

Parameters of the openQCD main programs

Martin Lüscher

February 2012

1. Introduction

The openQCD main programs have many adjustable parameters. Most parameter values are passed to the programs at run time through a human-readable parameter file. The only parameters that must be specified at compilation time are the lattice sizes and the MPI process grid (see `include/global.h` and `main/README.global`). Each program also has a few command-line options.

Some of the parameters are of a fairly trivial kind and are omitted in this note (see, instead, the `README` files that come with each program). Details about the simulation algorithm and the exact normalization conventions employed can be found in the documentation files in the `doc` directory.

2. Preliminaries

The input parameter files are divided into sections such as

```
[Solver 3]
solver CGNE
nmx 256
res 1.0e-10
```

In this case, the section describes the solver number 3 for the Dirac equation. The solver program is `CGNE`, the conjugate-gradient (CG) algorithm for the normal Dirac equation, while `nmx` specifies the maximal number of CG iterations that may be performed and `res` the desired relative residue of the calculated solutions.

Sections start with a title in square brackets [...] and the lines within a section can be ordered arbitrarily. The section quoted above, for example, could equivalently be written as

```
[Solver 3]
res 1.0e-10
solver CGNE    # Should try a better solver
nmix 256
```

Any text appearing to the right of the number sign # is considered to be a comment. A section is delimited by its headline (the text in square brackets) and the headline of the next section. Blank lines are ignored and sections can appear in any order in the input file.

Once the program has read the parameter file, the data are entered into a parameter data base and can, from there, be retrieved by the subprograms that depend on some of the parameter values. The data base is administered by a set of modules in the directory `modules/flags`. In case the exact meaning of a particular parameter is unclear, it may be helpful to read the explanations on the top of these program files.

3. Actions

The lattice theory is defined by the gauge and the pseudo-fermion actions. In the case of the simulation programs, a description of all actions must be included in the parameter file.

3.1 Gauge action

Currently the supported gauge actions include the Wilson plaquette action, the tree-level improved Symanzik action and, more generally, any linear combination of these [1]. The associated parameter section

```
[Action 0]
action ACG
```

has a single line containing the tag `ACG`, while the inverse gauge coupling `beta` must be specified in a separate section

```
[Lattice parameters]
beta 5.3
```

```

c0 1.6667
kappa_u 0.13635
kappa_s 0.13620
kappa_c 0.13550
csw 1.9078
cG 1.10
cF 0.95

```

together with the coefficient `c0` of the plaquette part of the action [1] and the hopping parameters `kappa_u`, `kappa_s`, `kappa_c` of the up, strange and charm quark. The last three parameters in this section are $O(a)$ improvement coefficients (see refs. [2,3]).

3.2 Twisted-mass pseudo-fermion actions

Pseudo-fermion actions

$$S_{\text{pf}} = (\phi, (D^\dagger D + \mu_0^2)^{-1} \phi) \quad (3.1)$$

with a single twisted-mass parameter μ_0 are described by an action section

```

[Action 1]
action ACF_TM1
ipf 0
im0 0
imu 0
isp 3

```

The number in the headline is an index that serves as a tag for the different actions. It can be chosen arbitrarily in the range $0, 1, \dots, 31$, but must be unambiguous, i.e. there may be at most one action section in the parameter file with a given index.

The action (3.1) depends on the pseudo-fermion field ϕ , the bare quark mass m_0 in the Dirac operator and the twisted mass μ_0 . Pseudo-fermion fields are distinguished by an index `ipf`= $0, 1, \dots, \text{npf}-1$, where `npf` is the total number of these fields. The bare sea-quark masses $m_{0,u}$, $m_{0,s}$ and $m_{0,c}$ are similarly referred to by an index `im0`= $0, 1, 2$, while twisted masses are specified by giving their index `imu` in an array of values defined elsewhere (see subsect. 4.1). The last parameter in the section, `isp`, is the index of the solver for the Dirac equation, which is to be used by the program that computes the action (3.1) (see sect. 5 for the list of the available solvers).

