr. **Andrea Bussone**
Institut für Physik, Humboldt Universität, Newtonstr. 15, 14289 Berlin

Dr. **Isabel Campos**
Instituto de Física de Cantabria & IFCA-CSIC, Avda. de Los Castros s/n, 39005 Santander, Spain

Madeleine Dale
Università di Roma Tor Vergata, Dip. di Fisica, Via della Ricerca Scientifica 1, 00133 Rome, Italy
INFN, Sezione di Tor Vergata, Via della Ricerca Scientifica 1, 00133 Rome, Italy

Dr. **Patrick Fritzsch**
Theoretical Physics Department, CERN, CH-1211 Geneva 23, Switzerland

Dr. **Martin Hansen**
SDU eScience Center, University of Southern Denmark, Campusvej 55, DK-5230 Odense M, Denmark

Dr. **Marina Krstic Marinkovic**
Fakultät für Physik, Ludwig-Maximilians-Universität Munich, Schellingstrasse 4, 80799 Munich

Jens Lücke
Institut für Physik, Humboldt Universität, Newtonstr. 15, 14289 Berlin
DESY, Platanenallee 6, 15738 Zeuthen

Prof. Dr. **Nazario Tantalo**
Università di Roma Tor Vergata, Dip. di Fisica, Via della Ricerca Scientifica 1, 00133 Rome, Italy
INFN, Sezione di Tor Vergata, Via della Ricerca Scientifica 1, 00133 Rome, Italy

1 Overview of the application

Isospin-breaking and QED radiative corrections to hadronic observables are generally rather small but they become phenomenologically relevant when the target precision is at the percent level. For example, a recent review [**?**] of the results obtained by the different lattice groups shows that leptonic and semileptonic decay rates of π and K mesons are presently known at the sub–percent level of accuracy. At the same time, QED radiative corrections to these quantities are estimated to be of the order of a few percent, by means of chiral perturbation theory [**?**], and must be included at this level of precision.

This project is the first step of a long-term research programme aiming at calculating isospin-breaking and QED radiative corrections in hadronic quantities from first principles in QCD+QED. The signature of the proposed project is the use of C* boundary conditions [**? ? ? ?**] which allow for a local and gauge-invariant formulation of QED in finite volume and in the charged sector of the theory [**? ?**]. In particular, full QCD+QED configurations will be generated at various values of the fine-structure constant α_R (including $\alpha_R = 0$) in such a way that physical observables can be interpolated at the physical value of $\alpha_R \simeq 1/137$. Important technical outcome of this project will be a first study of autocorrelations and finite-volume effects in the proposed setup, as well as a quantitative measurement of the effectiveness of the proposed tuning strategy to identify lines of constant physics. The generated configurations will be used to explore a variety of observables: from baryon correlators and masses to decay rates of mesons in QCD+QED.

The open-source `openQ*D-1.0` code [**?**] will be used to generate gauge configurations. This code has been developed by the RC* collaboration (several investigators of this application are among its developers). It is an extension of the `openQCD-1.6` code [**?**] for QCD.

We plan to generate 2 ensembles (`QCD1,2`) of QCD configurations with lattice spacing $a \simeq 0.054$ fm, and 3+1 flavours, i.e. three light degenerate quarks with a mass corresponding to $M_\pi = M_K \simeq 410$ MeV and a charm quark with a mass corresponding to $M_D \simeq 1870$ MeV. The two ensembles will differ by their volume, one with $M_\pi L \simeq 3.6$ and one with $M_\pi L \simeq 5.4$. We plan to generate also 2 ensembles (`Q*D1,2`) of $N_f = 1{+}2{+}1$ QCD+QED configurations with similar lattice spacing, unphysically large fine-structure constant $\alpha_R \simeq 0.05$ (in order to amplify isospin-breaking effects), and quark masses satisfying $m_u < m_d = m_s < m_c$ which correspond roughly to $M_{\pi^\pm} = M_{K^\pm} \simeq 425$ MeV, $M_{K^0} \simeq 390$ MeV, and $M_D \simeq 1870$ MeV. The most important parameters for the 4 runs are summarized in table 1. The strategy behind the choice of these parameters is motivated in section 3. All runs will start from (almost) thermalized configurations which will be generated on different machines.[1] Notice that some of the parameters for the `Q*D1` and `Q*D2` runs need to be tuned. This will be done in the electroquenched setup (i.e. assuming electrically neutral see quarks) on the `QCD1` and `QCD2` ensembles. The observables that are needed for the tuning are significantly less expensive than the generation of the configurations and will be calculated on different machines. The workflow is described by the GANTT diagram in figure 1.

