

Deflation acceleration of lattice QCD simulations

Martin Lüscher

*CERN, Physics Department, TH Division,
CH-1211 Geneva 23, Switzerland
E-mail: luscher@mail.cern.ch*

ABSTRACT: Close to the chiral limit, many calculations in numerical lattice QCD can potentially be accelerated using low-mode deflation techniques. In this paper it is shown that the recently introduced domain-decomposed deflation subspaces can be propagated along the field trajectories generated by the Hybrid Monte Carlo (HMC) algorithm with a modest effort. The quark forces that drive the simulation may then be computed using a deflation-accelerated solver for the lattice Dirac equation. As a consequence, the computer time required for the simulations is significantly reduced and an improved scaling behaviour of the simulation algorithm with respect to the quark mass is achieved.

KEYWORDS: Lattice Gauge Field Theories, Lattice QCD.

Contents

1. Introduction	1
2. Preliminaries	2
2.1 Lattice formulation	2
2.2 DD-HMC algorithm	3
2.3 Domain-decomposed deflation subspaces	3
3. Acceleration of the DD-HMC algorithm	4
3.1 Subspace update procedure	4
3.2 Chronological inversion method	5
3.3 Reversibility of the molecular-dynamics evolution	6
4. Tests of the accelerated algorithm	6
4.1 Lattice parameters and field ensembles	6
4.2 HMC parameters	7
4.3 Choice of the solver tolerances	7
4.4 Deflation-subspace and other parameters	8
4.5 Test results	8
5. Concluding remarks	9
A. Subspace generation and update procedures	10
A.1 Field initialization and first steps	10
A.2 Subspace refinement	10
A.3 Subspace updates	11
B. Automatising the subspace updates	11

1. Introduction

Numerical simulations of lattice QCD have become a powerful tool for quantitative studies of the properties of the strongly interacting particles. The systematic errors in these calculations are not easy to assess, however, and most results published to date must for this reason be regarded as preliminary. In order to remove this deficit, simulations in a wide range of lattice spacings, lattice volumes and sea-quark masses will have to be performed, thus requiring larger and larger lattices to be considered.

The use of advanced simulation techniques that scale well with the quark masses and the lattice parameters is likely to be crucial for the success of this programme. In the

last few years, a lot of progress has already been made in this direction. In particular, preconditioned [1]–[7] and other variants [8] of the Hybrid Monte Carlo (HMC) algorithm [9] were introduced which allowed simulations to be performed at much smaller quark masses than was possible before (see refs. [10]–[15], for example).

Important advances have also been made in the area of variance-reduction methods and quark propagator calculations. The progress here is often based on deflation ideas, where the low modes of the lattice Dirac operator are treated separately from the bulk of the modes [16]–[24]. Usually the mode separation is achieved by calculating the low-lying eigenvalues and the associated eigenvectors of the Dirac operator. When implemented in this way, low-mode deflation however tends to become inefficient or even impractical on large lattices, because the number of eigenvectors that must be computed grows proportionally to the lattice volume.

This limitation can be overcome through the use of domain-decomposed deflation subspaces and projection techniques that do not require the deflation subspace to be spanned by eigenvectors of the Dirac operator [24]. Domain-decomposed deflation subspaces can be generated with a modest computational effort and were found to be highly effective. The numerical solution of the lattice Dirac equation, for example, can be accelerated by an impressive factor using such deflation subspaces together with a suitable preconditioner [24].

In the present paper it will be shown that domain-decomposed deflation subspaces can be easily propagated along the molecular-dynamics trajectories generated by the HMC algorithm. At each step of the molecular-dynamics evolution, the subspace may then be used to accelerate the computation of the quark forces that drive the simulation. The simulation algorithm itself thus remains the same, but the required computer time is reduced significantly, particularly so at small quark masses. Moreover, a further acceleration can be achieved through the simultaneous use of the “chronological inversion method” of Brower et al. [25].

2. Preliminaries

The deflation acceleration of the HMC algorithm will be worked out for a definite choice of the lattice formulation of the theory and of the preconditioned form of the algorithm. However, low-mode deflation and the subspace propagation do not depend on the particular choices made and are expected to be widely applicable.

2.1 Lattice formulation

The standard $O(a)$ -improved Wilson theory [26, 27] will be considered, with a doublet of mass-degenerate sea quarks and non-perturbatively determined coefficient c_{sw} of the Sheikholeslami-Wohlert improvement term [28]. In this theory, the adjustable parameters are the lattice sizes T and L in the time and the space directions, the inverse bare coupling β and the bare sea-quark mass in units of the lattice spacing (or, equivalently, the hopping parameter κ_{sea}).