Hasenbusch pseudo-fermion actions

$$S_{\text{pf}} = (\phi, (D^\dagger D + \mu_1^2)(D^\dagger D + \mu_0^2)^{-1} \phi) \quad (3.2)$$

with two twisted-mass parameters μ_0 and μ_1 are described by an action section

```
[Action 2]
action ACF_TM2
ipf 1
im0 0
imu 0 1
isp 1 0
```

As explained in the notes [4], these actions arise when the light-quark determinant is split into several factors. The two masses correspond to the two entries on the lines with tag `imu` and `isp`. Note that one needs to specify two solvers, the first for the Dirac equation with twisted mass μ_0 and the second for the equation with twisted mass μ_1 (which must be solved when ϕ is generated at the beginning of the molecular-dynamics trajectories).

3.3 Even-odd preconditioned pseudo-fermion actions

When even-odd preconditioning is used, the pseudo-fermion action (3.1) gets replaced by

$$S_{\text{pf}} = (\phi, (\hat{D}^\dagger \hat{D} + \mu_0^2)^{-1} \phi) - 2 \ln \det D_{\text{oo}}, \quad (3.3)$$

where \hat{D} is the even-odd preconditioned Dirac operator and D_{oo} the odd-odd part of the Dirac operator [3]. The factor $\det D_{\text{oo}}$ is referred to as the “small determinant” and it is understood that the pseudo-fermion field ϕ vanishes on the odd sites of the lattice.

Apart from the action line, the parameter sections

```
[Action 3]
action ACF_TM1_EO_SDET
ipf 2
im0 0
imu 0
isp 3
```

describing the action (3.3) look the same as the ones describing the actions without even-odd preconditioning. The same comment applies to the sections

```
[Action 4]
action ACF_TM2_EO
ipf 3
```

```

im0 0
imu 0 1
isp 1 0

```

that the describe the even-odd preconditioned version

$$S_{\text{pf}} = (\phi, (\hat{D}^\dagger \hat{D} + \mu_1^2)(\hat{D}^\dagger \hat{D} + \mu_0^2)^{-1} \phi) \quad (3.4)$$

of the Hasenbusch pseudo-fermion action (there is no small-determinant contribution in this case).

3.4 Rational function actions

The RHMC algorithm for the strange and the charm quark is based on the Zolotarev rational approximation of the operator $(\hat{D}^\dagger \hat{D})^{-1/2}$ (see ref. [5] for detailed explanations). Zolotarev rational functions are defined by a section

```

[Rational 0]
degree 10
range 0.025 6.02

```

that specifies the degree (number of poles) of the function and the approximation range on the axis of eigenvalues of $(\hat{D}^\dagger \hat{D})^{1/2}$. The range should be sufficiently large to include, with probability practically equal to 1, the whole spectrum of the operator.

The poles and zeros of the Zolotarev rational functions are ordered such that the larger ones come first. For reasons of stability and performance, the rational function should be split into a product of factors, each including a range of poles and zeros [5]. The pseudo-fermion action associated with such a factor is described by a section

```

[Action 5]
action ACF_RAT
ipf 4
im0 1
irat 0 7 9
isp 1

```

The parameters `ipf`, `im0` and `isp` in this section have the same meaning as in the case of the twisted-mass pseudo-fermion actions, while `irat` specifies the index of the rational function and the range of poles included in the factor. Poles are counted from zero and pole ranges are inclusive, i.e. `irat 0 7 9` selects the poles number 7, 8 and 9 from the rational function with index 0.

Since even-odd preconditioning is used for the strange and the charm quark, the associated quark determinants contain the small determinant $\det D_{\text{oo}}$. It is convenient to combine this factor with the rational factor that contains the largest poles. The corresponding section is

```
[Action 6]
action ACF_RAT_SDET
ipf 5
im0 1
irat 0 0 6
isp 4
```

For a given rational function and mass index `im0`, the action sections must be such that all poles and the small determinant are included once. In the example discussed here, this is achieved with two factors, but one is free to split the rational functions in more factors. Note that for each factor a different pseudo-fermion field must be used.