A summary of the cost analysis for the proposed application is presented in table 2. We therefore apply for computer time on the Göttingen MPP system for a total amount of

$$\textbf{16.8M core}\times\textbf{hours on SLK node} = \textbf{0.42M (SLK node)}\times\textbf{hours} = \textbf{2.52M NPL}$$

with the following breakdown per quarter

$$\textbf{2.52M NPL} = \textbf{672k} + \textbf{573k} + \textbf{514k} + \textbf{761k NPL}\ .$$

After this initial study, we plan to apply for computer time in continuation projects to simulate at some intermediate value of α, and at values of the meson masses closer to the physical ones. We foresee that this phase of the project will take up to 3 years. After this is done, the proposed strategy will be reassessed and perfected before moving forward (e.g. to finer lattice spacings).

[1] Members of the collaboration have limited computer time on the following HPC systems: Altamira (provided by IFCA at the University of Cantabria), FinisTerrae II (provided by CESGA, Galicia Supercomputing Centre), Marconi (provided by CINECA, under the initiative INFN-LQCD123).

ensemble	volume	β	α	κ_u	$\kappa_d = \kappa_s$	κ_c	$c_{\mathrm{sw,SU(3)}}$	$c_{\mathrm{sw,U(1)}}$
QCD1	60×32^3	3.24	0	0.13440733	0.13440733	0.12784	2.18859	0
QCD2	80×48^3	3.24	0	0.13440733	0.13440733	0.12784	2.18859	0
Q*D1	60×32^3	3.24	0.05	to be tuned			2.18859	1
Q*D2	80×48^3	3.24	0.05	to be tuned			2.18859	1

Table 1: Simulation parameters for the four proposed runs. The masses in the Q*D runs must be tuned as described in section 3. All ensembles will be generated with C* boundary conditions in space and periodic boundary conditions in time.

ensemble	volume	nodes	proc. grid	local lattice	time per MDU	total time	total cost
QCD1	60×32^3	32	$10 \times 8 \times 4^2$	6×8^3	800 s	37 d	1.2 Mch
QCD2	80×48^3	32	$10 \times 8 \times 4^2$	8×12^3	3550 s	164 d	5.1 Mch
Q*D1	60×32^3	32	$10 \times 8 \times 4^2$	6×8^3	1350 s	63 d	2 Mch
Q*D2	80×48^3	64	$10 \times 16 \times 4^2$	$8 \times 6 \times 12^3$	3000 s	139 d	8.5 Mch
							16.8 Mch

Table 2: Estimate of time and cost for each run. The reported volume corresponds to the physical volume. In the openQ*D code, C* boundary conditions are implemented by means of an orientifold construction which effectively doubles the simulated volume. For instance the 60×32^3 physical volume is simulated with a $60 \times 64 \times 32^2$ volume which matches the product of the local lattice times the processor grid. The total time and cost are calculated by considering 50 configurations separated by 80 molecular dynamics units (MDUs). In the last column, Mch = Mcore-hours.