The utility of any particular deflation method in general depends on the physical situation one is interested in. Here it will be assumed that the theory is considered in

the large-volume regime of QCD, where, say, $T \geq L \geq 2 \text{ fm}$ and $M_\pi L \geq 3$, M_π being the mass of the “pion” at the specified sea-quark mass. Moreover, in order to guarantee the stability of the HMC simulations, the lattice parameters must be such that the low end of the spectrum of the Dirac operator is safely separated from the origin [29].

2.2 DD-HMC algorithm

The HMC algorithm is nowadays practically only used in preconditioned form, where the quark determinant is split into several factors before the pseudo-fermion fields are introduced [1]–[8]. The associated forces must then be computed at regular intervals along the generated trajectories in field space, the larger forces more often than the smaller ones, according to a scheme introduced by Sexton and Weingarten [30].

In the case of the DD-HMC algorithm (which is the algorithm chosen here),¹ the lattice is decomposed into non-overlapping blocks of lattice points and the quark determinant is factorised into the product of the determinants of the block Dirac operators times another factor that couples the gauge fields on the different blocks. There are thus two kinds of quark forces, the block forces and the block-interaction force, the latter being the one that includes the contributions of the low modes of the Dirac operator.

The computation of the block-interaction force requires the Dirac equation on the full lattice to be solved for two source fields. This calculation usually consumes the dominant fraction of the computer time spent for the simulation. The equation will here be solved using the Schwarz-preconditioned GCR algorithm [5], or its deflated version [24] once the deflation subspace is available along the molecular-dynamics trajectories.

2.3 Domain-decomposed deflation subspaces

For the reader’s convenience and in order to set up the notation, the definition of domain-decomposed deflation subspaces is recalled in the following paragraphs and it is explained how they are obtained in practice [24].

The starting point is again a regular decomposition of the lattice into non-overlapping blocks Λ of fixed size (equal to 4^4 or $6^2 \times 4^2$, for example). Similar block decompositions of the lattice are also used for the preconditioning of the HMC algorithm and the GCR solver, but all these block grids are logically unrelated and it is better to think of them as being separate structures even if the block sizes happen to be the same.

Once a particular block decomposition is selected, a set $\psi_l(x)$, $l = 1, \dots, N_s$, of quark fields is generated through the smoothing procedure outlined below. The fields are then projected to the blocks Λ through

$$\psi_l^\Lambda(x) = \begin{cases} \psi_l(x) & \text{if } x \in \Lambda, \\ 0 & \text{otherwise,} \end{cases} \quad (2.1)$$

and the linear space spanned by the set of all block fields $\psi_l^\Lambda(x)$ is taken to be the deflation subspace. The latter thus has dimension equal to N_s times the number of blocks. Typical

¹The acronym DD-HMC stands for “Domain Decomposition Hybrid Monte Carlo”, which is the now commonly adopted name of the Schwarz-preconditioned HMC algorithm that was introduced in ref. [6].

values of N_s range from 12 to 24, but as explained in ref. [24], some tuning is normally required to find the best values of N_s .

The smoothing procedure starts from some random fields $\psi_l(x)$, $l = 1, \dots, N_s$, and consists in applying a number of approximate inverse iteration steps

$$\psi_l(x) \rightarrow "D^{-1}"\psi_l(x) \tag{2.2}$$

to them, where D denotes the lattice Dirac operator at a value of the bare quark mass close to the critical mass (the inverse of D is put in quotes in this formula in order to make it clear that an approximation to the inverse is being applied). As a result the components of the fields along the high modes of the Dirac operator are suppressed and the fields will therefore have a strong overlap with the subspace spanned by the low modes of the operator. The local coherence of the latter then guarantees that the associated domain-decomposed deflation subspace is highly effective [24].

Exactly which algorithm is used for the approximate solution of the Dirac equation in the recursion (2.2) should not matter too much. The procedure employed here is described in appendix A.

3. Acceleration of the DD-HMC algorithm

In the following, the basic strategy will be to generate a domain-decomposed deflation subspace at the beginning of each molecular-dynamics trajectory and to preserve its deflation efficiency along the trajectory by applying a suitable update procedure. The subspace is then used to speed up the computation of the block-interaction force and thus to accelerate the simulation. This section also includes a short description of the chronological inversion method of Brower et al. [25], which is recommended to be used together with the deflation acceleration.

