

Quick reference manual
for the SPHInX DFT program
Release 3.2

Christoph Freysoldt, freysoldt@mpi-susmat.de

July 6, 2026



Contents

1	Introduction	3
1.1	About	3
1.2	sphinx invocation	3
1.3	Parallelism	3
1.4	Units	5
2	The SPHInX parser	5
2.1	Overview	5
2.2	Expressions	6
2.3	Included files	6
2.4	The format statement, std files.	6
3	Overview on the input file for DFT calculations	7
3.1	The <code>structure</code> group	7
3.1.1	The <code>species</code> group	8
3.1.2	The <code>atom</code> group	9
3.1.3	<code>symmetry</code> group	9
3.2	The <code>basis</code> group	10
3.2.1	The <code>kPoint</code> group	11
3.2.2	The <code>kPoints</code> group	11
3.2.3	The <code>from</code> group	12
3.2.4	The <code>to</code> group	12
3.3	The <code>pawPot</code> group	13
3.4	The <code>PAWHamiltonian</code> group	14
3.4.1	The <code>vExt</code> group	17
3.4.2	The <code>xcMesh</code> group	17
3.4.3	The <code>vdwCorrection</code> group	17
3.4.4	The <code>HubbardU</code> group	18
3.4.5	The <code>site</code> group	18
3.4.6	The <code>A0</code> group	19
3.4.7	The <code>M0</code> group	20
3.4.8	The <code>orbital</code> group	22
3.5	The <code>spinConstraint</code> group	22
3.6	The <code>initialGuess</code> group	23
3.6.1	The <code>waves</code> group	24
3.6.2	The <code>lcao</code> group	24
3.6.3	The <code>rho</code> group	24
3.6.4	The <code>atomicSpin</code> group	25
3.6.5	The <code>aspherical</code> group	26
3.6.6	The <code>charged</code> group	28
3.6.7	The <code>occupations</code> group	29
3.6.8	The <code>exchange</code> group	30
3.7	The <code>pseudoPot</code> group	30
3.8	The <code>PWHamiltonian</code> group	31

3.9	The main group	33
4	Electronic loop	34
4.1	scfDiag: iterative diagonalization + density mixing	34
4.1.1	The CCG group	36
4.1.2	The blockCCG group	36
4.1.3	The preconditioner group	37
4.2	CCG: direct minimization	38
5	Geometry optimizers	39
5.1	The QN group	39
5.1.1	The bornOppenheimer group	40
5.2	The linQN group	40
5.3	The ricQN group	41
5.3.1	The ric group	42
5.4	The ricTS group	43
5.4.1	The transPath group	45
5.5	The extControl group	46
6	Extra computations (in main group)	48
6.1	evalForces: force evaluation	48
6.2	pDos: PAW projected DOS evaluation	48
7	Output	49
7.1	energy.dat: total energies	49
7.2	residue.dat: density residues	50
7.3	eps.dat: eigenvalues	50
7.4	rho.sxb: density (binary)	50
7.5	waves.sxb: wave functions (binary)	50
7.6	relaxedStr.sx: final atomic structure	51
7.7	relaxHist.sx: geometry optimization history	51
7.8	energy-structOpt.dat: geometry optimization energies	51
7.9	fftwisdom.dat: FFTW plans	51
7.10	hubbardOcc.dat: DFT+U occupations	52
	Index	54

1 Introduction

1.1 About

SPHInX is a C++ programming library for developing efficient physics codes [1]. Its dominant feature is a plane-wave density-functional theory (DFT) code named `sphinx`. This document is a short, non-tutorial manual for the DFT program `sphinx`. Other features of `sphinx` are not described here. In particular, $\mathbf{k} \cdot \mathbf{p}$, EXX, and tight-binding, are not covered.

1.2 `sphinx` invocation

`sphinx` reads its input from a main input file, usually called ‘`input.sx`’. Its format is described in Sec. 3. `sphinx` produces output in several dedicated files (see Sec. 7), but the main output comes on stdout (i.e. the screen). `sphinx` can be asked to duplicate (or redirect) its output into a log file. `sphinx` – like all SPHInX executables – can take options on the command line that control some of its global behavior.

Option	argument	description
<code>--help</code>		show all the available options
<code>--log</code>		create a log file
<code>--quiet</code>		don’t produce output on stdout
<code>--input</code>	filename	input file (default: <code>input.sx</code>)

If SPHInX has been compiled with the FFTW library for the fast Fourier transforms (FFTs), there are two additional options to access the planning mode of FFTW (`--fastfft` and `--wisdom`).

1.3 Parallelism

SPHInX uses MPI (message passing interface) high-level parallelism and openMP low-level parallelism. Both features have to be enabled at compile time!

MPI parallelism allows to distribute memory and workload over independent tasks (processes) which may run even on different computers. At present, SPHInX uses MPI mostly for \mathbf{k} -point parallelization, parallelization over atoms in the computation of the PAW one-center contributions, and for exact-exchange workload. MPI parallelism is enabled by starting SPHInX with `mpirun`.

Example:

Run `sphinx` with 20 MPI tasks:

```
mpirun -n 20 sphinx
```

openMP parallelism uses multiple threads within the same process, allowing to share the workload among multiple CPU cores on the same computer. openMP is used for low-level parallelization within the numerical libraries (BLAS,

FFT) and for some internal routines (e.g. exchange-correlation potential, non-local projectors on x86-64-AVX machines). openMP is enabled by setting the `SX_THREADS` environment variable.

Example:

Run `sphinx` with 4 openMP threads:

```
export SX_THREADS=4
sphinx
```

Note that environment variables in c-shell variants (`csh`, `tcsh`) are set differently

```
setenv SX_THREADS 4
sphinx
```

Efficiency. Parallelism always comes at a cost, i.e., the calculation never runs proportionally faster if you provide more cores. This is because there are always some parts that do not run in parallel, AND there is an intrinsic overhead in communication, AND we are wasting some time when waiting for the slowest task to finish. It is impossible to predict how efficiently parallelization works for your calculation – you must test it for a specific system. Yet, some general rules apply

- Efficiency stops when there's nothing to parallelize. If you have 10 **k**-points, no more than 10 MPI tasks can do work.
- We don't use internal load balancing. Efficiency is best if task can be distributed equally (bad example: 15 **k**-points on 7 MPI tasks: a single MPI task must do three **k**-points).
- openMP efficiency levels off well below the total core count. 4-8 openMP threads work reasonably well on modern machines, but not much more.
- Hyperthreading within the CPU is bad for performance. Only use physical cores.
- If you do care about parallelization efficiency, use the detailed timings at the end of the output to determine bottlenecks and how well they parallelize.

MPI and openMP parallelism can be combined when you run into efficiency limits with only one type of parallelization. However, the total core efficiency (i.e. how much of the serial performance of each core you get) is the product of the individual efficiency. So if you have 70% efficiency for MPI and 70% efficiency for openMP, the combination runs at $0.7 \cdot 0.7 = 49\%$ efficiency.

1.4 Units

SPHInX uses mostly atomic units: bohr for coordinates, Hartree atomic units for energies. Some exceptions exist, e.g. for plane-wave cutoffs (Rydberg), electronic temperature (eV), eigenvalues (eV), ...

2 The SPHInX parser

2.1 Overview

The input file format is best explained starting from an **example**:

```
basis {
  eCut = 20; // Ry
  kPoint {
    coords = [1/2,1/2,1/2];
    weight = 1;
    relative;
  }
  folding = 4 * [1,1,1];
}
```

The SPHInX input format is a structured, hierarchical format with a C-like syntax. It consists of named groups and parameters. Parameter and group names are case-sensitive.

Groups. The content of a **group** (such as **basis**) is enclosed in curly brackets {}. Groups may contain parameters, and other groups.

The order of unique groups is *not* important. Groups that may appear multiple times, or groups that describe an action of the code (notably in the **main** group) are processed in the order of appearance.

Parameters. **Parameters** (such as **eCut**) are assigned values with the equal sign. Parameter assignments must be followed by a semicolon (;). Some parameters/variables may be a vector (or a list), for instance **coords** in the example above. The vector elements are comma-separated and enclosed by square brackets []. Parameters may expect string values. The string is enclosed in double quotes ("). Parameters that expect filenames may be enclosed by <> to look for the files in the search path. Enclosure by double quotes looks for them in the current directory.

Flags. Flags are special parameters, that normally do not carry a value. A flag (e.g. **relative**) is set by specifying its name, followed by a semicolon. Flags are unset by assigning the value 0.

White space. White space (including tabs and newlines) can be added quite freely, except within words, numbers, etc.

Comments. **Comments** can be added by the `//` and `/* */` syntax. A `//` comment extends until the rest of the line.¹ A `/* */` comment omits everything between the `/*` and `*/` markers.

2.2 Expressions

The SPHInX parser supports basic algebraic expressions, such as adding, subtracting, multiplying etc. All numbers are changed to double precision when doing so, i.e., $1/2$ is equivalent to `0.5`. Some mathematical functions (`sqrt`, `cbrt` (cubic root), `sin`, `cos`, `exp`, `log`) are available. At the top level, additional variables may be set and used in algebraic expressions. Strings can be concatenated with `+`.

2.3 Included files

The input file may include other files via the include statement.

```
include <parameters.sx>;  
include "startStructure.sx";
```

Similar to C/C++, double quotes indicate that the file is expected in the current directory (of execution), while `<>` indicates to look for the file in the parser's search path. The file `parameters.sx` (located in `share/sphinx/`) contains a large number of predefined keys that may offer mnemonic names for some (numeric) settings. It should always be included.

2.4 The format statement. std files.

Each input file should begin with a `format` statement, e.g. for PAW

```
format paw;
```

What is the role of this statement? The SPHInX input file is read by the SPHInX parser, and validated against a syntax type definition (`std`). The relevant `std` file is selected by the `format` statement in the input file (here, it would load `'paw.std'`). The `std` files are located² in `share/sphinx/std/`. The `std` files themselves have a similar format as the actual input file, but define the type expected for each parameter, as well as mutual exclusions, range limits, etc. The `std` files can be an additional resource for discovering what SPHInX offers. The `std` files make heavy use of the include feature.

¹Instead of `//` one can also use `#` (like in shell) or `%` (like in LaTeX).

²within the `sphinx/src` folder in the source tree, and within the installation path when installed

3 Overview on the input file for DFT calculations

A typical input file for PAW will look like this:

```
format paw;  
include <parameters.sx>;  
  
structure { ... }  
basis { ... }  
pawPot { ... }  
PAWHamiltonian { ... }  
initialGuess { ... }  
main { ... }
```

The individual groups are described in the other sections of this manual. They define the atomic structure (`structure`), the plane-wave basis set and **k**-points (`basis`), the PAW potentials to be used (`pawPot`), other settings of the PAW Hamiltonian such as the xc functional (`PAWHamiltonian`), how to set up the starting density and wavefunctions (`initialGuess`), as well as the type of calculation (geometry optimization, single-point calculation, band structure, ...) in the `main` group.

A typical input file for norm-conserving pseudopotentials will look like this:

```
format sphinx;  
include <parameters.sx>;  
  
structure { ... }  
basis { ... }  
pseudoPot { ... }  
PWHamiltonian { ... }  
initialGuess { ... }  
main { ... }
```

Quite similar, no? The difference is in the format, the `pseudoPot` and `PWHamiltonian` groups.

3.1 The structure group

The structure group specifies the atomic positions in SPHInX format.