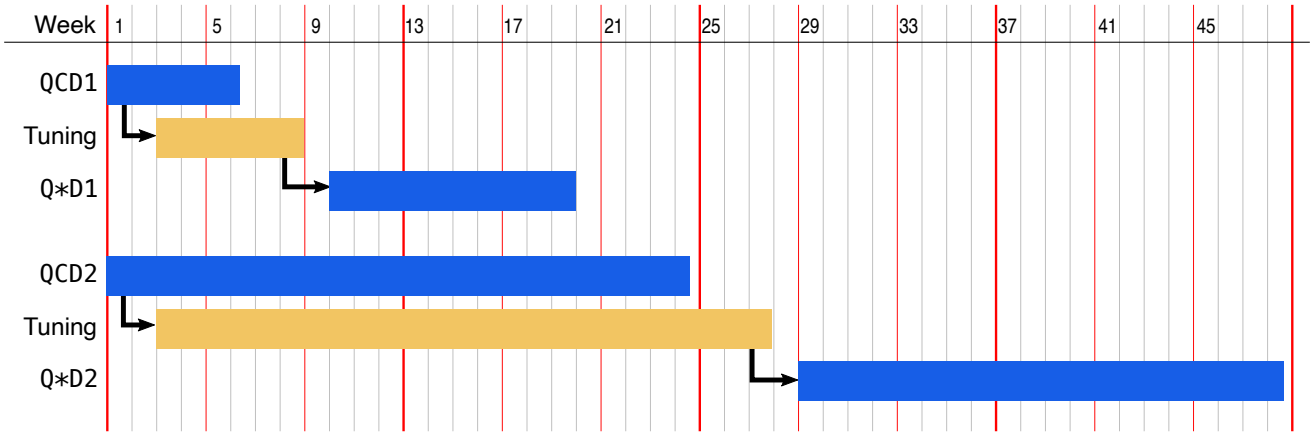

Figure 1: Gantt diagram of the planned workflow, with dependencies. The blue bars correspond to jobs to be run on the HLRN machines, while the yellow bars correspond to jobs to be run on different machine (which are already available). Configurations will be generated on the HLRN machines. The tuning process needs measurement of mesonic correlators in the electroquenched setup on the QCD1,2 configurations. The calculation of the electroquenched observables can start as soon as the first QCD configuration is available. However the corresponding Q*D run can start only after the tuning has terminated.

2 Preparatory work

A. Patella and N. Tantalo have worked on the theoretical foundation of QCD+QED in finite volume with C* boundary conditions [**?** **?**], showing that this provides a sound setup to calculate isospin-breaking and QED radiative corrections to hadronic quantities by means of lattice simulations. M. Hansen, A. Patella and N. Tantalo have demonstrated [**?**] how electrically charged states are interpolated with a class of operators, originally suggested by Dirac and built as functionals of the photon field, that are invariant under local gauge transformations. It turns out that the quality of the numerical signal of charged–hadron masses is the same as in the neutral sector. A preliminary strategy to describe states of charged hadrons with real photons in a fully gauge–invariant way is also discussed, providing a first evidence that the proposed strategy is numerically viable.

I. Campos, P. Fritzsch, M. Hansen, M. Marinkovic, A. Patella, N. Tantalo (as part of the RC* collaboration) are among the developer of the open-source `openQ*D-1.0` code [**?**], which will be used in this project. The features of this code are described in [**?**], together with an analysis of performance and a discussion of some non-trivial tests. Earlier versions of the code have been discussed in [**?** **?**].

P. Fritzsch has worked on the determination of the improvement coefficient $c_{\rm sw}$ in $N_f = 3+1$ QCD [**?** **?**], which is used also in the simulations proposed in this project.

N. Tantalo has worked extensively on isospin-breaking and radiative corrections to hadronic observables in the context of lattice simulations: from the development of the so-called Rome123 method [**?**], to pioneering calculations of radiative corrections to leptonic $\pi^\pm$ and $K^\pm$ decay rates [**?** **?**]. These works are closely related to the proposed project, even though the used methods are quite different.

3 Choice of parameters

The QCD+QED action with four flavours of $O(a)$-improved Wilson fermions depends on 10 parameters: the SU(3) bare coupling constant β, the bare fine-structure constant α, the bare masses m_f with $f = u, d, s, c$, and the improvement coefficients $c^q_{\rm sw,SU(3)}$ and $c^q_{\rm sw,U(1)}$ for $q = 2/3, -1/3$. We describe here how these parameters are chosen. A summary of the parameters of the proposed runs can be found in table 1.