3.1 Subspace update procedure

The effectiveness of any particular update procedure can be determined by observing the number N_{GCR} of deflated solver iterations that are required for the computation of the block-interaction force. The typical behaviour of the iteration number along a molecular-dynamics trajectory is shown in figure 1. One actually observes two iteration numbers in such an experiment, represented by the lower and upper series of points in the figure, since the Dirac equation must be solved two times when the force is calculated.

In the example shown here, each update of the deflation subspace consists in applying a single deflated inverse iteration step to the quark fields $\psi_l(x)$, $l = 1, \dots, N_s$, that define the subspace via the block projection (2.1) (deflated inverse iteration is explained in detail in appendix A). After every subspace update, the solver iteration numbers drop to a lower value and then slowly increase again as the gauge field evolves along the trajectory. The behaviour is different in the initial period of the molecular-dynamics integration for a reason explained in the next subsection.

As it turns out, this simple procedure works well at all values of the lattice spacing and the sea-quark mass considered so far. The optimal number of subspace updates varies,

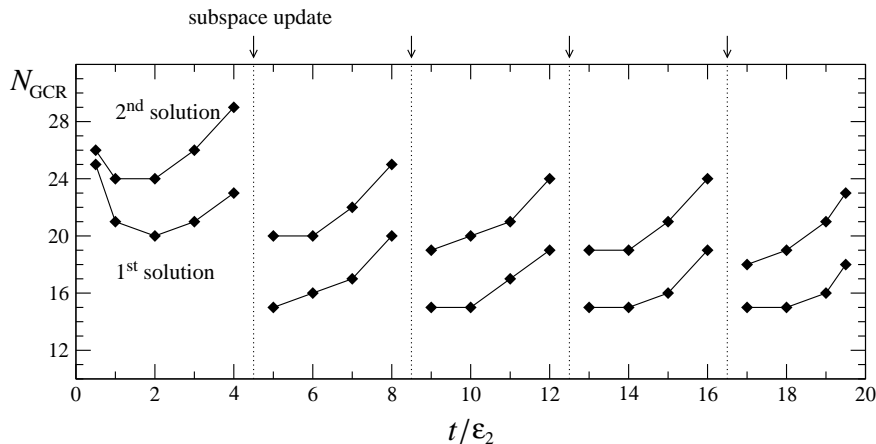


Figure 1: Typical history of the iteration numbers N_{GCR} (data points) of the deflated Schwarz-preconditioned GCR solver along a molecular-dynamics trajectory on a 64×32^3 lattice at $\kappa_{\text{sea}} = 0.13625$, plotted against the molecular-dynamics time t given in units of the integration step size ε_2 (the lattice and algorithmic parameters are fully specified in section 4).

but in order to maintain a high deflation efficiency of the subspace it proves to be entirely sufficient to update the fields $\psi_l(x)$ through the application of a single deflated inverse iteration step.

The frequency of the subspace updates along the molecular-dynamics trajectories could be taken to be a static parameter whose value is chosen based on experience. However, it is more elegant and probably also more efficient to let the simulation algorithm choose when the subspace is to be updated. Such an automatic update procedure is described in appendix B.

3.2 Chronological inversion method

The solutions of the Dirac equation calculated in the course of a molecular-dynamics trajectory evolve smoothly with time and one can try to forecast the solution at the next integration step from the previous solutions. The forecast will normally not be as accurate as required, but the solver can start from the proposed solution and will obtain the solution to the specified precision faster than when it starts from zero.

The number of previous solutions that should be kept in memory is an adjustable parameter whose optimal value depends on the lattice and algorithmic parameters as well as on which Dirac equation is being considered. In the case of the DD-HMC algorithm, there are altogether three equations, the normal even-odd preconditioned equation on the blocks which must be solved when calculating the block forces and the two equations on the full lattice which must be solved when calculating the block-interaction force. The corresponding numbers of saved solutions will be denoted by p_1 , p_2 and \tilde{p}_2 , respectively. On the blocks the extrapolation method is taken to be the “minimal residual extrapolation” recommended by Brower et al. [25], while for the solutions on the full lattice a polynomial extrapolation is used.