Example:

```

structure {
  cell = 10.2 * [[0, 1/2, 1/2],
                 [1/2, 0, 1/2],
                 [1/2, 1/2, 0]];

  species {
    element="Si";
    atom { coords = [0,0,0]; relative; }
    atom { coords = [1/4,1/4,1/4]; relative; }
  }
}

```

The following parameters may be set:

parameter	description
movable	(flag) allow atoms to move. Default: all atoms are movable, unless any movable flag is used for any species/atom.
movableX	(flag) allow atoms to move in the x direction. Default: movable, unless movableY or movableZ are used.
movableY	(flag) allow atoms to move in the x direction. Default: movable, unless movableX or movableZ are used.
movableZ	(flag) allow atoms to move in the x direction. Default: movable, unless movableX or movableY are used.
cell	(required) The unit cell (in bohr). This is a list of the three basis vectors in Cartesian coordinates.

The **movable** flags are applied hierarchically. Settings at the structure level can be overridden at the species or atom level. Disabling a movable flag of a surrounding group is achieved by setting to 0, e.g. **movableY** = 0;

The **structure** group must contain at least one **species** group. It may contain a **symmetry** group.

3.1.1 The species group

The **species** group defines atomic positions for one chemical species. Atoms must be sorted by their chemical species.

```

species {
  element="Al ";
  atom { ... }
  atom { ... }
}
species {
  element="O ";
  atom { ... }
  atom { ... }
  atom { ... }
}

```

The `species` group may contain the `movable`, `movableX`, `movableY`, `movableZ` parameters as specified above. In addition it may contain the `element` parameter to indicate the chemical symbol, enclosed by double quotes.

The `species` group must contain at least one `atom` group.

3.1.2 The atom group

The `atom` group defines atomic positions for one atom. Atoms must be sorted by their chemical species.

The `atom` group may contain the `movable`, `movableX`, `movableY`, `movableZ` parameters as specified above in Sec. 3.1. In addition, the following parameters may be set:

parameter	description
<code>coords</code>	(required) The atomic coordinates as a 3-vector. Unless the <code>relative</code> flag is employed, the coordinates are Cartesian (in bohr).
<code>relative</code>	(flag) The coordinates are given relative to the unit cell vectors.
<code>movableLine</code>	(optional) The movement of the atom is restricted to a line. The value gives the direction of the line as a 3-vector.
<code>label</code>	(optional string) Assign a label (or rather a tag) to this atom. If labels are used, atoms with different labels are considered inequivalent. Think of spin configurations for a use-case.

3.1.3 symmetry group

The `symmetry` group (within the structure group) defines the rotational symmetries of the system around the origin of the coordinate system. If not given, the symmetries are determined automatically. However, non-chemical degrees of freedom (such as spins) may break the symmetry. As all forces / displacements are symmetrized, such a situation may require to set the symmetries

by hand. Alternatively, **giving an empty symmetry group switches off symmetrization.**

The **symmetry** group contains multiple **operator** groups. Each operator group contains the parameter **S**, a Cartesian rotation matrix given row-wise. The symmetries must form a group.

Example:

```
symmetry {
  // Symmetry Number 1, E
  operator { S = [[1.000000,0.000000,0.000000],
                  [0.000000,1.000000,0.000000],
                  [0.000000,0.000000,1.000000]]; }
  // Symmetry Number 2, m [0,0.707107,-0.707107]
  operator { S = [[1.000000,0.000000,0.000000],
                  [0.000000,0.000000,1.000000],
                  [0.000000,1.000000,0.000000]]; }
}
```

The best way to set up a reduced symmetry group is to use

```
sxstructsym --printsym -i input.sx
```

and then grep the symmetry group from the output and remove the unwanted symmetries.

3.2 The basis group

The **basis** group defines the plane-wave basis and the **k**-points, and optionally adjusts FFT mesh sizes.

Example:

```
basis {
  eCut = 20; // Ry
  kPoint {
    coords = [1/2,1/2,1/2];
    weight = 1;
    relative;
  }
  folding = 4 * [1,1,1];
}
```

The **basis** group must contain a **kPoint** (Monkhorst-Pack meshes) or a **kPoints** (bandstructures) group.

The following parameters may be set:

parameter	description
<code>eCut</code>	(required) Plane-wave cutoff for wavefunctions (in Rydberg)
<code>gCut</code>	(optional) Override cutoff for the density G-basis (defaults to $4 \times \text{eCut}$)
<code>folding</code>	Monkhorst-Pack folding. 3 -vector of integers. Defaults to [1,1,1]
<code>mesh</code>	(optional) FFT mesh. 3 -vector of integers. Default: automatic
<code>meshAccuracy</code>	(optional) FFT mesh accuracy, relative to the default. Default: 1, i.e. sufficient to hold $2 \times$ the wavefunction plane-waves to avoid wrap-around errors
<code>saveMemory</code>	(flag) Do not store phase-factor for each atom. This may make the code marginally slower, but save some RAM.

3.2.1 The `kPoint` group

The `kPoint` group defines the offset of the Monkhorst-Pack mesh. For this, one (or more) `k`-points are defined within the Brillouin zone. The complete zone with all the `k`-points in it (usually 1) is then shrunk by a factor `folding` in each direction; and the originally Brillouin zone is filled with copies of this shrunk zone. Afterwards, symmetry is used to remove redundant `k`-points.

Alternatively, all `k`-points can be specified via `kPoint` groups, assigning weights as desired.

Example:

```
kPoint { coords = [1/2, 1/2, 1/2]; relative; }
```

The following parameters may be set:

parameter	description
<code>relative</code>	(flag) Coordinates are relative. This should be usually set.
<code>coords</code>	(required) The coordinates. If relative flag is not given, the units are 1/bohr.
<code>weight</code>	(optional) Give a weight to this <code>k</code> -point. Weights must sum to 1. Default: 1/number of <code>k</code> -points.

Typical offsets are 0 or 1/2. For decoupling across vacuum regions with `folding=1`, use 1/4 to minimize band dispersion effects.

3.2.2 The `kPoints` group

The `kPoints` group is used to conveniently define band structure paths. It contains a sequence of `from` and `to` groups. **Note:** Band structure paths are subject to Monkhorst-Pack folding, so set `folding` to [1,1,1] or leave it out.

Example:

```

kPoints {
  dK = 0.01;
  // — fcc lattice
  // L point
  from { coords=PI/aLat * [1,1,1]; label="L"; }
  // Gamma point
  to { coords=[0,0,0]; label="\xG"; }
  // X point
  to { coords=[2 * PI/aLat, 0,0]; label="X"; }
}

```

The group must contain a **from**, and at least one **to** group.

The following parameters may be set:

parameter	description
relative	(flag) Coordinates are relative.
dK	(optional) Set the number of intermediate k -points such that the distance is at most dK (in 1/bohr).

3.2.3 The from group

The **from** group (within the **kPoints** group) adds a single **k**-point at the desired position. It may be used multiple times.

The following parameters may be set:

parameter	description
relative	(flag) Coordinates are relative.
coords	(required) The coordinates. If relative flag is not given, the units are 1/bohr.
label	(optional) String. Give a label to this k -point.

3.2.4 The to group

The **to** group (within the **kPoints** group) adds a line of **k**-points from the previous one to a new position. The number of points is set directly with **nPoints** or indirectly via **dK**.

The following parameters may be set:

parameter	description
relative	(flag) Coordinates are relative.
coords	(required) The coordinates. If relative flag is not given, the units are 1/bohr.
label	(optional) String. Give a label to this k -point.
dK	Set maximum k -point distance.
nPoints	Specify number of points to add. The final one will be at coords .

3.3 The pawPot group

The `pawPot` group defines the PAW potentials, by a sequence of `species` groups. The order of species must agree with the `structure` group.

Example:

```
pawPot {
  species {
    name = "Nitrogen";
    potType = "AbInit";
    element = "N";
    potential="N_LDA_abinit.paw";
    lMaxRho=2;
    angularGrid=4;
  }
  species {
    name = "Hydrogen";
    element="H";
    ...
  }
}
```

The group must contain one `species` group per species. Within the `species` group, the following parameters are commonly set:

parameter	description
<code>name</code>	(optional string) English name of the element
<code>element</code>	(optional string) Chemical symbol
<code>potential</code>	(required filename) Name of the potential file.
<code>potType</code>	(required string) Type of the potential file, see table below

Possible potential file formats:

potType	description
"AbInit"	conventional abinit format (non-xml)
"AtomPAW"	potentials from the atompaw generator (output 2, <code>.atomicdata</code>)
"CPPAW"	potentials of Blöchl's original PAW code
"VASP"	vasp potential file
"xml"	PAW-XML format [2] (abinit, atompaw, gpaw)

The following parameters within the `species` group allow experimenting with some details.

parameter	description
<code>lMaxRho</code>	(optional) Truncate the spherical expansion of densities (and compensation charges) at this l .
<code>angularGrid</code>	(optional integer) Choose a different angular grid for xc calculation in the PAW sphere. Larger is finer. Default is 7 (110 points).
<code>nRadGrid</code>	(optional integer) Interpolate to a different radial grid.
<code>checkOverlap</code>	(flag) Check that PAW norm is guaranteed to be positive definite in the limit of large cutoffs. This is on by default. Some problematic PAW potentials may fail the check, but work normally in some circumstances, so you can switch off the check here.
<code>rPAW</code>	(optional) Change the PAW radius used for atomic quantities "inside the PAW sphere".

3.4 The PAWHamiltonian group

The PAWHamiltonian group defines the DFT functional, the number of empty states, the smearing, and some other settings.

Example:

```
PAWHamiltonian {
  nEmptyStates = 10;
  ekt = 0.02;
  xc = LDA_PW;
}
```

The following common parameters may be set:

parameter	description
<code>xc</code>	(required) xc functional to be used. Constants are defined in <code>parameters.sx</code> . See below.
<code>libxc</code>	(optional) Use an xc functional from libxc. See below.
<code>ekt</code>	(optional) smearing in eV. Should be 0 for semiconductors.
<code>MethfesselPaxton</code>	(optional) If ≥ 0 , use Methfessel-Paxton smearing of indicated order. Order 0 is same as Gaussian smearing.
<code>FermiDirac</code>	(optional) If ≥ 0 , use FermiDirac smearing of indicated order. Order 0 is the default; order 1 corresponds to first-order corrections. Higher orders are not yet implemented.
<code>nEmptyStates</code>	(optional) number of empty states. Remember to scale with system size!
<code>nExcessElectrons</code>	(optional) Number of extra electrons (= minus charge). Can be fractional.
<code>spinPolarized</code>	(flag) Run a collinear spin-polarized calculation.
<code>dipoleCorrection</code>	(optional) Use the dipole correction for slab systems. The in-plane lattice must be perpendicular to the z-axis, and the third basis vector must be aligned with the z-axis. For charged calculation, this requests the generalized dipole correction [3], which may need some care for initializing the charge (see <code>charged</code> in the <code>initialGuess</code> group).
<code>zField</code>	(optional) Use an additional electric field along z when using the dipole correction (eV/bohr).

The PAWHamiltonian group may additionally contain the `vExt` and `xcMesh` groups to set an external potential and the mesh for xc calculation, respectively. It may also contain a `HubbardU` group for DFT+U calculations. Adding a `vdwCorrection` group enables van-der-Waals corrections.