For the proposed simulations, we choose the Lüscher-Weisz gauge action for the SU(3) field with

$$\beta = 3.24 \ . \tag{1}$$

According to [**?**], this corresponds to a lattice spacing $a \simeq 0.054$ fm in pure QCD. This result assumes the physical value of the auxiliary observable t_0 obtained from the CLS $N_f = 2+1$ simulations [**?**], i.e. $(8t_0)^{1/2} = 0.415(4)(2)$ fm. The use of the value of t_0 from $N_f = 2+1$ QCD is well motivated on the basis of the decoupling theorem, as shown in [**?**]. Perhaps a better estimate could be obtained by using the MILC $N_f = 2+1+1$ result for t_0 [**?**], i.e. $(8t_0)^{1/2} = 0.4005^{(23)}_{(-14)}$ fm, which yields a marginally coarser lattice spacing $a \simeq 0.056$ fm. We remind that t_0 is defined by solving the equation

$$t_0^2 \langle {\rm tr}\, G_{\mu\nu} G_{\mu\nu}(0, t_0)\rangle = 0.3 \ , \tag{2}$$

where $G_{\mu\nu}(x,t)$ is some discretization of the SU(3) field tensor calculated in terms of the gauge field at positive flow time t [**?**]. Of course the physical value of t_0 should be obtained by setting the scale with an experimentally measurable dimensionful observable, e.g. the mass of the Ω baryon in a full QCD+QED simulation. However this step will not be done here and will be postponed to a later stage. In this document, we stick to the convention used in [**?**] and we use the $N_f = 2+1$ CLS value $(8t_0)^{1/2} = 0.415$ fm with no error, keeping in mind that this value contains an irreducible $O(\alpha)$ ambiguity.

Isospin-breaking corrections are expected to be of the order of the percent, and hence particularly hard to resolve. As done in [**?** **?** **?** **?**], we want to simulate at unphysical values of the fine-structure

constant α (including $\alpha = 0$) in order to interpolate to the physical value. In the proposed simulations we will choose two values for the bare fine-structure constant, i.e.

$$\alpha = 0 \ , \quad \alpha = 0.5 \ . \tag{3}$$

However, in order to make contact with physics, one needs to look at a renormalized coupling α_R at some fixed renormalization scale, e.g. $\mu = (8t_0)^{-1/2}$. In practice we define

$$\alpha_R = \frac{8\pi}{3} t_0^2 \langle F_{\mu\nu} F_{\mu\nu}(0, t_0) \rangle \ , \tag{4}$$

where $F_{\mu\nu}(x,t)$ is some discretization of the U(1) field tensor calculated in terms of the gauge field at positive flow time t. Up to $O(\alpha_R^2)$ terms, the physical point is defined by $\alpha_R = 1/137$.

As usually done in pure-QCD simulations, we want to define *renormalized trajectories*, in parameter space which may correspond to unphysical hadron masses, typically with a larger gap. Several renormalized trajectories should be simulated in order to approach the physical point (i.e. the point defined by $a = 0$ and the physical values of the hadron masses). Our renormalized trajectories are defined by keeping the following quantities

$$\phi_0 = 8t_0 (M_{K^\pm}^2 - M_{\pi^\pm}^2) \ , \tag{5}$$

$$\phi_1 = 8t_0 (M_{\pi^\pm}^2 + M_{K^\pm}^2 + M_{K^0}^2) \ , \tag{6}$$

$$\phi_2 = 8t_0 (M_{K^0}^2 - M_{K^\pm}^2) \alpha_R^{-1} \ , \tag{7}$$

$$\phi_3 = \sqrt{8t_0} (M_{D_s} + M_{D^0} + M_{D^\pm}) \ , \tag{8}$$

constant, while α and β are varied.

- Notice that $\phi_0 = 0$ if and only if $m_d = m_s$, where the theory is invariant under an SU(2) flavour symmetry which rotates down and strange (which is often called *U-spin symmetry*). We will refer to the $\phi_0 = 0$ as the *U-symmetric point*. At the physical point

$$\phi_0^{\rm phys} \simeq 0.992 \ . \tag{9}$$

- The quantity ϕ_1 has been already used in other contexts [**?** **?**]. In χPT, one easily shows that

$$\phi_1 = A(m_{u,R} + m_{d,R} + m_{s,R}) + B\alpha_R + \text{NLO} + O(\alpha_R^2) \ , \tag{10}$$