In the case of the block interaction force, the combination of the deflation acceleration and the chronological inversion method leads to the peculiar behaviour seen in the initial period of the iteration number history shown in figure 1. What happens there is that the solution forecast rapidly improves in the first few steps and more than compensates the loss of efficiency of the deflation subspace. The accuracy of the solution forecast then saturates and the solver iteration numbers start to grow until the subspace is updated for the first time.

3.3 Reversibility of the molecular-dynamics evolution

In order to guarantee the correctness of the HMC algorithm, the approximate integration of the molecular-dynamics equations must be reversible. The multiple-time integration scheme used here fulfils this requirement, although in practice the reversibility can be compromised by rounding errors and the chosen solver tolerances.

When the deflation acceleration is switched on, the situation is complicated by the fact that the propagation of the deflation subspace is not reversible. The propagation of the solutions of the Dirac equation is also not reversible, but as both methods only serve to speed up the computation of the quark forces, the reversibility violations caused by them are proportional to the solver tolerances [25]. These must therefore be chosen so that the reversibility is guaranteed to high precision (see subsection 4.3).

4. Tests of the accelerated algorithm

In this section the DD-HMC and the accelerated DD-HMC algorithm are submitted to a speed test. The HMC parameters are set to the same values in the two cases so that any observed speed-up factors can be unambiguously attributed to the deflation acceleration and the chronological inversion method.²

4.1 Lattice parameters and field ensembles

All tests reported below were performed on a 64×32^3 lattice at inverse bare coupling $\beta = 5.3$, coefficient $c_{\text{sw}} = 1.90952$ [28] of the Sheikholeslami-Wohlert improvement term and four values of the sea-quark hopping parameter κ_{sea} (see table 1). At these points in parameter space, the lattice spacing in physical units is estimated to be 0.0784(10) fm, while the “pion” mass M_π ranges from about 618 to 282 MeV as κ_{sea} increases from 0.13590 to 0.13635 [10].

In order to obtain average timings with statistical errors of at most a few percent, it suffices to perform a series of short simulations, starting from a set of statistically independent representative gauge-field configurations. Ensembles of such field configurations were generated on the specified lattice as part of another project by the authors of ref. [10] and were made available for the tests conducted here.

²An updated version of the DD-HMC code is available under the terms of the GNU Public License (GPL) at <http://cern.ch/luscher/DD-HMC>.

κ_{sea}	τ	N_0	N_1	N_2	p_1	p_2	\tilde{p}_2
0.13590	0.50	4	5	11	8	6	5
0.13610	0.50	4	5	14	8	6	5
0.13625	0.50	4	5	20	8	6	5
0.13635	0.25	4	5	16	8	4	3

Table 1: Parameter values used in the test runs

4.2 HMC parameters

Following previous work [10], the DD-HMC algorithm is set up on a division of the lattice into blocks of size 8^4 . The length τ of the molecular-dynamics trajectories and the integration step numbers N_0 , N_1 and N_2 are then set to the values quoted in table 1. With these choices, trajectory acceptance rates of 80–85% are achieved. Recall that the step numbers refer to the gauge, block and block-interaction forces, respectively [6]. In particular, the latter is evaluated at molecular-dynamics times separated by steps of size $\varepsilon_2 = \tau/N_2$.

The simulations performed at the lightest quark mass considered ($\kappa_{\text{sea}} = 0.13635$) are close to the edge of the stability range of the HMC algorithm [29]. Following a suggestion of the PACS-CS collaboration [14, 15], the trajectory length has been set to half the usual value in this case in order to reduce the rate of trajectories with large energy deficits at the end of the molecular-dynamics evolution. The experience so far is that this choice does not lead to larger autocorrelation times (if measured in units of molecular-dynamics time) so that the efficiency of the simulation remains practically the same.

4.3 Choice of the solver tolerances

The computation of the pseudo-fermion action and the fermion forces requires the Dirac equation $D\psi = \eta$ (or the associated normal equation) to be solved numerically. In each case the algorithm used for this task is stopped as soon as the calculated solution satisfies $\|\eta - D\psi\| \leq \omega\|\eta\|$ for a specified tolerance ω .

For lattices of size 32×24^3 , the tolerances recommended in ref. [5] were 10^{-8} for the block forces, 10^{-7} for the block-interaction force and 10^{-11} and 10^{-10} for the corresponding pseudo-fermion actions. The same tolerances were also used in the simulations reported in ref. [10], some of which were performed on 64×32^3 lattices.