4. Molecular-dynamics parameters

The simulation algorithm implemented in the openQCD package involves a numerical integration of the molecular-dynamics equations that derive from the chosen action. Apart from the action, this requires the integration scheme (the “integrator”) to be specified as well as the solvers used for the computation of the various pseudo-fermion forces.

4.1 HMC parameters

Some basic parameters of the algorithm are collected in the section

```
[HMC parameters]
actions 0 1 2
npf 6
mu 0.015 0.2 1.0
nlv 2
tau 1.0
```

The numbers on the `action` line are the indices of the actions that are to be included in the molecular-dynamics Hamilton function. If some of these actions depend on

twisted mass parameters, their values must be listed on the line with tag `mu`. As already mentioned, the action sections refer to these values by quoting their position `imu` (counting from 0) in the array. The parameter `npf` specifies the total number of pseudo-fermion fields that must be allocated.

The last two parameters, `nlv` and `tau`, define the number of integrator levels and the molecular-dynamics trajectory length (see ref. [4] for the normalization of the latter; the integration scheme is discussed in subsect. 4.3).

4.2 Forces

The forces that derive from the actions `ACG`, `...`, `ACF_RAT_SDET` discussed in sect. 3 are distinguished by the symbols `FRG`, `...`, `FRF_RAT_SDET`. The corresponding sections are

```
[Force 0]
force FRG

[Force 1]
force FRF_TM1
isp 6
ncr 4

[Force 2]
force FRF_TM2
isp 7
ncr 0

[Force 3]
force FRF_TM1_EO_SDET
isp 6
ncr 4

[Force 4]
force FRF_TM2_EO
isp 7
ncr 0

[Force 5]
force FRF_RAT
isp 8
```

```

[Force 6]
force FRF_RAT_SDET
isp 9

```

In each case, the index in the headline must match the index of the associated action. The force tags (such as `FRF_TM2_E0`) could be inferred from the corresponding action sections, but are required in order to improve the readability of the input parameter files.

Most parameters of the forces are inherited from those of the associated actions. The solver parameter sets selected by the indices `isp` may however be different. In the case of the forces `FRF_TM2` and `FRF_TM2_E0`, the specified solver is used for the solution of the Dirac equation with twisted mass μ_0 . No solver is needed here for the second twisted mass μ_1 .

The parameter `ncr` controls the chronological propagation of the solutions of the Dirac equation along the molecular-dynamics trajectories. If `ncr` is set to a positive value, the simulation program attempts to reduce the number of solver iterations required for the computation of the specified force by extrapolating the previous `ncr` solutions in molecular-dynamics time. The feature is switched off if `ncr` is set to zero.

4.3 Integration scheme

The numerical integration of the molecular-dynamics equations makes use of a hierarchical integrator that can be specified on the input parameter file. Currently three elementary integration schemes are supported, the leapfrog, the 2nd order Omelyan–Mryglod–Folk (OMF) and a 4th order OMF integrator. They are distinguished by the symbols `LPFR`, `OMF2` and `OMF4`.

Hierarchical integrators have several levels with increasing integration step sizes. They are described by the total integration time `tau`, the number `nlv` of levels, the elementary integrators used at each level and the forces integrated there. The parameters `tau` and `nlv` are part of the HMC parameter set (see subsect. 4.1). Each level is then described by a section like

```

[Level 2]
integrator LPFR
nstep 12
forces 2 4 5

```

In this example, the simulation program is instructed to use the leapfrog integrator at the third level (levels are counted from 0). The elementary leapfrog integration step is to be applied 12 times and the forces integrated at this level are the ones

with index 2,4 and 5. If the level is the top level, the integration step size for these forces is thus $\tau/12$.