$$t_0 = t_0^\chi + C(m_{u,R} + m_{d,R} + m_{s,R}) + D\alpha_R + \text{N}^2\text{LO} + O(\alpha_R^2) \ , \tag{11}$$

  where for t_0 one can use [**?**]. One of the advantages of keeping ϕ_1 fixed is that this eliminates the dependence on the mass of t_0 (and implicitly of the lattice spacing) at leading order χPT with $\alpha = 0$. When electromagnetic interaction are switched on, as long as ϕ_1 is kept constant, one expects a small variation in t_0 of $O(\alpha_R)$. At the physical point

$$\phi_1^{\rm phys} \simeq 2.26 \ . \tag{12}$$

- In χPT, the quantity

$$8t_0 (M_{K^0}^2 - M_{K^\pm}^2) = E(m_{d,R} - m_{u,R}) + F\alpha_R + \text{NLO} + O(\alpha_R^2) \tag{13}$$

  receives two contributions: a term proportional to $m_{d,R} - m_{u,R}$ from strong-isospin effects and a term from QED corrections proportional to α_R. At the physical point, the two effects have the same order of magnitude. We choose to keep this feature along our renormalized trajectories by scaling the above quantity (and therefore $m_{d,R} - m_{u,R}$) proportionally to α_R. This corresponds to the choice of keeping ϕ_2 fixed and equal to its value at the physical point

$$\phi_2^{\rm phys} \simeq 2.37 \ . \tag{14}$$

  Notice that with this choice, at $\alpha = 0$ one has $M_{K^0} = M_{K^\pm}$ which corresponds to the isospin symmetric limit $m_u = m_d$.

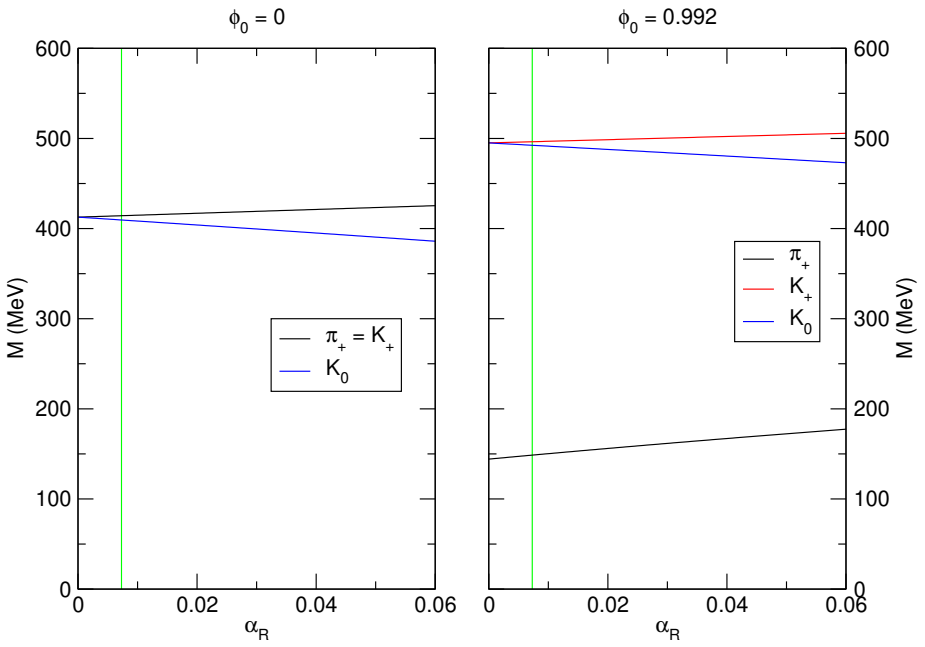

Figure 2: The light mesons masses are plotted as a function of the fine-structure constant for two choices of renormalized trajectories. In both renormalized trajectories we have chosen $\phi_{1,2,3} \simeq \phi_{1,2,3}^{\rm phys}$. The left pane corresponds to the choice $\phi_0 = 0$ (U-symmetric trajectory), and the right pane corresponds to the choice $\phi_0 = \phi_0^{\rm phys}$ (physical trajectory). The green vertical line corresponds to $\alpha_R = 1/137$. Notice that the π_0 is never considered. This is because it decays under electromagnetic interactions.