However, as already mentioned, the deflation acceleration and the chronological inversion method tend to compromise the reversibility of the numerical integration of the molecular-dynamics equations. Extensive reversibility tests on the 64×32^3 lattice at $\kappa_{\text{sea}} = 0.13625$ actually show that the tolerances should better be set to 10^{-10} for the forces and 10^{-11} for the pseudo-fermion actions if the accelerations are switched on. After a return trajectory, the absolute value of the energy deficit is then always less than 10^{-5} and the components of the initial and final gauge-field configurations deviate by at most 10^{-9} .

Similar reversibility violations were observed without acceleration and the previously recommended values of the tolerances. In the speed tests, the tolerances were therefore

κ_{sea}	DD-HMC			Accelerated DD-HMC		
	N_{CG}	N_{GCR}	t [min]	N_{CG}	N_{GCR}	t [min]
0.13590	124	41	22	59	17	17
0.13610	126	60	36	53	18	22
0.13625	127	96	75	48	20	32
0.13635	128	158	93	40	22	28

Table 2: Average solver iteration numbers N_X and execution times t per trajectory.

taken to be the old ones in the case of the ordinary DD-HMC simulations and the smaller ones specified above in the case of the accelerated simulations.

4.4 Deflation-subspace and other parameters

At all values of the sea-quark mass, the deflation subspace was generated by applying 9 inverse iteration steps to $N_s = 20$ random quark fields and by projecting them to a division of the lattice into blocks of size 4^4 . This choice of parameters practically coincides with the one suggested in ref. [24]. Some experimenting was however required in order to find the best values of the numbers p_1 , p_2 and \tilde{p}_2 of old fields that are to be used for the chronological propagation of the solutions of the Dirac equation (see table 1).

The parameters of the Schwarz preconditioned GCR solver were set to the values previously recommended in refs. [5, 24]. In particular, the size of the blocks on which the Schwarz preconditioner operates was taken to be 8×4^3 in all cases.

4.5 Test results

The figures listed in table 2 show that the average solver iteration numbers N_{GCR} required for the computation of the block-interaction force are significantly reduced when the deflation acceleration and the chronological inversion method are switched on. Most of the reduction is achieved through the deflation of the Dirac equation, the additional reduction through the solution forecast being at the level of 10–20%.

Also shown in the table are the average conjugate-gradient iteration numbers N_{CG} needed to compute the block forces. Here the observed reduction in the iteration numbers is a consequence of the solution forecast alone. Note that a reduction by a factor 2-3 is achieved even though the solver tolerance had to be lowered in order to preserve the reversibility of the integration of the molecular-dynamics equations.

The average execution times t per trajectory quoted in table 2 were measured on a recent PC cluster with 32 (single-core) dual-processor nodes connected through a switched Infiniband network. While these figures depend on program and hardware details, they clearly show, in a realistic case, that the algorithmic accelerations do result in important speed-up factors. They also lead to a softer scaling of the timings with respect to the sea-quark mass (see figure 2), as was to be expected given the flat scaling behaviour of the deflated Schwarz-preconditioned GCR solver [24].

The deflation acceleration tends to reduce the computer time needed for the calculation of the block-interaction forces to a level where the time spent in other parts of the program




Figure 2: Plot of the average execution times t (data points) required per molecular-dynamics trajectory of length $\tau = 0.5$ as a function of the inverse of the bare current-quark mass m_{sea} given in lattice units. At the smallest quark mass considered, the time needed for two trajectories of length $\tau = 0.25$ is plotted. The lattice and algorithmic parameters are specified in subsections 4.1–4.4 and timings were taken on a PC cluster with 64 processing units.

is not completely negligible anymore. In the test runs, the generation and propagation of the deflation subspace consumed some 4–5 min per trajectory, while the bulk of the time was divided roughly like 2:1 among the subprogram that computes the block-interaction force and the remaining subprograms.

5. Concluding remarks

Domain-decomposed deflation subspaces are technically attractive for many reasons. One of them certainly is the fact that high deflation efficiencies can be achieved on large lattices using fairly low numbers of modes per domain. For the acceleration of the HMC algorithm, another very important property is that these subspaces can be obtained with a modest computational effort. Moreover, little extra work is required to maintain their efficiency along the molecular-dynamics trajectories generated by the HMC algorithm.