Empty states The number of empty states is calculated relative to the minimum number of states needed to accommodate the electrons, i.e., the smallest integer above $n\text{Electrons}/2$. **Note:** Versions prior to 3.2 had set the base-line for spin-polarized systems to 5 extra states; this has changed now. If `nEmptyStates` is not provided, the following defaults apply (since version 3.2): If `ekt` is zero and the system is spin-compensated, semiconducting behaviour is assumed, and `nEmptyStates` defaults to zero. Otherwise, the number of empty states is set to $3 + n\text{Atoms}/2$ for spin-compensated systems, and $3 + 1.5 * n\text{Atoms}$ for spin-polarized systems.

If empty states are needed, the highest states should always have zero occupation. More states may help to improve convergence speed, notably for direct minimization (see CCG) or spin constraints.

Smearing schemes SPHInX supports Fermi-Dirac (the default), Gaussian, and Methfessel-Paxton smearing. Methfessel-Paxton smearing of order > 0 tends to yield lower electronic entropy at a comparable smearing width compared to Fermi-Dirac and Gaussian smearing and is therefore believed to produce free energies and forces closer to the $T=0\text{K}$ values. On the downside, it may produce occupation numbers outside the physical range $[0..1]$ and negative entropies. Note also that high orders introduce wiggles in the occupation function, which may couple to features in the DOS and produce artifacts.

Note that smearing parameters between different smearing schemes are *not* comparable on a 1:1 basis. Taking the full-width at half-maximum (FWHM) value of the underlying distribution as an approximate scaling factor, equivalent smearings require the following relative smearing parameters.

scheme	Fermi-Dirac	FD1	Gaussian(MP0)	MP1	MP2	MP3
FWHM [$k_B T$]	3.53	2.63	1.67	1.25	1.04	0.914
factor	1	1.35	2.11	2.82	3.39	3.9

MP n = Methfessel-Paxton order n , FD1 = Fermi-Dirac order 1.

Example: A Fermi-Dirac smearing of 0.1 eV is approximately equivalent to a Gaussian smearing of 0.21 eV.

Available xc functionals The xc parameter can be LDA_PW (10), i.e. the Perdew-Wang parametrization of LDA, or PBE (1). Perdew-Zunger LDA (0) works, but is not recommended because the parametrization is discontinuous, limiting the convergence in some cases. If the code is compiled with libxc support [4], you can omit the xc parameter and use libxc instead. libxc implements and names exchange functionals and correlation functionals separately. You must identify the libxc names (e.g. via their web site [5], or via the `sxlistxc` tool) and then set libxc to a string containing the two names with a '+' sign in between.

Example:

```
PAWHamiltonian {
  nEmptyStates = 10;
  ekt = 0.02;
  libxc = "gga_x_pbe+gga_c_pbe";
}
```

Meta-GGAs using the kinetic-energy density are implemented, but not well tested yet.

Hybrid functionals are still experimental, and slow. No libxc support yet, only self-implemented HSE and PBE0.

The following parameters allow experimenting with some details.

parameter	description
omegaHSE	(optional) Change the ω screening length of HSE
alphaHybrid	(optional) Change the non-local exchange mixing parameter of hybrid functionals.

3.4.1 The vExt group

The `vExt` group in the `PAWHamiltonian` or `PWHamiltonian` group defines an external potential.

Example:

```
PAWHamiltonian {  
  ...  
  vExt { file="sawtooth.sxb"; }  
}
```

It contains a single parameter, `file`, which contains the filename of the netcdf-type potential file to be used. The format of that file is like for a density (see Sec. 7.4).

3.4.2 The xcMesh group

The `xcMesh` group defines a specific FFT mesh for the calculation of the xc functional. It may be set with either of the following parameters (cf. the mesh definition in the `basis` group). The default is to double the density mesh, roughly equivalent to `meshAccuracy=2`.

The `xcMesh` group may appear in some groups other than the `Hamiltonian` to temporarily override the xc mesh: `scfDiag.CCG`, `blockCCG`, `CCG`.

parameter	description
<code>eCut</code>	Plane-wave cutoff for wavefunctions (in Rydberg). The xc mesh will be sufficient to hold $2\times$ this cutoff.
<code>mesh</code>	FFT mesh. 3-vector of integers.
<code>meshAccuracy</code>	FFT mesh accuracy, relative to the default. Default: 1, i.e. sufficient to hold $2\times$ the wavefunction plane-waves to avoid wrap-around errors

3.4.3 The vdWCorrection group

The `vdWCorrection` group in the `PAWHamiltonian` or `PWHamiltonian` specifies van-der-Waals corrections. At present, only the Grimme D2 implementation has been tested. The Tkatchenko-Scheffler implementation should be considered experimental.

Example:

```
PAWHamiltonian {  
  ...  
  vdWCorrection { method="D2"; }  
}
```

The following parameters may be set:

parameter	description
<code>method</code>	(required) The correction method; can be "D2" (Grimme D2) or "TS" (Tkatchenko-Scheffler).
<code>combinationRule</code>	(optional) "Tang" or "GB" (Gould-Bucko) for D2.

The default combination rule is due to Tang [6]:

$$C_6^{AB} = \left[\frac{1}{2} \left(\frac{\alpha^B}{C_6^B \alpha^A} + \frac{\alpha^A}{\alpha^B C_6^A} \right) \right]^{-1},$$

where C_6 are the C_6 coefficients and α is the polarizability. The alternative is Gould-Bucko [7]

$$C_6^{AB} = 1.43 \text{ Hartree bohr}^{-6} (\alpha^A \alpha^B \text{ bohr}^6)^{0.725}.$$

3.4.4 The HubbardU group

The HubbardU group defines on-site correlation corrections from the Hubbard model [8], commonly known as DFT+U. It uses the rotationally invariant formulation for symmetry equivalent orbitals.

Example:

```
PAWHamiltonian {
  ...
  HubbardU {
    site { ... } // 1st element with some U
    site { ... } // 2nd element with some U
    site { ... } // 3rd element with some U
  }
}
```

SPHInX offers four type of sites: deep atomic orbitals using the PAW projector defined via the `site` group, atomic orbitals inside the PAW radius of a specific l-channel as seen by the PAW projectors (also via the `site` group) atomic orbitals of a given radial shape via the `A0` group, or molecular orbitals of homonuclear diatomics via the `M0` group. They are described below.

The following parameters may be set:

parameter	description
<code>verbose</code>	Request more verbose output (and some test files).

3.4.5 The site group

The `site` group within the HubbardU group defines on-site correlation corrections using PAW projectors.

Example:

```

PAWHamiltonian {
  ...
  HubbardU {
    site { element="Fe"; U = 6; l=2}
  }
}

```

The following parameters may be set:

parameter	description
<code>element</code>	defines the element via its name
<code>species</code>	defines the element via its species number (1,2,3...) within the input file
<code>label</code>	defines the relevant atoms via their label. All atoms must belong to the same species. See also <code>label</code> in the <code>atom</code> group
<code>projectorType</code>	use a single projector from the PAW potential. Counting starts at 1. <i>l</i> -channel is chosen automatically.
<code>l</code>	use all projectors of a particular <i>l</i> -channel for projection. A radial truncation inside the PAW sphere is applied. Use <code>rPAW</code> in the <code>pawPot</code> group to change the radius.
<code>U</code>	The actual U value (in eV)
<code>shift</code>	An additional energy shift of the projector (in eV)

The species must be uniquely defined via `element`, `species`, or `label`.

If a single `projectorType` is chosen, it is assumed that the PAW projector projects a normalized valence AO to unity.

3.4.6 The AO group

The AO group within the HubbardU group defines on-site correlation corrections using AO orbital projectors. The projectors are constructed from atomic orbitals (AOs) of given radial shape. This shape is defined in the `orbital` group. The AO projectors include also the atomic PAW normalization.

Example:

```

PAWHamiltonian {
  ...
  HubbardU {
    AO {
      element="O";
      orbital { file="quamol-O.sxb "; is=1; iot=1; }
      rCut = 3;
      cutWidth=0.5;
      U = 0.5;

      nRadGrid=100;
    }
  }
}

```

The group must contain an `orbital` group.

The following parameters may be set:

parameter	description
<code>element</code>	defines the element via its name
<code>species</code>	defines the element via its species number (1,2,3...) within the input file
<code>label</code>	defines the relevant atoms via their label. All atoms must belong to the same species. See also <code>label</code> in the <code>atom</code> group
<code>nRadGrid</code>	number of radial points to represent atomic orbital projector. Default: 200
<code>rCut</code>	cutoff radius for atomic orbital (in bohr)
<code>cutWidth</code>	(optional) smoothening parameter for cutoff (in bohr). Default 0.7 bohr.
<code>U</code>	The actual U value (in eV)
<code>shift</code>	An additional energy shift of the projector (in eV)

The species must be uniquely defined via `element`, `species`, or `label`.

3.4.7 The MO group

The MO group within the HubbardU group defines on-site correlation corrections using MO orbital projectors. The molecular orbitals (MOs) are constructed from atomic orbitals (AOs) of given radial shape. This shape is defined in the `orbital` group. The MO projectors are constructed from AO projectors such that a normalized MO is projected to unity. The AO projectors include also the atomic PAW normalization.

Example:

```

PAWHamiltonian {
  ...
  HubbardU {
    MO {
      element="O";
      orbital { file="quamol-O.sxb "; is=1; iot=1; }
      mMO = 1;
      sign=-1;
      rCut = 3;
      cutWidth=0.5;
      U = 0.5;
      minDist = 2.0;
      maxDist = 3.5;

      nRadGrid=100;
      nInterpolate=10;
    }
  }
}

```

The group must contain an `orbital` group.

The following parameters may be set:

parameter	description
<code>element</code>	defines the element via its name
<code>species</code>	defines the element via its species number (1,2,3...) within the input file
<code>label</code>	defines the relevant atoms via their label. All atoms must belong to the same species. See also <code>label</code> in the atom group
<code>maxDist</code>	(required) maximum distance (in bohr) of two atoms to be considered a molecule
<code>minDist</code>	(optional) minimum distance (in bohr) of two atoms to be considered a molecule. Defaults to 50% of <code>maxDist</code> .
<code>nInterpolate</code>	(optional) number of distance points used to interpolate orbital normalization. Default: 100
<code>nRadGrid</code>	number of radial points to represent atomic orbital projector. Default: 200
<code>rCut</code>	cutoff radius for atomic orbital (in bohr)
<code>cutWidth</code>	(optional) smoothing parameter for cutoff (in bohr). Default 0.7 bohr.
<code>mMO</code>	rotational constant of orbital symmetry ($\sigma=0$, $\pi=1$)
<code>sign</code>	relative sign of orbitals on both atoms. Can be +1 or -1.
<code>U</code>	The actual U value (in eV)
<code>shift</code>	An additional energy shift of the projector (in eV)

The species must be uniquely defined via `element`, `species`, or `label`.

3.4.8 The orbital group

The `orbital` group within the `HubbardU { MO {} }` or `HubbardU { AO {} }` group defines where the radial shape of the atomic orbital is found. This is usual a `quamol` [9], an optimized orbital for the molecular species at hand.

Example:

```
PAWHamiltonian {
  ...
  HubbardU {
    MO {
      ...
      orbital { file="quamol-O.sxb "; is=1; iot=1; }
    }
  }
}
```

The following parameters may be set:

parameter	description
<code>fromPotential</code>	(flag) Get orbital shape from potential. This is not recommended.
<code>file</code>	(filename) Get orbital shape from this <code>quamol</code> -type <code>sxb</code> file.
<code>is</code>	(optional) species id within file (starts at 0). If not given, assumes same species ordering in <code>sxb</code> file as in input file.
<code>iot</code>	(required) Which orbital to take. Starts at 0.