- The quantity ϕ_3 is used essentially to fix the charm quark mass [**?**]. We will take it equal to its value at the physical point

$$\phi_3^{\rm phys} \simeq 12.0 \ . \tag{15}$$

The aim of this project is to simulate on the unphysical renormalized trajectory defined by

$$\phi_0 = 0 \ , \quad \phi_1 \simeq \phi_1^{\rm phys} \ , \quad \phi_2 \simeq \phi_2^{\rm phys} \ , \quad \phi_3 \simeq \phi_3^{\rm phys} \ . \tag{16}$$

The light stable mesons for this trajectory (together with the physical one) are represented in fig. 2.

For $\alpha = 0$ this corresponds to the QCD SU(3)-symmetric point $m_u = m_d = m_s$. The parameters for this point have been determined in [**?**] (see table 1), and produce $\phi_1 = 2.174(12)$ and $\phi_3 = 12.059(20)$ (notice that there is a trivial mismatch of definitions between this application and [**?**]).

For $\alpha = 0.05$, the bare masses need to be tuned in order to keep the $\phi_{0,1,2,3}$ variables constant. This can be done in the electroquenched setup on the QCD configurations.[2] We have tested in a $N_f = 2$ simulation that, if the bare masses are tuned in this way and then used as parameters in a full QCD+QED simulation with $\alpha = 0.05$, the measured mesons differ by not more than 3% from the electroquenched ones, while at this value of α, isospin effects should be of about 8%. The parameters found in this way surely constitute an very good initial guess for any more accurate tuning. An important goal of this project is to quantify the accuracy of this strategy.

Two volumes are considered, 60×32^3 and 80×48^3, with C* boundary conditions in space and periodic boundary conditions in time. At the proposed lattice spacing, topology is not expected to be frozen, c.f. [**?**], which justifies the use of periodic boundary conditions in time. The two volumes 60×32^3 and 80×48^3 correspond to $M_{\pi^\pm} L \simeq 3.6$ and 5.4 respectively at $\alpha_R = 0$, and $M_{\pi^\pm} L \simeq 3.4$ and

[2]The electroquenched setup corresponds to consider electrically charged valence quarks and electrically neutral see quarks. In practice this means that electroquenched observables are measured on pure QCD configurations.

5.1 respectively for $\alpha_R = 0.05$. This will allow a first study of the finite-volume effect in QCD+QED with C* boundary conditions.

We finally comment on the SW improvement coefficients: $c_{\mathrm{sw,SU(3)}}$ is associated to the operator $\bar{\psi}\sigma_{\mu\nu}G_{\mu\nu}\psi$, and $c_{\mathrm{sw,U(1)}}$ is associated to the operator $\bar{\psi}\sigma_{\mu\nu}F_{\mu\nu}\psi$. Each of these coefficients depends on the electric charge of the quark, for a total of four improvement coefficients in a QCD+QED simulation. We take the two SU(3) coefficients equal to the non-perturbatively determined ones in pure-QCD [**?**], i.e.

$$c^{q=2/3}_{\mathrm{sw,SU(3)}} = c^{q=-1/3}_{\mathrm{sw,SU(3)}} = 2.18859 \tag{17}$$

and the two U(1) coefficients equal to the tree-level ones

$$c^{q=2/3}_{\mathrm{sw,U(1)}} = c^{q=-1/3}_{\mathrm{sw,U(1)}} = 1 \ . \tag{18}$$

In practice this means that $O(a)$ corrections are not entirely eliminated, however they are suppressed with one power of α.

4 Code and algorithms

The production of gauge field ensembles and measurement of physical observables is based on the new openQ*D code base, recently made available under the GNU General Public License []. The programs are highly optimized for machines with current x86-64 processors, but will run correctly on any system that complies with the ISO C89 and the MPI 1.2 standards. The code is structured to ensure a very good data locality. Nevertheless, the performance of the programs is mainly limited by data movement, i.e., the memory-to-processor bandwidth and network latency.