The speed and excellent scaling behaviour of the accelerated DD-HMC algorithm are encouraging and make simulations of lattice QCD with light Wilson quarks more feasible than ever before. Deflation acceleration is, however, not limited to the DD-HMC algorithm nor does one have to follow the lines of this paper in all detail. In particular, instead of the Schwarz alternating procedure, other relaxation algorithms can conceivably be used, both as preconditioner for the deflated GCR solver and for the generation of the domain-decomposed deflation subspaces.

Acknowledgments

I am indebted to Peter Weisz for his critical comments on a first draft of the paper. Thanks also go to Björn Leder and Rainer Sommer for contributing machine-specific accelerations for IBM Blue Gene/L computers to the generic DD-HMC code and to Björn and Carlos Pena for their help in validating the new version of the program. The gauge-field configurations used for the numerical studies were generated by the authors of ref. [10]. All computations reported here were performed on a dedicated PC cluster at CERN.

A. Subspace generation and update procedures

As explained in the main text, the domain-decomposed deflation subspaces are generated and updated using approximate inverse iteration. In this appendix, the exact procedures that were used in the test runs are described in some greater detail.

A.1 Field initialization and first steps

The inverse iteration (2.2) operates on the quark fields $\psi_l(x)$, $l = 1, \dots, N_s$, that define the domain-decomposed deflation subspace via the block projection (2.1). At the beginning of each molecular-dynamics trajectory, the components of these fields are initialized to uniformly distributed random values in the range $[-1, 1]$. The fields are then updated three times according to

$$\psi_l(x) \rightarrow M_{\text{sap}}\psi_l(x), \tag{A.1}$$

where M_{sap} denotes the multiplicative Schwarz preconditioner that was introduced in ref. [5]. M_{sap} depends on several adjustable parameters, but their choice is not critical and good results are obtained using similar parameter values as in the case of the Schwarz-preconditioned GCR solver. However, as already mentioned in section 2, the bare quark mass should be set to a value close to (or even equal to) the critical mass in this calculation, independently of the value of the sea-quark mass.

A.2 Subspace refinement

After the initial phase of the subspace generation, the fields $\psi_l(x)$ are updated several more times using a deflated variant of approximate inverse iteration. A deflated inverse iteration step begins by constructing the domain-decomposed deflation subspace from the current set of fields through the block projection (2.1). This space is left unchanged until all fields are updated once according to the rule

$$\psi_l(x) \rightarrow M_{\text{sap}}\{\psi_l(x) - D\zeta_l(x)\} + \zeta_l(x), \tag{A.2}$$

where $\zeta_l(x) = P\tilde{\zeta}_l(x)$ is an approximate solution of the equation

$$PDP\tilde{\zeta}_l(x) = \psi_l(x) \tag{A.3}$$

in which P denotes the orthonormal projector to the deflation subspace.

Note that the fields $\psi_l(x)$ are contained in the deflation subspace by construction. Equation (A.3) is thus a well-defined, square linear system that actually coincides with the “little Dirac equation” of ref. [24]. In appendix A of that paper, a preconditioned GCR solver for the little equation is described which is also used here. Since inverse iteration is anyway approximately implemented, there is no point in solving the equation very accurately. In the test runs reported in this paper, for example, the solver tolerance was set to 10^{-3} .

The total number of inverse iteration steps that must be applied to generate an effective deflation subspace depends on the lattice parameters, the chosen block size and the number N_s of deflation modes per block. In the cases considered so far, good deflation subspaces were obtained after 9-11 steps (3 ordinary plus 6-8 deflated inverse iteration steps).

A.3 Subspace updates

As explained in section 3, the deflation subspace is updated along the molecular-dynamics trajectories by applying a deflated inverse iteration step from time to time. The procedure is the same as the one described above, i.e. the quark fields $\psi_l(x)$, $l = 1, \dots, N_s$, that define the current subspace are updated according to eq. (A.2).

However, it is recommended to orthonormalize the fields before they are updated, as otherwise it can happen that they become more and more aligned to each other as one moves along the trajectory. After many steps, inverse iteration actually projects any field to the few very lowest modes of the Dirac operator, and although the gauge field changes along a trajectory, the subspace updates can have a similar effect.

B. Automatising the subspace updates

The basic idea of the method proposed here is to observe the solver iteration numbers along the trajectories (as in figure 1) and to update the subspace when the numbers have grown beyond a certain level since the last update.