An example of a complete integrator description is provided by the following three sections:

```
[Level 0]
integrator OMF4
nstep 1
forces 0

[Level 1]
integrator OMF4
nstep 1
forces 1 2 3

[Level 2]
integrator OMF2
lambda 0.16667
nstep 6
forces 4
```

There are three levels in this case, with average step sizes equal to $\tau/300$, $\tau/60$ and $\tau/12$ at level 0, 1 and 2, respectively (note that the OMF2 and OMF4 integrators update the gauge field 2 and 5 times per application). The OMF2 integrator depends on a parameter `lambda`, whose value must be given at each level where this integrator is used. See `modules/update/README.mdint` for further explanations.

5. Solver parameters

All solvers referred to in the action and force sections must be described in the input parameter file. The solvers are labeled by an index `isp` in the range $0, 1, \dots, 31$ in much the same way as the actions and the forces. There is no one-to-one correspondence between solver programs and solver sections in the parameter file. One may, for example, have two sections

```
[Solver 3]
solver CGNE
nmix 256
```

```

res 1.0e-10

[Solver 4]
solver CGNE
nmx 256
res 1.0e-8

```

where the only difference is the value of the desired residue `res`. By setting `isp=3` or `isp=4` in an action or force section, the numerical accuracy of the action and force computations can then be individually controlled.

5.1 Conjugate gradient solvers

There are two conjugate gradient solvers, the ordinary CG algorithm for the normal Dirac equations,

$$(D^\dagger D + \mu^2)\psi = \eta \quad \text{and} \quad (\hat{D}^\dagger \hat{D} + \mu^2)\psi = \eta, \quad (5.1)$$

and the mult-shift CG algorithm for the even-odd preconditioned simultaneous equations [5,6]

$$(\hat{D}^\dagger \hat{D} + \mu_k^2)\psi_k = \eta \quad k = 0, 1, \dots, n-1. \quad (5.2)$$

Examples of parameter sections describing solvers based on the former have already appeared in this note. The parameter sections

```

[Solver 3]
solver MSCG
nmx 256
res 1.0e-10

```

for the multi-shift solver look practically the same. Note that this solver can only be used for rational function actions and forces.

5.2 SAP preconditioned solvers

There are currently two solvers that make use of the Schwarz Alternating Procedure (SAP) as preconditioner for the GCR algorithm. When these solvers are used, the block size `bs` of the SAP block grid needs to be specified in a separate section

```

[SAP]
bs 4 6 6 4

```

The ordinary SAP preconditioned GCR algorithm is then described by a parameter section

```
[Solver 7]
solver SAP_GCR
nkv 16
isolv 1
nmr 4
ncy 5
nmx 24
res 1.0e-8
```

where **nkv** is the maximal number of Krylov vectors that may be generated before the GCR algorithm is restarted, while **nmx** is the maximal total number of generated Krylov vectors and **res** the desired relative residue of the calculated solutions. All other parameters in the section describe the particular SAP preconditioner to be used, namely **ncy** and **nmr** specify the number of SAP cycles and block solver iterations to be applied and the flag **isolv** indicates whether the block solver is even-odd preconditioned (**isolv=1**) or not (**isolv!=1**).

The other solver that makes use of the SAP is the deflated SAP preconditioned GCR solver. This solver is described by a section

```
[Solver 8]
solver DFL_SAP_GCR
nkv 16
isolv 1
nmr 4
ncy 5
nmx 24
res 1.0e-8
```

that coincides with the section for the ordinary SAP preconditioned solver except for the solver symbol.

5.3 Deflation parameters

When the deflated solver is used, the parameters related to the deflation subspace must be specified in a few further sections. The block size **bs** of the deflation block grid and the number **Ns** of deflation modes per block are set by including the section

```
[Deflation subspace]
bs 8 4 4 4
```

Ns 20

For the generation of the deflation subspace, a bootstrapped inverse iteration algorithm is employed. In this process, the hopping parameter `kappa` and the twisted mass `mu` are set to the values specified in the section

```
[Deflation subspace generation]
kappa 0.1350
mu 0.01
ninv 9
nmr 4
ncy 4
nkx 16
nmx 96
res 1.0e-3
```

The other parameters set in this section include the number `ninv` of inverse iteration steps to be applied and the numbers `ncy` and `nmr` of SAP cycles and block solver iterations to be used when the SAP preconditioner is invoked. The last three parameters, `nkx`, `nmx` and `res`, refer to the GCR algorithm used for the solution of the little Dirac equation and have exactly the same meaning as the corresponding parameters of the SAP preconditioned GCR algorithm (see subsect. 5.2).