The simulation program implements an (rational-)HMC algorithm which evolves the physical fields in phase space. Each trajectory starts with momentum fields randomly chosen from a normal distribution. Then the fields are evolved according to the molecular dynamics (MD) of the Hamiltonian equations. The equations are integrated for a fixed molecular dynamics time (trajectory length), using nested hierarchical symplectic integrators such as the 4th-order Omelyan–Mryglod–Folk (OMF) integrator []. At the end of each trajectory, the fields are submitted to an accept-reject step that corrects for the integration errors.

The code has a highly-optimised lattice Dirac operator (e.g. even-odd preconditioning) and implements frequency-splitting for the quark determinant. The use of the rational approximation and twisted masses requires standard reweighting techniques which are supported by calculating the corresponding factors a posteriori. During the MD evolution, the Dirac operator has to be inverted multiple times which is accelerated using modern techniques like deflation, multi-shift and chronological solvers. The choice of solvers (CGNE, MSCG, SAP+GCR, DFL+SAP+GCR) is separately configurable for each force component and pseudo-fermion action. Additional features are an implementation of the Fast Fourier Transform as well as Fourier acceleration for the electromagnetic gauge fields. A corresponding acceleration technique for the strong gauge field does not exist.

5 Code performance

The `openQ*D` code is an extension of the `openQCD` code [**?**] which has been extensively used to generate QCD configurations e.g. by the ALPHA collaboration, c.f. [**?**] and by the CLS network, c.f. [**?**].

The inversion of the Dirac operator constitutes the bulk of the calculation in the proposed runs. In particular the Dirac operator for light quarks is inverted with a deflated SAP-preconditioned GCR solver. In order to illustrate the scalability of the code, we have studied the strong scaling of the SAP preconditioner on the HLRN-IV system. Results are reported in the left pane of fig. 3. The code shows

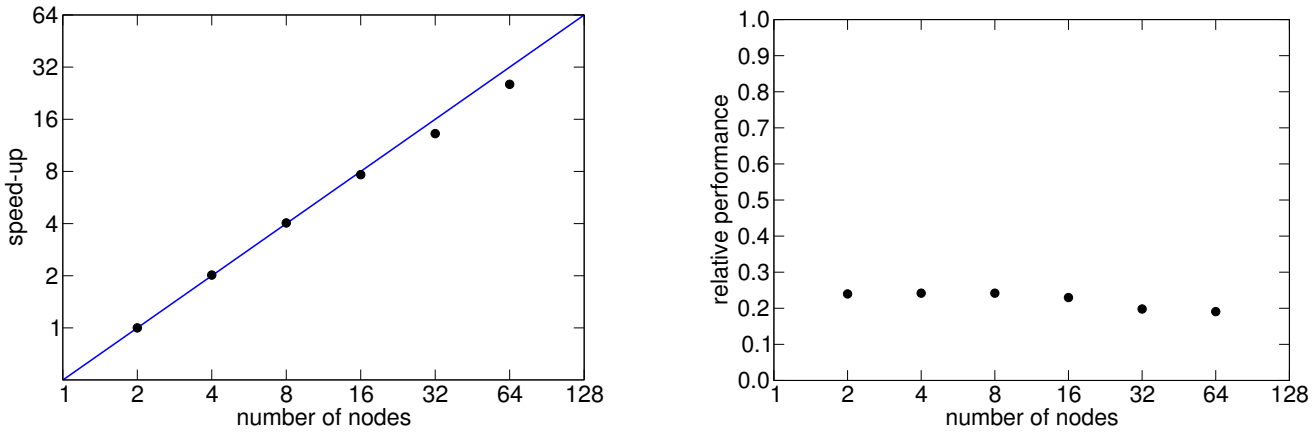

Figure 3: Strong-scaling and performance analysis of the SAP preconditioner. For this study we have chosen a QCD+QED setup, a $160 \times 64 \times 32 \times 32$ lattice, and we have varied the number of nodes. In the left pane, we plot the speed-up as a function of the number of nodes n. The speed-up is defined as the ratio of the running time on 2 nodes divided by the running time on n nodes. In the right pane, we plot the relative performance as a function of the number of nodes n. The relative performance is defined as the achieved Flops per processor divided by the theoretical peak performance.

almost perfect scaling up to 16 nodes, and a slight degradation for 32 and 64 nodes. In order to be able to conclude this project within a year, we choose to generate the `QCD1`, `QCD2`, `Q*D1` ensembles with 32 nodes, and the `Q*D2` ensemble with 64 nodes (compare with table 2 and figure 1). In principle we could have chosen to generate the `QCD1` and `Q*D1` with 16 nodes. However our choice makes the use of resources more uniform over the 4 quarters, and allows us to have the `QCD1` and `Q*D1` ready for presentation in the yearly Lattice conference.