3.5 The spinConstraint group

The optional `spinConstraint` group defines atomic spin constraints for PAW calculations. They can be set via atomic labels, or read in for all atoms from an external file.

Note: Atomic spins refer to the spin density inside the PAW sphere and are not identical to the definition used for setting up spin configurations from complete atoms in `initialGuess.atomicSpin`. The sphere radius can be changed via the `rPAW` parameter in the `pawPot` group.

Note: The direct minimization algorithm `CCG` cannot (yet) be used with atomic spin constraints.

Note: The speed of convergence improves if the plenty of empty states via the `nEmptyStates` are provided. This is because the Lagrangian parameters are optimized using a constrained search in the subspace of wavefunctions. As a rule of thumb, ensure that at least the entire valence shell for each atom is included. Example: for Fe (`[Ar]4s23d6`, 8 valence electrons) include empty orbitals for the

3d-shell (+2 states) and the 4p-shell (+3 states), i.e. 5 empty states per iron atom.

Example:

```
spinConstraint {
  label = "No.1 ";
  constraint = 2;
}
spinConstraint {
  label = "No.2 ";
  constraint = -2.5;
}
```

The following parameters may be set:

parameter	description
label	The present constraint applies to atoms with the given label.
constraint	Value of the desired atomic spin
file	Read all spin constraints from this file.

If the constraints are read from a file, there must be one number per atom. If not all atoms are constrained, the number is to be replaced by 'X', followed by a newline. For instance,

```
1
1
X
-2
```

for a four-atom calculation would constrain the spins of the first two atoms to a value of 1, the last one (no. 4) to a value of -2, and have no spin constraint on atom 3.

3.6 The initialGuess group

In order to start a DFT calculations, one must set up an initial guess for the density and for the wave functions. The `initialGuess` group defines how this is done, as well as a few other settings (such as keeping the waves on disk to save RAM).

Example:

```
initialGuess {
  waves { lcao {} }
  rho { atomicOrbitals; }
}
```

The default is to set up the density from a superposition of atomic densities, and the wave-functions from a single-step LCAO calculation, using the atomic

valence orbitals [10]. This works exceptionally well. If you want to finetune the behavior, the `initialGuess` group must contain a `waves` or a `rho` group. Otherwise, you may omit the `waves` and `rho` groups to get the default behavior.

Additionally, the `initialGuess` group may contain an occupations group to set up initial occupations (notably when keeping them fixed), and an `exchange` group for hybrid functionals.

The following parameters may be set:

parameter	description
<code>noWavesStorage</code>	(flag) Do not save wavefunctions after the initial guess
<code>noRhoStorage</code>	(flag) Do not save the density after the initial guess.

3.6.1 The waves group

The `waves` group defines the strategy for setting up the initial waves. The major strategies are LCAO (best), reading from a file (excellent, if you have one), or random. To use LCAO, you must include an `lcao` group.

The following parameters may be set:

parameter	description
<code>file</code>	(optional) File name for reading in the initial guess from a previous calculation.
<code>random</code>	(flag) Initialize with random numbers.
<code>keepWavesOnDisk</code>	(flag) Keep waves on disk, load only a single k -point at each time. May save a lot of RAM, but can be quite a bottleneck on small systems.

3.6.2 The lcao group

The `lcao` group within the `initialGuess.waves` group finetunes the LCAO calculation, if necessary. Notably, you can iterate the LCAO calculation to self-consistency. This is generally **no good idea**.

parameter	description
<code>maxSteps</code>	(optional) Max. number of steps. If 0 or 1, the initial density will <i>not</i> be updated.
<code>dEnergy</code>	(optional) Energy convergence criterium

In addition to these, the `lcao` group supports all the mixer settings (`mixingMethod`, `nPulaySteps`, `spinMixing`, `rhoMixing`, `preconditioner` group) of the `scfDiag` group (Sec. 4.1).

3.6.3 The rho group

The `rho` group defines the strategy for the initial density. This can be a superposition of atomic densities, from a file, from the wave functions (if read from a file), or random.

The following parameters may be set:

parameter	description
<code>file</code>	Read density from this file.
<code>fromWaves</code>	(flag) Compute from the wave functions (which must be from file in this case).
<code>random</code>	(flag) Request random density
<code>atomicOrbitals</code>	(flag) Superposition of atomic densities
<code>spinMoment</code>	(optional) When from atomic densities, apply a global spin polarization.

If `atomicOrbitals` is chosen, one may set up a spin configuration by using `atomicSpin` groups.

For charged calculations, one may specify an initial localization of charge in the charged group.

3.6.4 The `atomicSpin` group

The `atomicSpin` group (within the `initialGuess.rho` group) defines spin polarization for the initial guess on an per-atom basis when the initial density is set up from atoms. Post-polarizing an external density is not possible at present. Atoms can be identified per label (one `atomicSpin` group per label), or from an external file with spins for all atoms.

Example:

```
aLat = 5.35;
structure {
  include <structures/sc.sx>;
  species { element="Fe";
    atom { coords=[0,0,0]; relative; label="A"; }
    atom { coords=[1/2, 1/2, 1/2]; relative;
      label="B"; }
  }
}
initialGuess {
  rho {
    atomicOrbitals;
    atomicSpin { label="A"; spin=2; }
    atomicSpin { label="B"; spin=-2; }
  }
}
```

The following parameters may be set:

parameter	description
<code>spin</code>	The desired spin moment value.
<code>label</code>	For which atoms does this spin apply?
<code>file</code>	(excludes spin and label) Read atomic spins from this file (one spin per line), one per atom, in sequential order.

3.6.5 The aspherical group

The `aspherical` group within the `initialGuessrho` group allows to set up anisotropic charge densities on selected atoms (only PAW). Other atoms will be treated like in `atomicOrbitals`.

Example: See next page.

```

structure {
  cell = [[ 5.692374, 0.000000, 0.000000],
            [ 2.846187, 4.929740, 0.000000],
            [ 2.846187, 1.643247, 9.295609]];
  movable;
  species {
    element="Fe";
    atom {coords = [ 0.0, 0.0, 0.0]; relative;
           label = "up"; }
    atom {coords = [ 0.5, 0.5, 0.5]; relative;
           label = "down"; }
  }
  species {
    element="O";
    atom {coords = [ 0.25, 0.25, 0.25]; relative; }
    atom {coords = [ 0.75, 0.75, 0.75]; relative; }
  }
  symmetry {}
}

initialGuess {
  rho {
    aspherical {
      label="up";
      l=2;
      sphericalFocc = [5,0];
      orbital {
        focc=1;
        iSpin=1;
        eigenvector=[0,0,1,0,0];
      }
    }
    aspherical {
      label="down";
      l=2;
      sphericalFocc = [0,5];
      orbital {
        focc=1;
        iSpin=0;
        eigenvector=[0.8,0,0,-0.6,0];
      }
    }
  }
}

```

The following parameters may be set:

parameter	description
<code>label</code>	selection of atoms by label (single element!)
<code>element</code>	(excludes label) selection of atoms by element
<code>l</code>	(optional) select single l channel (must correspond to a unique, occupied valence orbital)
<code>sphericalFocc</code>	(optional) occupation of spherical part. Must provide occupations for all partial wave types (occupations up to $2L + 1$) or the active l channel as a vector. For spin-polarized, provide first all spin-up, then all spin-down occupations.

If the aspherical occupations go beyond a single l -channel (identified via $l=0,1,2,\dots$), no l channel should be given. The eigenvectors (see below) then must contain entries for all partial wave types.

For `spherical0cc`, one can either only specify occupations for the active l channel only, or occupations for all partial waves if electron occupations must be redistributed. Note that un-bound partial waves cannot be normalized in full space, but normalization is enforced within the definition range. It is strongly advised to not use unbound waves for initialization, i.e. set their occupations to zero. The order of partial waves is the one from the potential file. If `sphericalFocc` is not given, default values are assumed for the total occupation, i.e., any aspherical occupation weights are removed from the spherical ones.

The aspherical part is specified via the `orbital` groups, in which the following parameters are set.

parameter	description
<code>iSpin</code>	(optional) the desired spin channel (0 or 1). 0 if not provided.
<code>focc</code>	occupation in this orbital
<code>eigenvector</code>	Normalized eigenvector. If l is provided, refers to that subspace. Otherwise, this uses the full partial wave basis; coefficients for non-occupied partial waves should be set to zero to avoid normalization issues.

The total occupation matrix is set up from the aspherical and the spherical part. If `sphericalFocc` is not provided, it is assumed that the aspherical occupation is redistributed, i.e., the spherical occupations are chosen such that the total aspherical occupations sum up to the default.

`aspherical` cannot be combined with `atomicSpin`. The actual atomic spin can be set up by choosing the spin-dependent occupation accordingly.

3.6.6 The charged group

The `charged` group (within the `initialGuess.rho` group) defines the localization of initial charge. This may be a Gaussian charge, or a sheet-layer charge

along z (for charged slabs) with a Gaussian profile.

Example:

```
initialGuess {
  rho {
    atomicOrbitals;
    charged { charge=1; coords=[0,0,13.2]; beta=2; }
  }
}
```

The following parameters may be set:

parameter	description
charge	(required) The classical charge (i.e. -nExcessElectrons from the PAWHamiltonian or PWHamiltonian group).
beta	(optional) Gaussian broadening
z	Request a sheet charge at this z
coords	Request a Gaussian charge at this position (in bohr).

Multiple charges can be set up. There is no check if the resulting density has the correct number of electrons.

3.6.7 The occupations group

The `occupations` group within the `initialGuess` group defines the initial occupations. This makes sense if the density is computed from wave functions, or if the occupations are going to be fixed at these values.

Example:

```
initialGuess {
  ...
  occupations {
    bands { range=[1,10]; focc=2; }
    kPoints { values=[1,3];
      bands { values=[10]; focc=0; }
      bands { values=[11]; focc=2; }
    }
  }
}
```

One can specify selected sets of states by nesting the following groups: `kPoints`, `spin`, `bands`. Each of these groups selects one or more indices (starting at 1) from the specified index type (i.e. `k`-points, spins, states) by one of the following two parameters:

parameter	description
values	(list of integers) Specifically list the indices affected.
range	(list of 2 integers) Specify start and end index.

The possible nestings are

- `kPoints { spin { bands { ... focc=...; } } }`
- `kPoints { bands { ... focc=...; } }`
- `spin { bands { ... focc= ...; } }`
- `bands { ... focc=...; }`

The innermost loop must set the `focc` parameter to the desired occupation. If an outer group is omitted, the settings apply to all indices of that type (e.g. no `kPoints` group = all `k`-points). The groups are processed in order of appearance, i.e., later groups override the previous settings, if the index ranges overlap.

In the above example, the first `bands` group sets the occupations of bands 1–10 at all `k`-points to 2. Then, we change the occupation at `k`-points 1 and 3, to have zero occupation in band 10 and an occupation of 2 for band 11.

3.6.8 The exchange group

Note: hybrid functionals are experimental and slow.

The `exchange` group allows to set waves for the non-local exchange operator at the initialization stage. This is necessary if you want to initialize the waves from an LCAO calculation. The exchange group contains a single parameter, `file`, which contains the filename of the waves file to be used.

Example:

```
initialGuess {
  exchange { file="waves-pbe.sxb"; }
  rho { file="rho-pbe.sxb"; }
  waves { lcao {} }
}
```

3.7 The pseudoPot group

The `pseudoPot` group defines the norm-conserving pseudopotentials by a sequence of species groups. The order of species must agree with the structure group.