Let $n(t)$ be the sum of the GCR solver iteration numbers at molecular-dynamics time t and suppose that the last subspace update was at time t_0 . In the following steps along the trajectory, $n(t)$ will normally increase monotonically up to some time t_1 , where the sum of the iteration number differences $n(t) - n(t_0)$ satisfies the bound

$$\sum_{t=t_0}^{t_1} \{n(t) - n(t_0)\} \geq \frac{1}{2} N_s. \tag{B.1}$$

When this point is reached, the subspace is updated at the next step of the molecular-dynamics integration and the procedure then repeats itself to the end of the trajectory.

The inequality (B.1) balances the effort required for the update of the subspace against the additional work that is required for the solution of the Dirac equation with respect to what it would be if the deflation subspace were always in good condition. However, when the deflation acceleration is combined with the chronological inversion method, the left-hand side of the inequality should be modified so as to take into account the fact that

the iteration numbers may not increase monotonically between the subspace updates. A simple possibility is to replace the differences $n(t) - n(t_0)$ in eq. (B.1) by

$$n(t) - \min_{t_0 \leq s \leq t} n(s) + \begin{cases} \frac{1}{16}n(t) & \text{if } t < \max\{p_2, \tilde{p}_2\}\varepsilon_2, \\ 0 & \text{otherwise,} \end{cases} \quad (\text{B.2})$$

where p_2 and \tilde{p}_2 are the numbers of old fields used for the propagation of the solutions of the first and the second Dirac equation on the full lattice (cf. subsection 3.2).

References

- [1] M. Hasenbusch, *Speeding up the Hybrid-Monte-Carlo algorithm for dynamical fermions*, *Phys. Lett.* **B 519** (2001) 177 [[hep-lat/0107019](#)].
- [2] M. Hasenbusch and K. Jansen, *Speeding up lattice QCD simulations with clover-improved Wilson fermions*, *Nucl. Phys.* **B 659** (2003) 299 [[hep-lat/0211042](#)].
- [3] ALPHA collaboration, M. Della Morte et al., *Simulating the Schrödinger functional with two pseudo-fermions*, *Comput. Phys. Commun.* **156** (2003) 62 [[hep-lat/0307008](#)].
- [4] M. Lüscher, *Lattice QCD and the Schwarz alternating procedure*, *JHEP* **05** (2003) 052 [[hep-lat/0304007](#)].
- [5] M. Lüscher, *Solution of the Dirac equation in lattice QCD using a domain decomposition method*, *Comput. Phys. Commun.* **156** (2004) 209 [[hep-lat/0310048](#)].
- [6] M. Lüscher, *Schwarz-preconditioned HMC algorithm for two-flavour lattice QCD*, *Comput. Phys. Commun.* **165** (2005) 199 [[hep-lat/0409106](#)].
- [7] C. Urbach, K. Jansen, A. Shindler and U. Wenger, *HMC algorithm with multiple time scale integration and mass preconditioning*, *Comput. Phys. Commun.* **174** (2006) 87 [[hep-lat/0506011](#)].
- [8] M.A. Clark and A.D. Kennedy, *Accelerating dynamical fermion computations using the Rational Hybrid Monte Carlo (RHMC) algorithm with multiple pseudofermion fields*, *Phys. Rev. Lett.* **98** (2007) 051601 [[hep-lat/0608015](#)].
- [9] S. Duane, A.D. Kennedy, B.J. Pendleton and D. Roweth, *Hybrid Monte Carlo*, *Phys. Lett.* **B 195** (1987) 216.
- [10] L. Del Debbio, L. Giusti, M. Lüscher, R. Petronzio and N. Tantalo, *QCD with light Wilson quarks on fine lattices. I: first experiences and physics results*, *JHEP* **02** (2007) 056 [[hep-lat/0610059](#)]; *QCD with light Wilson quarks on fine lattices. II: DD-HMC simulations and data analysis*, *JHEP* **02** (2007) 082 [[hep-lat/0701009](#)].
- [11] ETM collaboration, P. Boucaud et al., *Dynamical twisted mass fermions with light quarks*, *Phys. Lett.* **B 650** (2007) 304 [[hep-lat/0701012](#)].
- [12] JLQCD and TWQCD collaboration, H. Fukaya et al., *Two-flavor lattice QCD simulation in the epsilon-regime with exact chiral symmetry*, *Phys. Rev. Lett.* **98** (2007) 172001 [[hep-lat/0702003](#)]; *Two-flavor lattice QCD in the epsilon-regime and chiral random matrix theory*, *Phys. Rev.* **D 76** (2007) 054503 [[arXiv:0705.3322](#)].
- [13] RBC AND UKQCD collaboration, C. Allton et al., *2+1 flavor domain wall QCD on a $(2fm)^3$ lattice: light meson spectroscopy with $L_s = 16$* , *Phys. Rev.* **D 76** (2007) 014504 [[hep-lat/0701013](#)].