The base clock frequency of the Intel Skylake Gold 6148 processor with AVX-512 instruction is 1.6 GHz, see e.g. `https://en.wikichip.org/wiki/intel/xeon_gold/6148`, yielding a theoretical peak performance of 51.2 GFlops/core (in double precision).[3] The SAP preconditioner has achieved a performance corresponding to 24% and 19% of the theoretical peak performance, in the cases of 2 and 64 nodes respectively, see right pane of fig. 3.

6 Justification of resources

Dedicated QCD and QCD+QED simulation with $N_f = 2$, pion masses of about 400 MeV, a 64×32^3 lattice with C* boundary conditions in the spatial directions, and the RHMC algorithm have been performed to estimate the cost of this project. The cost for generating a thermalized configuration has been measured on the HLRN-IV system in Goettingen, and appropriately rescaled to provide the cost estimates in table 2.

We plan to generate 50 independent configurations per run, and we have used the estimate of the autocorrelation time of 80 molecular-dynamics units (MDU) provided in [**?**]. Notice that [**?**] uses the HMC for the light quarks, while we need to use the RHMC (even in QCD runs, because of the C* boundary conditions). Since the action is different, the autocorrelation time may turn out to be different as well. Autocorrelations may also turn out to be different between QCD and QCD+QED runs. However it is worth noticing that the use of Fourier acceleration for the U(1) field is expected to reduce the effect on the autocorrelations due to the U(1) field. An analysis of the autocorrelations is integral part of this project.

[3]Notice that the nominal base frequency 2.4 GHz can be achieved only without AVX-512 instructions, yielding a lower theoretical performance of 38.4 GFlops/core (in double precision). We also assume in this estimate that the turbo boost is not used.

The rescaling factor applied to the $N_f = 2$ sample simulations has been calculated by assuming linear dependence with the volume, and by rescaling with the number of pseudofermion actions needed to generate the configurations (the cost of the charm quark has been estimated to be about half the cost of a light quark). Notice that the proposed simulations are significantly more expensive than QCD simulations in the isospin limit and periodic boundary conditions, because of the use of RHMC for the light quarks. The `QCD1` and `QCD2` runs have $m_u = m_d = m_s$ which means that a single pseudofermion action can be chosen to simulate the three light quarks. Since the up-type and down-type quarks have different electric charges, two pseudofermion actions must be used to simulate the three light quarks in the `Q*D1` and `Q*D2` runs, which yields a twofold increase in the computational cost of the light-quark sector. As observed in [**?**], the time needed to invert the Dirac operator is not affected by the fact that QED is turned on, provided that the pion/Kaon mass is constant. The `Q*D1` and `Q*D2` runs have a slightly lighter K^0 with respect to the `QCD1` and `QCD2` runs because of isospin-splitting effects, however we have checked explicitly that the time needed for the deflated solver implemented in `openQ*D` to solve the Dirac equation is essentially independent of the pion/kaon mass in the range that is relevant for the proposed simulations. These observations allow to scale the measured times from our $N_f = 2$ QCD and QCD+QED sample simulations and provide an estimate for the time needed to generate a single MD trajectory in the proposed runs on the HLRN-IV system.

This project requires very limited amount of storage. We will need to store 50 gauge configurations for each run described in table 2, for a total of 597 Gb, which will be written in the WORK filesystem. These configurations will be moved to available storage space at the Institut für Physik, Humboldt-Universität zu Berlin, via ssh, at the end of the project. The configurations will be available to the members of the RC* collaboration only for a period of 2 years, and they will be made publicly available after 2 years.