Note: PAW and norm-conserving pseudopotentials cannot be mixed. Using `pseudoPot` requires to use `PWHamiltonian` to define the Hamiltonian.

Example:

```

pseudoPot {
  species {
    potential      = "al.cpi";
    name           = "Aluminum";
    element        = "Al";
    valenceCharge  = 3;
    lMax           = d;
    lLoc           = d;
    lcaoOrbitals   = [s,p,d];
    atomicRhoOcc   = [2,1,0];
    rGauss         = 1;
    reciprocalMass = 26.98;
    dampingMass    = 0.7;
    ionicMass      = 26.98;
  }
}

```

The group must contain one `species` group per species. Within the `species` group, the following parameters may be set:

parameter	description
<code>name</code>	(required string) English name of the element
<code>element</code>	(optional string) Chemical symbol
<code>potential</code>	(required filename) Name of the potential file.
<code>valenceCharge</code>	(required) Valence charge.
<code>lMax</code>	(required) Max. l-component of the pseudopotential.
<code>lLoc</code>	(required) Select the local l-component.
<code>lcaoOrbitals</code>	(required) Which orbitals should be used for lcao initialization. Note: <code>s,p,d</code> , and <code>f</code> are predefined constants in <code>parameters.sx</code> .
<code>rGauss</code>	(required) Broadening of compensation charge, usually 1.
<code>atomicRhoOcc</code>	(required) Occupation numbers for charge initialization.
<code>reciprocalMass</code>	(required) Mass of the ion.
<code>dampingMass</code>	(required, currently unused). Damping for damped-Newton geometry optimization.
<code>ionicMass</code>	(required) Mass of the ion.

Note: The appearance of two masses, as well as the lack of default parameters is a historic feature (mistake?) and may be changed in future versions.

3.8 The PWHamiltonian group

The PWHamiltonian group defines the DFT functional, the number of empty states, the smearing, and some other settings.

Example:

```

PWHamiltonian {
  nEmptyStates = 10;
  ekt = 0.02;
  xc = LDA;
}

```

The following common parameters may be set:

parameter	description
xc	(required) xc functional to be used. Constants are defined in <code>parameters.sx</code> . See below.
ekt	(optional) smearing in eV. Should be 0 for semiconductors.
MethfesselPaxton	(optional) If ≥ 0 , use Methfessel-Paxton smearing of indicated order. Order 0 is same as Gaussian smearing.
FermiDirac	(optional) If ≥ 0 , use FermiDirac smearing of indicated order. Order 0 is the default; order 1 corresponds to first-order corrections. Higher orders are not yet implemented.
nEmptyStates	(optional) number of empty states. Remember to scale with system size!
nExcessElectrons	(optional) Number of extra electrons (= minus charge). Can be fractional.
spinPolarized	(flag) Run a collinear spin-polarized calculation.
dipoleCorrection	(optional) Use the dipole correction for slab systems. The in-plane lattice must be perpendicular to the z-axis, and the third basis vector must be aligned with the z-axis. For charged calculation, this requests the generalized dipole correction, which may need some care for initializing the charge (see <code>charged</code> in the <code>initialGuess</code> group).
zField	(optional) Use an additional electric field along z when using the dipole correction (eV/bohr).

For empty states, please see notes on empty states in the `PAWHamiltonian` section on page 15.

For smearing, please see notes on smearing schemes in the `PAWHamiltonian` section on page 16.

Available xc functionals The `xc` parameter can be `LDA_PW` (10), i.e. the Perdew-Wang parametrization of LDA, or `PBE` (1). Perdew-Zunger LDA (0) works, but is not recommended because the parametrization is discontinuous, limiting the convergence in some cases. If you need other functionals, contact `freysoldt@mpi-susmat.de`.

The `PWHamiltonian` group may additionally contain the `vExt` and `xcMesh` groups to set an external potential and the mesh for xc calculation, respectively. Adding a `vdwCorrection` group enables van-der-Waals corrections.

3.9 The main group

The `main` group selects and controls the algorithms after startup during the main part of the calculation. The can be divided into three categories: electronic minimization (Sec. 4), geometry optimization (Sec. 5), and additional computations (Sec. 6) The main group can contain a single algorithm, or a sequence of them, which are then executed in order of appearance.

All self-consistent DFT calculations must employ some algorithm to compute the electronic wavefunctions and density for a given geometric structure. If additionally the atomic positions are altered according to the calculated forces, electronic minimization becomes the inner loop. Consequently, the electronic minimizers appear in the `bornOppenheimer` group of the geometry optimizers. Electronic minimizers may appear as a direct subgroup of `main`. It may sometimes be useful to select a different set of electronic minimization settings for the initial phase, if the initial guess is not good enough, and then use other settings (less robust, more efficient) later on.

Example:

```
main {
  // — get a better starting guess with
  //    small rhoMixing
  scfDiag {
    dEnergy=1e-4;
    rhoMixing=0.3;
    maxSteps=100;
  }
  // — optimize geometry
  linQN {
    dEnergy = 1e-6;
    maxSteps=50;
    bornOppenheimer {
      scfDiag {
        dEnergy = 1e-8;
        rhoMixing=1;
      }
    }
  }
}
```

The group must contain one or more electronic minimizer (`scfDiag`, `CCG`) or geometry optimizer (`QN`, `linQN`, `ricQN`, `ricTS`) group.

The main group has no parameters of its own.

4 Electronic loop

4.1 scfDiag: iterative diagonalization + density mixing

The `scfDiag` group selects and controls the iterative diagonalization + density mixing algorithm for the solution of the Kohn-Sham DFT equations.

Example:

```
scfDiag {  
  blockCCG { blockSize=32; maxStepsCCG=4; }  
  dEnergy = 1e-7;  
  preconditioner { type=KERKER; scaling=1;  
                  kerkerDamping=1; }  
  mixingMethod = PULAY;  
  nPulaySteps = 7;  
}
```

The group must contain one of the iterative diagonalization groups: `CCG` (conjugate-gradient³), `blockCCG` (block conjugate-gradient, recommended). It may contain a `xcMesh` group (to override the xc mesh) and a `preconditioner` group (for the density). If the density preconditioner is not specified, Kerker preconditioning with a default damping is used.

The following parameters may be set:

³The first *C* stands for *Complex*.

parameter	description
<code>dEnergy</code>	(optional) Free energy convergence criterium (in Hartree). Default is 10^{-8} .
<code>maxSteps</code> <code>maxResidue</code>	(optional) Max. number of steps (density updates). (optional) Additional requirement for convergence: density residue must fall below this threshold.
<code>printSteps</code>	(optional) Eigenvalues are printed by default every 10 steps. This interval can be changed by <code>printSteps</code> .
<code>mixingMethod</code>	(optional) Method for the density mixing. Constants defined in <code>parameters.sx</code> . Can be PULAY (default, 2) or LINEAR (0).
<code>nPulaySteps</code>	(optional) Number of previous steps (densities) to use in Pulay mixing. Default: 7.
<code>rhoMixing</code>	(optional) Additional linear mixing factor for density updates (1=full update (default), 0=no change). Low values may lead to a more stable convergence, but will slow down the calculation if set too low.
<code>spinMixing</code>	(optional) Linear mixing parameter for spin densities.
<code>keepRhoFixed</code>	(flag) Do not update the density (for band structures).
<code>keepOccFixed</code>	(flag) Do not update the occupation numbers.
<code>keepSpinFixed</code>	(flag) Do not change the spin moment.
<code>spinMoment</code>	(optional) Keep the spin moment at this value.
<code>ekt</code>	(optional) Override electronic temperature setting in the Hamiltonian group.
<code>dipoleCorrection</code>	(flag) Override the dipole correction setting in the Hamiltonian group.
<code>dSpinMoment</code>	accuracy of iterative enforcement of spin constraints. Default: 10^{-8} .
<code>noRhoStorage</code>	(flag) Do not write <code>rho.sxb</code> .
<code>noWavesStorage</code>	(flag) Do not write <code>waves.sxb</code> .

Multiple iterative diagonalizer groups may be used to change the xc mesh during the minimization. For this, each diagonalizer groups must contain the `dEnergy` parameter. **Example:**

```

scfDiag {
  blockCCG { xcMesh { meshAccuracy=0.7; }
             dEnergy=1e-4; }
  blockCCG { xcMesh { meshAccuracy=1; }
             dEnergy=1e-8; }
  blockCCG { xcMesh { meshAccuracy=2; }
             dEnergy=1e-9; }
  dEnergy=1e-9;
}

```

This example would use an xc-mesh with a relative accuracy of 0.7 until the energy has converged to 10^{-4} Hartree, a standard $2\times$ mesh (accuracy=1) until

the energy change is below 10^{-8} Hartree, and a double-accuracy mesh thereafter. The following parameters allow messing around with some details.

parameter	description
<code>dRelRes</code>	(optional) The energy convergence criterium for eigenvalues is calculated from the density residue adaptively. This parameter allows to set the initial (and maximum) value of the scaling factor.
<code>calcForces</code>	(flag) Calculate and print forces in each step (slow!). This can be used to determine the influence of convergence criteria on the precision in forces.
<code>dumpTime</code>	(optional) Set interval between dumps of density/wave functions. The dumps may allow for restarts when a calculation crashes. Since the dumping happens at the end of a SCF cycle, the actual dumping intervals may be longer. Default: 12 hours.

4.1.1 The CCG group

The CCG group (within the `scfDiag` group) selects conjugate-gradient algorithm for (inner-loop) iterative diagonalization. After all states have been updated, a subspace diagonalization is performed. This algorithm works best for very small systems. For larger systems (> 5 states), the `blockCCG` is superior.

Note: CCG is also the name for the conjugate-gradient-based direct minimization algorithm, see Sec. 4.2.

The following parameters may be set:

parameter	description
<code>dRelEps</code>	(optional) Stop iterating when the change in eigenvalue falls below this fraction of the change in the first (steepest-descent) step.
<code>maxStepsCCG</code>	(optional) Max. number of steps to perform. Default: 5.
<code>dEnergy</code>	(optional) Use these settings until energy change fall below this threshold.

The CCG group may contain a `xcMesh` group to override the xc mesh settings from the Hamiltonian or the `scfDiag` group.

4.1.2 The blockCCG group

The `blockCCG` group (within the `scfDiag` group) selects the block conjugate-gradient algorithm for (inner-loop) iterative diagonalization. After all states have been updated, a subspace diagonalization is performed. For very small systems (≤ 5 states), the unblocked CCG is superior.

The following parameters may be set:

parameter	description
<code>dRelEps</code>	(optional) Stop iterating when the change in eigenvalue falls below this fraction of the change in the first (steepest-descent) step.
<code>maxStepsCCG</code>	(optional) Max. number of steps to perform. Default: 5.
<code>blockSize</code>	(optional) Block size. Default: 64
<code>nSloppy</code>	(optional) Don't try to converge the highest nSloppy states (useful for bandstructures).
<code>dEnergy</code>	(optional) Use these settings until energy change fall below this threshold.
<code>verbose</code>	(flag) Produce debug output.
<code>numericalLimit</code>	(flag) Stop iterating when approaching the numerical limit.

The `blockCCG` group may contain a `xcMesh` group to override the keyxc mesh settings from the Hamiltonian or the `scfDiag` group.