- [14] PACS-CS collaboration, Y. Kuramashi, *Dynamical Wilson quark simulations toward the physical point*, in *Proceedings of the 25th international symposium on lattice field theory*, Regensburg (2007), PoS(LAT2007)017.
- [15] PACS-CS collaboration, N. Ukita et al., *Light hadron spectrum with 2 + 1 flavor dynamical $O(a)$ -improved Wilson quarks*, in *Proceedings of the 25th international symposium on lattice field theory*, Regensburg (2007), PoS(LAT2007)138.
- [16] Ph. de Forcrand, *Progress on lattice QCD algorithms*, *Nucl. Phys.* **47** (*Proc. Suppl.*) (1996) 228 [[hep-lat/9509082](#)].
- [17] H. Neff, N. Eicker, T. Lippert, J.W. Negele and K. Schilling, *On the low fermionic eigenmode dominance in QCD on the lattice*, *Phys. Rev.* **D 64** (2001) 114509 [[hep-lat/0106016](#)].
- [18] R.B. Morgan and W. Wilcox, *Deflation of eigenvalues for GMRES in lattice QCD*, *Nucl. Phys.* **106** (*Proc. Suppl.*) (2002) 1067 [[hep-lat/0109009](#)].
- [19] L. Giusti, C. Hoelbling, M. Lüscher and H. Wittig, *Numerical techniques for lattice QCD in the epsilon-regime*, *Comput. Phys. Commun.* **153** (2003) 31 [[hep-lat/0212012](#)].
- [20] L. Giusti, P. Hernández, M. Laine, P. Weisz and H. Wittig, *Low-energy couplings of QCD from current correlators near the chiral limit*, *JHEP* **04** (2004) 013 [[hep-lat/0402002](#)].
- [21] T.A. DeGrand and S. Schaefer, *Improving meson two-point functions in lattice QCD*, *Comput. Phys. Commun.* **159** (2004) 185 [[hep-lat/0401011](#)].
- [22] SESAM collaboration, G.S. Bali, H. Neff, T. Duessel, T. Lippert and K. Schilling, *Observation of string breaking in QCD*, *Phys. Rev.* **D 71** (2005) 114513 [[hep-lat/0505012](#)].
- [23] J. Foley et al., *Practical all-to-all propagators for lattice QCD*, *Comput. Phys. Commun.* **172** (2005) 145 [[hep-lat/0505023](#)].
- [24] M. Lüscher, *Local coherence and deflation of the low quark modes in lattice QCD*, *JHEP* **07** (2007) 081 [[arXiv:0706.2298](#)].
- [25] R.C. Brower, T. Ivanenko, A.R. Levi and K.N. Orginos, *Chronological inversion method for the Dirac matrix in Hybrid Monte Carlo*, *Nucl. Phys.* **B 484** (1997) 353 [[hep-lat/9509012](#)].
- [26] B. Sheikholeslami and R. Wohlert, *Improved continuum limit lattice action for QCD with Wilson fermions*, *Nucl. Phys.* **B 259** (1985) 572.
- [27] M. Lüscher, S. Sint, R. Sommer and P. Weisz, *Chiral symmetry and $O(a)$ improvement in lattice QCD*, *Nucl. Phys.* **B 478** (1996) 365 [[hep-lat/9605038](#)].
- [28] ALPHA collaboration, K. Jansen and R. Sommer, *$O(a)$ improvement of lattice QCD with two flavors of Wilson quarks*, *Nucl. Phys.* **B 530** (1998) 185 [*Erratum ibid.* **643** (2002) 517] [[hep-lat/9803017](#)].
- [29] L. Del Debbio, L. Giusti, M. Lüscher, R. Petronzio and N. Tantalo, *Stability of lattice QCD simulations and the thermodynamic limit*, *JHEP* **02** (2006) 011 [[hep-lat/0512021](#)].
- [30] J.C. Sexton and D.H. Weingarten, *Hamiltonian evolution for the Hybrid Monte Carlo algorithm*, *Nucl. Phys.* **B 380** (1992) 665.