In the course of the deflated solver iterations, the single-precision and the double-precision deflation projectors are frequently applied. Both operations require the little Dirac equation to be solved. The section

```
[Deflation projectors]
nkx 24
nmx 256
res 1.0e-6
resd 1.0e-12
```

sets the parameters of the GCR solver used for this task, `res` and `resd` being the desired residues of the calculated solutions when the equation is solved using single- and double-precision arithmetic, respectively.

Along the molecular-dynamics trajectories, the deflation subspace loses its efficiency and must therefore be refreshed from time to time. This feature is controlled by the parameter section

```
[Deflation update scheme]
dtau 0.091
nsm 1
```

where `dtau` sets the interval in molecular-dynamics time after which the subspace is updated and `nsn` the number of smoothing steps to be applied in this process. The subspace updates can be turned off by setting `nsn=0`. Clearly, the section is not needed in measurement programs, where the deflated solver may be used for quark propagator calculations, for example.

6. Reweighting factors

In general, the ensembles of field configurations generated in openQCD simulations must be reweighted. The rational approximation used for the strange and the charm quark determinants, for example, requires reweighting in order to correct for the approximation error. In the case of the light quark doublet, reweighting is required if a twisted mass term was added as infrared regulator [7].

For a given ensemble of field configurations, the program `ms1` computes stochastic estimates of a set of reweighting factors. Exactly which factors are to be computed can be specified in the input parameter file.

6.1 Light quark reweighting factors

In ref. [7], two kinds of twisted-mass regularizations of the light quark determinant were proposed which amount to replacing

$$\det\{D^\dagger D\} \rightarrow \det\{D^\dagger D + \mu^2\} \quad (6.1)$$

and

$$\det\{D^\dagger D\} \rightarrow \det\{(D^\dagger D + \mu^2)^2 (D^\dagger D + 2\mu^2)^{-1}\} \quad (6.2)$$

respectively. The twisted mass parameter $\mu > 0$ provides the desired infrared regularization and is usually set to a value on the order of light quark mass.

The associated reweighting factors,

$$W_1 = \det\{D^\dagger D (D^\dagger D + \mu^2)^{-1}\}, \quad (6.3)$$

$$W_2 = \det\{D^\dagger D (D^\dagger D + 2\mu^2) (D^\dagger D + \mu^2)^{-2}\}, \quad (6.4)$$

can be estimated stochastically through

$$W_{1,N} = \frac{1}{N} \sum_{k=1}^N \exp\{-\mu^2(\eta_k, (D^\dagger D)^{-1} \eta_k)\}, \quad (6.5)$$

$$W_{1,N} = \frac{1}{N} \sum_{k=1}^N \exp\{-\mu^4(\eta_k, (D^\dagger D)^{-1} (D^\dagger D + 2\mu^2)^{-1} \eta_k)\}, \quad (6.6)$$

where η_1, \dots, η_N are random quark fields with normal distribution. In the parameter files, these factors are described by sections of the form

```
[Reweighting factor 0]
rwfact RWTM1
im0 0
mu 0.003
isp 0
nsrc 24
```

As usual the index in the headline serves to distinguish different reweighting factors. The line with tag `rwfact` selects the reweighting factor type (RWTM1 for W_1 and RWTM2 for W_2), while the other parameters, `im0`, `mu`, `isp` and `nsrc`, are the index of the bare quark mass m_0 , the value of the twisted mass μ , the index of the solver for the Dirac equation to be used and the number N of random source fields to be generated for each gauge field configuration.