4.1.3 The preconditioner group

The `preconditioner` group defines the density preconditioner, i.e., a transformation of the observed (or predicted) difference between the input and output density to the applied changes to the input density. An ideal preconditioner models the screening behavior of the system and is able to include the expected screening response into the suggested density change. Selecting an appropriate preconditioner, that reflects the screening properties of the system at hand, is a key to an efficient (i.e. fast) convergence. The preconditioner does *not* affect the converged result.

The following parameters may be set:

parameter	description
<code>type</code>	(required) Type of preconditioner, see below
<code>scaling</code>	(optional) Additional scaling factor. Default 1.
<code>spinScaling</code>	(optional) Additional scaling factor for the spin density. Default 1.
<code>kerkerDamping</code>	(optional) Damping constant for Kerker preconditioner. Default: 0.6
<code>dielecConstant</code>	(optional) Dielectric constant for the CSR model.

The following preconditioner types are available (constants defined in `parameters.sx`).

- NONE (0). No preconditioning. Ideal for atoms/molecules in vacuum.
- KERKER (1). Kerker preconditioner. Ideal for metals.
- CSR (3). Preconditioner for semiconductors based on the Cappellini-del-Sole-Reining-Bechstedt model dielectric function. Requires `dielecConstant`.
- ELLIPTIC (5). An explicit-solver preconditioner. No screening in vacuum region, Thomas-Fermi screening (Kerker-like) elsewhere. Ideal for metallic

slabs.

4.2 CCG: direct minimization

The CCG group selects and controls the direct minimization algorithm for the solution of the Kohn-Sham DFT equations [11].

Example:

```
CCG {  
  dEnergy=1e-7;  
  finalDiag;  
}
```

It may contain a `xcMesh` group to override the xc mesh settings from the Hamiltonian group.

Note: The direct minimization algorithm cannot be used with atomic spin constraints (set by `spinConstraint`).

The following parameters may be set:

parameter	description
<code>dEnergy</code>	(optional) Free energy convergence criterium (in Hartree). Default is 10^{-8} .
<code>maxSteps</code>	(optional) Max. number of steps.
<code>printSteps</code>	(optional) Eigenvalues are printed by default every 10 steps. This interval can be changed by <code>printSteps</code> .
<code>initialDiag</code>	(flag) Perform iterative wave-function optimization based on the initial density (this is the default)
<code>finalDiag</code>	(flag) Perform subspace diagonalization at the end.
<code>kappa</code>	(optional) Initial mixing between subspace Hamiltonian and wave-function updates. If set to a negative value, the value of κ will be fixed at the absolute value. Otherwise, κ is adapted on the fly.
<code>keepOccFixed</code>	(flag) Do not update the occupation numbers.
<code>ekt</code>	(optional) Override electronic temperature setting in the Hamiltonian group.
<code>dipoleCorrection</code>	(flag) Override the dipole correction setting in the Hamiltonian group.
<code>noRhoStorage</code>	(flag) Do not write <code>rho.sxb</code> .
<code>noWavesStorage</code>	(flag) Do not write <code>waves.sxb</code>

Note about switching between `scfDiag` and CCG: The direct minimization algorithm requires a good initial guess for the wave functions. Therefore, the default is to run a iterative diagonalization (blockCGG algorithm) before starting direct minimization. This is unnecessary when the wavefunctions come from iterative diagonalization, and can be switched off via the `initialDiag` flag. For semiconductors with no empty states, on the other hand, the wavefunctions spanning the occupied subspace may not be diagonalizing the Hamiltonian. In

order to switch to `scfDiag` or to obtain eigenvalues, the `finalDiag` flag must be set. If partially occupied or empty states are computed, an approximate diagonalization takes place as part of the algorithm.

5 Geometry optimizers

This section describes the geometry optimization groups. We have experimented with geometry optimization and found BFGS quasi-Newton schemes to be most efficient. Therefore, you will see three quasi-Newton optimizers. For small systems, the performance is quite similar. For more complex systems, the `ricQN` variant is recommended.

For transition states, there is an experimental quasi-Newton optimizer for 1st-order saddle points, `ricTS`.

In addition to the geometry optimizers, there is a pseudo-optimizer `extControl`. This algorithm opens two communication channels (named pipes) and allows to control geometry updates (and other things) from external scripts.

In all cases, the inner electronic loop enjoys the same flexibility of sequencing as the `main` group. Therefore, each geometry optimizer group contains a `bornOppenheimer` group that contains one or more electronic minimizers.

5.1 The QN group

The QN group selects and controls the geometry optimization via quasi-Newton scheme with BFGS updates. **Note:** In general, `ricQN` is the faster algorithm.

Example:

```
QN {
  maxSteps = 20;
  dEnergy = 1e-5;
  dF = 1e-3;
  maxStepLength=0.2;
  bornOppenheimer {
    scfDiag {
      maxSteps = 50;
      blockCCG { }
      dEnergy = 1e-7;
    }
  }
}
```

The group must contain a `bornOppenheimer` group to specify the electronic loop.

The following parameters may be set:

parameter	description
<code>maxSteps</code>	(optional) max. number of steps, default: 50
<code>dX</code>	(optional) convergence reached only when maximum displacement (length of displacement vector of a single atom) is less than this value (in bohr). Default: 0.01
<code>dF</code>	(optional) convergence reached only when maximum force (length of force vector of a single atom) is less than this value (in Hartree/bohr). Default: 0.001
<code>dEnergy</code>	(optional) convergence reached only when change in energy is less than this value (in Hartree). Default: 10^{-4}
<code>maxStepLength</code>	maximum allowed displacement (length of displacement vector for a single atom) in bohr. Larger steps are reduced by scaling. Default: 0.3
<code>hessian</code>	(filename) Initialize Hessian from file.
<code>driftFilter</code>	(flag) Project out the average force and displacement. Default: yes, if no constraints are used.

5.1.1 The `bornOppenheimer` group

The `bornOppenheimer` group defines the electronic loop within a geometry optimization. It contains one or more of the electronic loop groups, see Sec. 4. If more than one minimizer is used, the complete electronic loop sequence is executed at each ionic step.

5.2 The `linQN` group

The `linQN` group selects and controls the geometry optimization via quasi-Newton scheme with BFGS updates for the inverse Hessian. **Note:** In general, `ricQN` is the faster algorithm.

Example:

```

linQN {
  maxSteps = 20;
  dEnergy = 1e-5;
  dF = 1e-3;
  maxStepLength=0.2;
  bornOppenheimer {
    scfDiag {
      maxSteps = 50;
      blockCCG { }
      dEnergy = 1e-7;
    }
  }
}

```

The group must contain a `bornOppenheimer` group to specify the electronic loop.

The following parameters may be set:

parameter	description
<code>maxSteps</code>	(optional) max. number of steps, default: 50
<code>dX</code>	(optional) convergence reached only when maximum displacement (length of displacement vector of a single atom) is less than this value (in bohr). Default: 0.01
<code>dF</code>	(optional) convergence reached only when maximum force (length of force vector of a single atom) is less than this value (in Hartree/bohr). Default: 0.001
<code>dEnergy</code>	(optional) convergence reached only when change in energy is less than this value (in Hartree). Default: 10^{-4}
<code>maxStepLength</code>	maximum allowed displacement (length of displacement vector for a single atom) in bohr. Larger steps are reduced by scaling. Default: 0.3
<code>nProjectors</code>	(optional) number of previous steps to use for BFGS updates. Default: 10
<code>hessian</code>	(filename) Initialize Hessian from file.
<code>driftFilter</code>	(flag) Project out the average force and displacement. Default: yes, if no constraints are used.

5.3 The `ricQN` group

The `ricQN` group requests a quasi-Newton optimization with BFGS updates [12] of an on-the-fly optimized internal-coordinate based initial guess for the Hessian.

The following parameters may be set:

parameter	description
<code>maxSteps</code>	(optional) max. number of steps, default: 50
<code>dX</code>	(optional) convergence reached only when maximum displacement (length of displacement vector of a single atom) is less than this value (in bohr). Default: 0.01
<code>dF</code>	(optional) convergence reached only when maximum force (length of force vector of a single atom) is less than this value (in Hartree/bohr). Default: 0.001
<code>dEnergy</code>	(optional) convergence reached only when change in energy is less than this value (in Hartree). Default: 10^{-4}
<code>nProjectors</code>	(optional) number of previous steps to use for BFGS updates. Default: 10
<code>maxStepLength</code>	maximum allowed displacement (length of displacement vector for a single atom) in bohr. Larger steps are reduced by the trust radius method to a value close to the maximum (within 1%). Default: 0.3
<code>softModeDamping</code>	(optional) Initial value for Hessian shift (in Hartree/bohr ²). This is overridden with the first successful fit of a positive shift parameter. Default: 1e-2.
<code>driftFilter</code>	(flag) Project out the average force and displacement. Default: yes, if no constraints are used.

The group must contain a `bornOppenheimer` group to specify the electronic loop. The `ricQN` group may contain a `ric` group (see Sec. 5.3.1 to define the internal coordinate generation. If left out, default parameters are used, and output from the internal coordinate setup is suppressed.

5.3.1 The `ric` group

The `ric` group defines the parameters for internal coordinate generation.

The following parameters may be set:

parameter	description
<code>maxDist</code>	maximum possible distance for considering neighbors (in bohr). Default: 10
<code>typifyThreshold</code>	minimum bond length separation of distinct bond types (the <i>f</i> parameter in [12]). After sorting the bond lengths, the logarithm of subsequent lengths are compared. If they differ by less than the threshold, the two bonds are assigned the same bond type. Default: 0.05
<code>rmsThreshold</code>	minimum distance between two bond length clusters in units of their root-mean-square displacements (the <i>R</i> parameter of [12]). Default: 3
<code>planeCutLimit</code>	Relative size of coordination polyhedra to separate the nearest neighbors from further atoms (the <i>P</i> parameter of [12]). Larger values allow for more neighbors. Default: 0.95
<code>withAngles</code> <code>bvkAtoms</code>	(flag) add bond angle coordinates for all bonds (optional, experimental) List of atom ids (starting from 1) for which born-von-Karman transversal force constants are added. The comma-separated list must be enclosed by square brackets []. This adds a bond-directional coordinate to each bond of the atoms in the list.

5.4 The ricTS group

The `ricTS` group requests a quasi-Newton optimization for 1st-order saddle points (transition states) using updates [13] of an on-the-fly optimized internal-coordinate based initial guess for the Hessian [12]. An initial guess for the reaction coordinate must be known.

Note: This is an experimental feature. The optimization should be started within the saddle point region (one negative eigenvalue of the Hesse matrix), otherwise, the algorithm may converge to a different stationary point (a minimum, or a higher-order saddle point).

Example:

```

ricTS {
  maxSteps = 20;
  dEnergy = 1e-5;
  dF = 1e-3;
  maxStepLength=0.2;
  // for a 4-atomic structure
  // = 12 coordinates (x1, y1, z1, x2, y2, ...)
  transDir=[0, 0, -1, 0, 0, 0,
            0, 0, 1, 0, 0, 0];
  bornOppenheimer {
    scfDiag {
      maxSteps = 50;
      blockCCG { }
      dEnergy = 1e-7;
    }
  }
}

```

The following parameters may be set:

parameter	description
<code>maxSteps</code>	(optional) max. number of steps, default: 50
<code>dX</code>	(optional) convergence reached only when maximum displacement (length of displacement vector of a single atom) is less than this value (in bohr). Default: 0.01
<code>dF</code>	(optional) convergence reached only when maximum force (length of force vector of a single atom) is less than this value (in Hartree/bohr). Default: 0.001
<code>dEnergy</code>	(optional) convergence reached only when change in energy is less than this value (in Hartree). Default: 10^{-4}
<code>nProjectors</code>	(optional) number of previous steps to use for BFGS updates. Default: 10
<code>maxStepLength</code>	maximum allowed displacement (length of displacement vector for a single atom) in bohr. Larger steps are reduced by the trust radius method to a value close to the maximum (within 1%). Default: 0.3
<code>transCurvature</code>	(optional) Initial curvature along transition direction, must be negative. Default: -0.1
<code>anyStationaryPoint</code>	(flag) Disable all mechanisms to discover a 1st-order transition state. Converge to a nearby stationary point.
<code>maxDirRot</code>	(optional) Control updates of transition direction. Parameter between 0 (no updates) and 1 (full updates). Default: 0.5.
<code>scheme</code>	(optional). Hesse update scheme (0=Murtagh-Sargent, 1=Powell symmetric Broyden, 2=Bofill, 3=Farkas-Schlegel, see [13], Eqs. 6-10). Default: Powell symmetric Broyden
<code>driftFilter</code>	(flag) Project out the average force and displacement. Default: yes, if no constraints are used.

The group must contain a `bornOppenheimer` group to specify the electronic loop. The `ricQN` group may contain a `ric` group (see Sec. 5.3.1 to define the internal coordinate generation. If left out, default parameters are used, and output from the internal coordinate setup is suppressed.

The initial direction of the transition can be specified either by giving a $3N_{\text{atom}}$ -dimensional displacement vector `transDir`, see example above. Alternative, the displacement pattern of selected atoms can be described by one or more `transPath` groups.

Note: The transition direction will be filtered (with respect to drift, movable constraints) and normalized.

5.4.1 The `transPath` group

The `transPath` group defines initial guess for the atomic displacement pattern along the negative curvature (transition direction).

The following parameters may be set:

parameter	description
<code>atomId</code>	Index of a single atom, $1 \dots N_{\text{atoms}}$.
<code>atomIds</code>	List (vector) of atom indices ($1 \dots N_{\text{atoms}}$)
<code>dir</code>	Vector of displacements

Either `atomId` or `atomIds` must be set. `dir` contains the x, y, z displacements for each atom (3 values for `atomId` or 3 values for each index in `atomIds`).

5.5 The `extControl` group

The `extControl` group allows to control parts of the SPHInX run by external scripts. For this, two communication channels are opened. Typically, these will be named pipes. For more info on the concept, see the `sxextopt` manual. The `extControl` group is the ‘DFT code’ side of the atomic structure algorithm protocol (ASAP).

The names of the communication files are specified via the environment, in `SX_EXT_CTRL` and `SX_EXT_RES`, respectively.

The group must contain one or more `bornOppenheimer` groups. Each `bornOppenheimer` group may contain an `id` parameter than allows to run the specified sequence of electronic minimizers (usually one) with the `run` command (see below).

If forces are not needed, the `noForces` flag can be used to suppress the calculation of forces.

The `extControl` algorithm understands the following commands:

command	description
confirm y	Confirm <i>ASAP</i> commands
confirm n	Do not confirm <i>ASAP</i> commands
end	End the calculation
run [<id>]	Run specified bornOppenheimer group (first one, if <id> is omitted).
onproblem=crash	Crash on wrong commands
onproblem=stop	Stop on wrong commands
onproblem=ignore	Ignore wrong commands (default)
get energy	Print free energy on result channel
get forces	Print forces on result channel
get natoms	Print number of atoms on result channel
get nspecies	Print number of species, and for each species number of atoms on result channel
get structure	Print atomic coordinates on result channel
get positions	Print atomic coordinates + chemical symbol on result channel
get cell	Print unit cell on result channel
set structure	Set atomic coordinates
set positions	Set atomic coordinates
shift atom	<id> <dx> <dy> <dz> shift specified atom. id starts at 1 and enumerates atoms across all species.
set cell	<i>not supported</i>
get stress	<i>not supported</i>
get elements	Print chemical symbols on result channel
get nspinconstraints	Print number of spin constraints on result channel
get nu	Print spin constraint Langrange parameters on result channel
set spinconstraint	Set target spins for spin constraints
get spinconstraint	Print target spins for spin constraints on result channel
get atom spin	Print atomic spins on result channel
enable spinconstraint	Switch on spin constraint
disable spinconstraint	Switch off spin constraint

6 Extra computations (in main group)

6.1 evalForces: force evaluation

The `evalForces` group is used to calculate forces and write them to a file in `sx-format`. This is useful for single-point calculations without a structure optimization. It should be used after an electronic loop.

Example:

```
CCG {  
  dEnergy=1e-7;  
}  
evalForces { file="forces.sx"; }
```

The following parameters may be set:

parameter	description
file	(optional) Filename.

If the group is encountered multiple times (or multiple groups write to the same file), the additional force output will be appended. The output format is in `relaxHist.sx` format.

6.2 pDos: PAW projected DOS evaluation

The `pDos` group is used to calculate the projected density of states from the PAW projections. The weight is the (PAW-reconstructed) wavefunction (squared) within different `l`-channels inside the PAW cutoff sphere. **Other codes use different spheres, so quantitative PDOS comparisons across different projection techniques make no sense.** It should be used after an electronic loop. If wave functions are written, there is a post-processing add-on, `sxpdos`, performing the same.

Example:

```
main {  
  CCG {  
    dEnergy=1e-7;  
    finalDiag;  
  }  
  pDos {  
    broadening=0.2;  
  }  
}
```

The following parameters may be set:

parameter	description
<code>file</code>	(optional) Filename base, defaults to "pdos".
<code>broadening</code>	(optional) Gaussian broadening in eV; default is 0.1 eV.
<code>range</code>	(optional) define energy range [min,max] for DOS.
<code>sumOnly</code>	Include all atoms, but don't write per-atom files.

By default, all atoms are selected. This can be changed with *one* of the following

parameter	description
<code>atoms</code>	List of atom indices, starting from 1.
<code>elements</code>	List of chemical element names.

or an `atomRange` group, that contains a `from` and `to` parameter. All atom indices start from 1.

Atom selection examples:

```
pDos {
  atomRange { from = 1; to = 10; }
}
pDos {
  elements = ["Fe", "O"];
}
```

The output is as follows. `fileAtomid.dat` lists the pdos sum for all selected atoms. Here, *id*, is the total atom index, starting from zero (**but input starts at 1**). *file* is the filename base set by the `file` parameter. Each row is schematically

```
<energy> <sum all channels> <l=0 (s)> <l=1 (p)> <l=2 (d)> ...
```

`filesum.dat` lists the pdos sum for all selected atoms, per species. Each row is schematically

```
<energy> <sum 1st species> <l=0 (s)> <l=1 (p)> <l=2 (d)> ...
<sum 2nd species> <l=0 (s)> <l=1 (p)> ...
```

For spin-polarized calculations, the file endings are `.up` and `.down` instead of `.dat`.

7 Output

7.1 energy.dat: total energies

This file lists the total energies in each electronic step. The format for `scfDiag` is (7 columns in one long line)

```
<step> <accumulated time> <total energy> <free energy>
  <T=0 energy> <band energy> <entropy>
```

Units are Hartree, except for dimension-less entropy S^{el} .

The entropy arises from the partial occupations (Fermi-Dirac distribution or Methfessel-Paxton) used for metals. The connection between the different energies (E for energy, F for free energy) is

$$\begin{aligned} F(T) &= E(T) - k_{\text{B}}T S^{\text{el}}(T) \\ E(T=0) &\approx \frac{1}{2}[E(T) + F(T)] + \mathcal{O}(T^3) \end{aligned}$$

Here, $k_{\text{B}}T$ is the electronic temperature parameter `ekt` parameter set in the Hamiltonian.

Note: Only the free energy is variational, and forces correspond to the derivatives of the free energy. Therefore, the Born-Oppenheimer surface at finite electronic temperature corresponds to the electronic free energy.

The format for CCG is

```
<step> <accumulated time> <total energy> <free energy>
```

Units are Hartree.

7.2 residue.dat: density residues

This file contains the density residue for each electronic step. The format is

```
<step> <residue>
```

For spin-polarized calculations, the third column lists the residue of the spin density.

7.3 eps.dat: eigenvalues

This file contains the computed eigenvalues. The format is

```
<ik> <eps_1> <eps_2> <eps_3> ...
```

Units are eV.

For spin-polarized calculations, the two spin channels have distinct eigenvalues contained separately in `eps.0.dat` and `eps.1.dat`, respectively.

7.4 rho.sxb: density (binary)

This file contains the final density in netcdf-format. The mesh order is z running fastest. For details on how to interpret the variables, please contact me.

7.5 waves.sxb: wave functions (binary)

This file contains the final wave functions in netcdf-format. It's pointless to try to process this file outside the SPHInX library. Several add-ons are available for processing the wavefunctions.

7.6 relaxedStr.sx: final atomic structure

This file lists the final atomic positions in SPHInX format. If the optimization does not converge, it contains the last configuration calculated.

7.7 relaxHist.sx: geometry optimization history

This file lists for each calculated configuration the atomic positions and forces, in SPHInX format.

Each configuration corresponds to a structure format, with an additional item `force`, which lists the force vector in atomic units (Hartree/bohr).

Example:

```
structure {
  cell = [[ 10.0000,  0.0000,  0.0000],
          [  0.0000, 10.0000,  0.0000],
          [  0.0000,  0.0000, 10.0000]];
  species {
    element="N";
    atom {coords = [ 5.00000,  5.00000,  5.00000];
          force  = [ 0.0027634, 0.0027634, 0.0027634]; }
  }
  species {
    element="H";
    atom {coords = [ 5.99349,  3.62616,  3.62616]; movable;
          force  = [-0.0191398, 0.0081741, 0.0081741]; }
    atom {coords = [ 3.62616,  5.99349,  3.62616]; movable;
          force  = [ 0.0081741, -0.0191398, 0.0081741]; }
    atom {coords = [ 3.62616,  3.62616,  5.99349]; movable;
          force  = [ 0.0081741, 0.0081741, -0.0191398]; }
  }
}
```

7.8 energy-structOpt.dat: geometry optimization energies

This file lists for each step the (electronic free) energy. The format is one step per line:

```
<it> <energy (in Hartree)>
```

The 0-th step is the initial structure.

7.9 fftwisdom.dat: FFTW plans

FFTW wisdom file. May be used to speed up FFTW planning in add-ons.

7.10 hubbardOcc.dat: DFT+U occupations

Occupations used for DFT+U occupations. For each site (with N_{orb} correlated orbitals), we list all N_{orb} eigenvalues and the corresponding eigenvectors of the occupation matrix. For atomic orbitals, $N_{\text{orb}} = 2l + 1$. For spin-polarized calculations, spin-up and spin-down occupations are given after each other. Different sites are separated by empty lines. The format for each line is

```
<eigenvalue> <c1> <c2> <c3> ... <cN>
```

Eigenvalues have units of electron number.

Site ordering for atoms is according to their ordering in the structure. If you specify multiple +U corrections (more than one `site`, `A0`, or `M0` group inside the `HubbardU` group), sites of different types come in blocks. For ordering of molecular sites, see the log file.