If even-odd preconditioning is used, eqs. (6.1) and (6.2) get replaced by

$$\det\{D^\dagger D\} \rightarrow \det\{(D_{\text{oo}})^2\} \det\{\hat{D}^\dagger \hat{D} + \mu^2\}, \quad (6.7)$$

$$\det\{D^\dagger D\} \rightarrow \det\{(D_{\text{oo}})^2\} \det\{(\hat{D}^\dagger \hat{D} + \mu^2)^2 (\hat{D}^\dagger \hat{D} + 2\mu^2)^{-1}\}. \quad (6.8)$$

The associated reweighting factors,

$$\hat{W}_1 = \det\{\hat{D}^\dagger \hat{D} (\hat{D}^\dagger \hat{D} + \mu^2)^{-1}\}, \quad (6.9)$$

$$\hat{W}_2 = \det\{\hat{D}^\dagger \hat{D} (\hat{D}^\dagger \hat{D} + 2\mu^2) (\hat{D}^\dagger \hat{D} + \mu^2)^{-2}\}, \quad (6.10)$$

can again be estimated stochastically. The corresponding parameter sections in the input files are as before except that the symbols `RWTM1` and `RWTM2` must be replaced by `RWTM1_EO` and `RWTM2_EO`.

6.2 Strange and charm quark reweighting factors

The reweighting factors required for the heavier quarks are given by

$$W_R = \det\{\hat{D}R\}, \quad (6.11)$$

where R is the rational function used to approximate the operator $(\hat{D}^\dagger \hat{D})^{-1/2}$ (see ref. [5] for further details). In the input parameter files, these factors are described by sections like

```
[Reweighting factor 1]
rwfact RWRAT
im0 1
irp 0
np 6 4
isp 1 0
nsrc 1
```

As before, `im0` is the index of the bare mass m_0 of the quark considered and `nsrc` the number of random source fields to be used for the stochastic estimation of W_R . The other parameters select the rational function R and determine exactly how the reweighting factor is to be computed.

Rational functions are defined by separate parameter sections (see subsect. 3.4). The parameter `irp` must be set to the index of the section that describes the function R . As already mentioned, the poles of R are ordered such that the larger ones come first. The program that computes the reweighting factor expands R into partial fractions and logically divides the pole terms into groups R_k of terms,

$$R = A\{1 + R_0 + R_1 \dots + R_{n-1}\}, \quad (6.12)$$

where R_0 contains the first few poles, R_1 the next few poles, and so on. The figures on the line with tag `np` specify the numbers of poles in these groups. In the example above, there are two parts, R_0 and R_1 , where R_0 includes the poles $0, 1, \dots, 5$ and R_1 the poles $6, \dots, 9$. Any sequence of positive numbers can appear on the line with tag `np` subject to the constraint that their sum coincides with the number of poles of R .

For each part R_k , the solver for the (twisted mass) Dirac equation can be chosen separately. The indices of the corresponding solver parameter sets must be given on the line with tag `isp`, where it is understood that the first solver is to be used when the operator R_0 is applied to the source field, the second when R_1 is applied, and so on. Currently the supported solvers are `MSCG`, `SAP_GCR` and `DFL_SAP_GCR` (see sect. 5).

References

- [1] *Gauge actions in openQCD simulations*, file `doc/gauge_action.pdf`
- [2] M. Lüscher, Stefan Schaefer, *Lattice QCD without topology barriers*, JHEP 1107 (2011) 036
- [3] *Implementation of the lattice Dirac operator*, file `doc/dirac.pdf`
- [4] *Molecular-dynamics quark forces*, file `doc/forces.pdf`
- [5] *Charm and strange quark in openQCD simulations*, file `doc/rhmc.pdf`
- [6] *Multi-shift conjugate gradient algorithm*, file `doc/mscg.pdf`
- [7] M. Lüscher, F. Palombi, *Fluctuations and reweighting of the quark determinant on large lattices*, PoS (LATTICE2008) 049