References

- [1] S. Boeck, C. Freysoldt, A. Dick, L. Ismer, and J. Neugebauer, *Comp. Phys. Commun.* **182**, 543 (2011), URL <https://doi.org/10.1016/j.cpc.2010.09.016>.
- [2] URL <https://esl.cecami.org/data/paw-xml/>.
- [3] C. Freysoldt, A. Mishra, M. Ashton, and J. Neugebauer, *Physical Review B* **102**, 045403 (2020), URL <https://dx.doi.org/10.1103/physrevb.102.045403>.
- [4] S. Lehtola, C. Steigemann, M. J. Oliveira, and M. A. Marques, *SoftwareX* **7**, 1–5 (2018), ISSN 2352-7110, URL <http://dx.doi.org/10.1016/j.softx.2017.11.002>.
- [5] Libxc web site, URL <https://libxc.gitlab.io/functionals/>.
- [6] K. T. Tang, *Physical Review* **177**, 108 (1969), URL <https://dx.doi.org/10.1103/physrev.177.108>.
- [7] T. Gould and T. Bučko, *J. Chem. Theory Computation* **12**, 36033 (2016), URL <https://dx.doi.org/10.1021/acs.jctc.6b00361>.
- [8] M. Cococcioni and S. de Gironcoli, *Physical Review B* **71**, 035105 (2005), URL <https://dx.doi.org/10.1103/physrevb.71.035105>.
- [9] B. Lange, C. Freysoldt, and J. Neugebauer, *Physical Review B* **84**, 085101 (2011), URL <https://doi.org/10.1103/physrevb.84.085101>.
- [10] J. Neugebauer and C. G. Van de Walle, in *Materials Theory, Simulations, and Parallel Algorithms, MRS Symposia Proceedings*, edited by E. Kaxiras, J. Joannopoulos, P. Vashisha, and R. K. Kalia (MRS, Pittsburgh, 1996), vol. 408.

- [11] C. Freysoldt, S. Boeck, and J. Neugebauer, *Phys. Rev. B* **79**, 241103(R) (2009), URL <https://doi.org/10.1103/PhysRevB.79.241103>.
- [12] C. Freysoldt, *Comp. Mat. Sci.* **133**, 71 (2017), URL <https://dx.doi.org/10.1016/j.commatsci.2017.03.001>.
- [13] B. Schlegel, *WIREs Comput. Mol. Sci.* **1**, 790 (2011), URL <https://dx.doi.org/10.1002/wcms.34>.

Index

- alphaHybrid, 16
- angular grid, 14
- angularGrid, 14
- anyStationaryPoint, 45
- A0, 18, 22, 52
 - definition*, 19
- ASAP, 46
- aspherical
 - definition*, 26
- atom, 9, 19–21
 - definition*, 9
- atomic positions, 7–9
- atomic structure, 7
- atomic structure algorithm protocol, 46
- atomicOrbitals, 25, 26
- atomicRhoOcc, 31
- atomicSpin, 22, 25
 - definition*, 25
- atomId, 46
- atomIds, 46
- atomRange
 - definition*, 49
- atoms, 49

- band structure path, 11
- basis, 7, 17
 - definition*, 10
- beta, 29
- blockCCG, 17, 34, 36
 - definition*, 36
- blockSize, 37
- bornOppenheimer, 33, 39, 41, 42, 45, 46
 - definition*, 40
- broadening, 49
- bvkAtoms, 43

- calcForces, 36
- CCG, 15, 17, 22, 33, 34, 36, 50
 - definition*, 36, 38
- charge, 15, 28, 32
- charge, 29
- charged, 15, 25, 32
 - definition*, 28
- checkOverlap, 14
- chemical element, 13, 31
- combinationRule, 18
- comment, 6
- conjugate-gradient, 36
- constraint, 23
- convergence, 40–42, 45
 - density residue, 35
 - energy, 24, 35, 38
- coords, 9, 11, 12, 29
- CSRB, 37
- cutoff, 11
- cutWidth, 20, 21

- dampingMass, 31
- dEnergy, 24, 35–38, 40–42, 45
- density preconditioner, 37
- dF, 40–42, 45
- DFT+U, 15
- diagonalization, 34, 36
- dielecConstant, 37
- dipole correction, 15, 32, 35
- dipoleCorrection, 15, 32, 35, 38
- dir, 46
- direct minimization, 38
- dK, 12
- dRelEps, 36, 37
- dRelRes, 36
- driftFilter, 40–42, 45
- dSpinMoment, 35
- dumping interval, 36
- dumpTime, 36
- dX, 40–42, 45

- eCut, 11, 17
- eigenvector, 28
- ekt, 15, 32, 35, 38, 50
- electric field, 15
- electronic minimizer, 33
- electronic temperature, 35, 38
- element, 13, 31
- element, 9, 13, 19–21, 28, 31

elements, 49
 empty states, 15, 32
 energy convergence, 24, 35, 38
 energy-structOpt.dat, 51
 energy.dat, 49
 eps.0.dat, 50
 eps.1.dat, 50
 eps.dat, 50
 evalForces
 definition, 48
 exchange, 24
 definition, 30
 extControl, 39
 definition, 46
 external potential, 17

 Fermi-Dirac distribution, 50
 FermiDirac, 15, 32
 FFT mesh, 10, 11, 17
 xc computation, 17
 fftwisdom.dat, 51
 file, 17, 22–25, 30, 48, 49
 finalDiag, 38, 39
 flag, 5
 focc, 28, 30
 folding, 11
 force, 51
 forces, 48
 dumping, 48
 format, 6
 from, 11, 12, 49
 definition, 12
 fromPotential, 22
 fromWaves, 25
 functional, 14, 31

 gCut, 11
 geometry
 input, 7
 optimization, 39

 hessian, 40, 41
 hubbard0cc.dat, 52
 HubbardU, 15, 18–20, 22, 52
 definition, 18

 id, 46

 initial guess, 23
 initialDiag, 38
 initialGuess, 7, 15, 22, 24–26, 28, 29,
 32
 definition, 23
 input file, 3
 comment, 6
 expression, 6
 flag, 5
 format, 5
 internal coordinate, 42
 ionicMass, 31
 iot, 22
 is, 22
 iSpin, 28

 kappa, 38
 keepOccFixed, 35, 38
 keepRhoFixed, 35
 keepSpinFixed, 35
 keepWavesOnDisk, 24
 Kerker, 37
 kerkerDamping, 37
 kPoint, 10
 definition, 11
 kPoints, 10, 12
 definition, 11

 l, 19, 28
 label, 9, 12, 19–21, 23, 25, 28
 LCAO, 23, 24, 30
 lcao, 24
 definition, 24
 lcaoOrbitals, 31
 LDA, 16, 32
 libxc, 15, 16
 linQN, 33
 definition, 40
 lLoc, 31
 lMax, 31
 lMaxRho, 14

 main, 7, 39
 definition, 33
 maxDirRot, 45
 maxDist, 21, 43

- maxResidue, 35
- maxStepLength, 40–42, 45
- maxSteps, 24, 35, 38, 40–42, 45
- maxStepsCCG, 36, 37
- memory
 - save, 11, 23
- mesh, 11, 17
- meshAccuracy, 11, 17
- Methfessel-Paxton, 50
- Methfessel-Paxton smearing, 16
- MethfesselPaxton, 15, 32
- method, 18
- minDist, 21
- mixer, 24
- mixing, 35
- mixingMethod, 35
- mMO, 21
- MO, 18, 22, 52
 - definition*, 20
- Monkhorst-Pack, 10, 11
- movable, 8, 9
- movableLine, 9
- movableX, 8, 9
- movableY, 8, 9
- movableZ, 8, 9
- MPI, 3
- MPI parallelism, 3
- mpirun, 3

- name, 13, 31
- nEmptyStates, 15, 22, 32
- nExcessElectrons, 15, 32
- nInterpolate, 21
- noForces, 46
- noRhoStorage, 24, 35, 38
- norm-conserving pseudopotentials, 30
- noWavesStorage, 24, 35, 38
- nPoints, 12
- nProjectors, 41, 42, 45
- nPulaySteps, 35
- nRadGrid, 14, 20, 21
- nSloppy, 37
- numericalLimit, 37

- occupation, 29
 - fixed, 35, 38

- occupations
 - definition*, 29
- omegaHSE, 16
- openMP, 3
- openMP parallelism, 3
- operator, 10
- orbital, 28
- orbital, 19–21
 - definition*, 22

- parallelism, 3
- parameter.sx
 - location, 6
- parameters.sx, 6, 15, 31, 32, 35, 37
- PAW potentials, 13
- PAW sphere, 14, 19, 22
- PAWHamiltonian, 7, 17, 29, 32
 - definition*, 14
- pawPot, 7, 19, 22
 - definition*, 13
- PBE, 16, 32
- pDos
 - definition*, 48
- plane-wave basis, 10
- plane-wave cutoff, 11
- planeCutLimit, 43
- potential
 - norm-conserving, 30
- potential, 13, 31
- potType, 13
- preconditioner, 37
- preconditioner, 24, 34
 - definition*, 37
- printSteps, 35, 38
- projected density of states, 48
- projectorType, 19
- pseudoPot, 7
 - definition*, 30
- Pulay mixing, 35
- PWHamiltonian, 7, 17, 29, 30
 - definition*, 31

- QN, 33
 - definition*, 39

- radial grid, 14

RAM
 save, 11, 23
random, 24, 25
range, 29, 49
range limits, 6
rCut, 20, 21
reciprocalMass, 31
relative, 9, 11, 12
relaxedStr.sx, 51
relaxHist.sx, 48, 51
residue.dat, 50
rGauss, 31
rho, 24–26, 28
 definition, 24
rho.sxb, 35, 38, 50
rhoMixing, 35
ric, 42, 45
 definition, 42
ricQN, 33, 39, 40
 definition, 41
ricTS, 33, 39
 definition, 43
rmsThreshold, 43
rPAW, 14, 19, 22

S, 10
saddle points, 43
saveMemory, 11
scaling, 37
scfDiag, 17, 24, 33, 36–39, 49
 definition, 34
scheme, 45
shift, 19–21
sign, 21
single-point calculations, 48
site, 18, 52
 definition, 18
slab, 15, 32
smearing, 15, 16, 32
softModeDamping, 42
species, 8, 13, 19–21, 31
 definition, 8
sphericalFocc, 28
spin, 15
 fixed, 35
spin, 25
 spin configuration, 25
 spin constraints, 22, 35
 spin polarization, 25
 spin-polarized, 15, 32
 spinConstraint, 38
 definition, 22
 spinMixing, 35
 spinMoment, 25, 35
 spinPolarized, 15, 32
 spinScaling, 37
 std, 6
 steps, 35, 38
 structure, 7, 13
 definition, 7
 subspace diagonalization, 36, 38
 sumOnly, 49
 SX_EXT_CTRL, 46
 SX_EXT_RES, 46
 SX_THREADS, 4
 symmetry, 8
 definition, 9
 symmetry group, 10

threads, 3
to, 11, 12, 49
 definition, 12
transCurvature, 45
transDir, 45
transition states, 43
transPath, 45
 definition, 45
type, 37
typifyThreshold, 43

U, 19–21
units, 5

valenceCharge, 31
values, 29
vdwCorrection, 15, 32
 definition, 17
verbose, 18, 37
vExt, 15, 32
 definition, 17

waves, 24

definition, 24
waves.sxb, 35, 38, 50
weight, 11
withAngles, 43

xc, 15, 32
xc functional, 15, 32
xc mesh, 34, 36, 38
xcMesh, 15, 32, 34, 36–38
definition, 17

z, 29
zField, 15, 32