

AiiDA tutorial

The AiiDA Team

May–June 2018

Contents

1 Preliminaries	1
2 Verdi command line	2
3 Verdi shell and AiiDA objects	8
4 Submit, monitor and debug calculations	12
5 Queries in AiiDA: The QueryBuilder	17
6 AiiDA Workflows: workfunctions, workchains	17
7 WorkChains, a real-world example: computing a band structure for a simple crystal structure	28
Appendices	28
A Calculation input validation	28
B Restarting calculations	31
C Queries in AiiDA - Optional examples and exercises	32
D More complex logic in workflows: while loops and conditional statements	34

1 Preliminaries

Before starting

Once connected to your machine, type in the remote terminal

```
workon aiiida
```

This will enable the virtual environment in which AiiDA is installed, allowing you to use AiiDA.

Now type in the same terminal

```
jupyter notebook --no-browser
```

This will run a server with a web application called `jupyter`, which is used to create interactive python notebooks. To connect to this application, copy the URL that has been printed to the terminal (it will be something like `http://localhost:8888/?token=2a3ba37cd1...`) and paste it into the URL bar of a web browser. You will see a list of folders: these are folders on the remote Amazon computer. We will use `jupyter` in section 5 and optionally in other sections as well.

Now launch an identical `ssh` connection (again, as explained above) in another terminal, and type `workon aiiida` here too. This terminal is the one you will actually use in this tutorial.

A final note: for details on AiiDA that may not be fully explained here, you can refer to the full AiiDA documentation, available online at <http://aiida-core.readthedocs.io/en/latest/>.

Troubleshooting tips (in case you have issues later)

- If you get an error like `ImportError: No module named aiida` or `No command 'verdi' found` double check that you have loaded the virtual environment with `workon aiida` before launching python, ipython or the jupyter server.
- If your browser cannot connect to the jupyter instance, check that you have correctly configured SSH tunneling/forwarding as described above. Also note that you should run the jupyter server from the terminal connected to the Amazon machine, while the web browser should be opened locally on your laptop or workstation.
- The Jupyter Notebook officially supports the latest stable versions of Chrome, Safari and Firefox. See <http://jupyter-notebook.readthedocs.io/en/4.x/notebook.html#browser-compatibility> for more information on browser compatibility (and update your browser if it is too old).

2 Using the verdi command line

This part of the tutorial will help to familiarize you with the command-line utility `verdi`, one of the most common ways to interact with AiiDA. `verdi` with its subcommands enables a variety of operations such as inspecting the status of ongoing or terminated calculations, showing the details of calculations, computers, codes, or data structures, access the input and the output of a calculation, etc. Similar to the `bash` shell, `verdi` command support Tab completion. Try right now to type `verdi` in a terminal and tap Tab twice to have a list of subcommands. Whenever you need the explanation of a command type `verdi help` or add `-h` flag if you are using any of the `verdi` subcommands. Finally, fields enclosed in angular brackets, such as `<pk>`, are placeholders to be replaced by the actual value of that field (an integer, a string, etc...).

2.1 The list of calculations

Let us try our first `verdi` commands. Type in the terminal

```
verdi calculation list
```

(Note: the first time you run this command, it might take a few seconds as it is the first time you are accessing the database in the virtual machine. Subsequent calls will be faster). This will print the list of ongoing calculations, which should be empty. The first output line should look like

PK	Creation	State	Type	Computer	Job state
----	-----	-----	-----	-----	-----

Total results: 0

Info: last time an entry changed state: never

In order to print a list with all calculations that finished correctly in the AiiDA database, you can use the `-s/--states` flag as follows:

```
verdi calculation list --states FINISHED
```

Another very typical option combination allows to get calculations in *any* state (flag `-a`) generated in the past `NUM` days (`-p <NUM>`): e.g., for calculation in the past 1 day: `verdi calculation list -p1 -a`. Since you have not yet run any calculations at the virtual machine that you currently use and all the existing calculations were imported and belong to a different user, you can type (flag `-A` shows the calculations of all the users):

```
verdi calculation list -A --states IMPORTED
```

Each row of the output identifies a calculation and shows some information about it. For a more detailed list of properties, choose one row by noting down its PK (primary key) number (first column of the output) and type in the terminal

```
verdi calculation show <pk>
```

The output depends on the specific pk chosen and should inform you about the input nodes (e.g. pseudopotentials, kpoints, initial structure, etc.), and output nodes (e.g. output structure, output parameters, etc.).

PKs/IDs vs. UUIDs: Beside the (integer) PK, very convenient to reference a calculation or data node in your database, every node has a UUID (Universally Unique ID) to identify it, that is preserved even when you share some nodes with coworkers—while the PK will most likely change. You can see the UUID in the output of `verdi calculation show` or `verdi node show`. Moreover, if you have already a UUID and you want to get the corresponding PK in your database, you can use `verdi node show -u <UUID>`, as we are going to do now.

Let us now consider the node with UUID = `ce81c420-7751-48f6-af8e-eb7c6a30cec3`, which identifies a relaxation of a BaTiO₃ unit cell run with Quantum Espresso pw.x. You can check the information on this node and get the PK with:

```
$ verdi node show -u ce81c420-7751-48f6-af8e-eb7c6a30cec3
```

Property	Value
-----	-----
type	PwCalculation
pk	4235
uuid	ce81c420-7751-48f6-af8e-eb7c6a30cec3
label	
description	
ctime	2014-10-27 17:51:21.781045+00:00
mtime	2018-05-16 11:19:39.848446+00:00
process state	
finish status	
computer	[2] daint
code	pw-SVN-piz-daint

Inputs	PK	Type
-----	----	-----
parameters	4236	ParameterData
kpoints	4526	KpointsData
pseudo_Ba	966	UpfData
pseudo_Ti	4315	UpfData
settings	4529	ParameterData
pseudo_0	4342	UpfData
structure	436	StructureData

Outputs	PK	Type
-----	----	-----
output_kpoints	3665	KpointsData
output_parameters	3670	ParameterData
output_structure	3666	StructureData
retrieved	3668	FolderData
output_trajectory_array	265	ArrayData
remote_folder	1977	RemoteData

Keep in mind that you can also use just a part (beginning) of the UUID, as long as it is unique, to show the node information. For example, to display the above information, you could also type `verdi node show -u ce81c420`. In what follows, we are going to mention only the prefixes of the UUIDs since they are sufficient to identify the correct node.

2.2 A typical AiiDA graph

AiiDA stores inputs required by a calculation as well as the its outputs in the database. These objects are connected in a graph that looks like Fig. 1. We suggest that you have a look to the figure before going ahead.

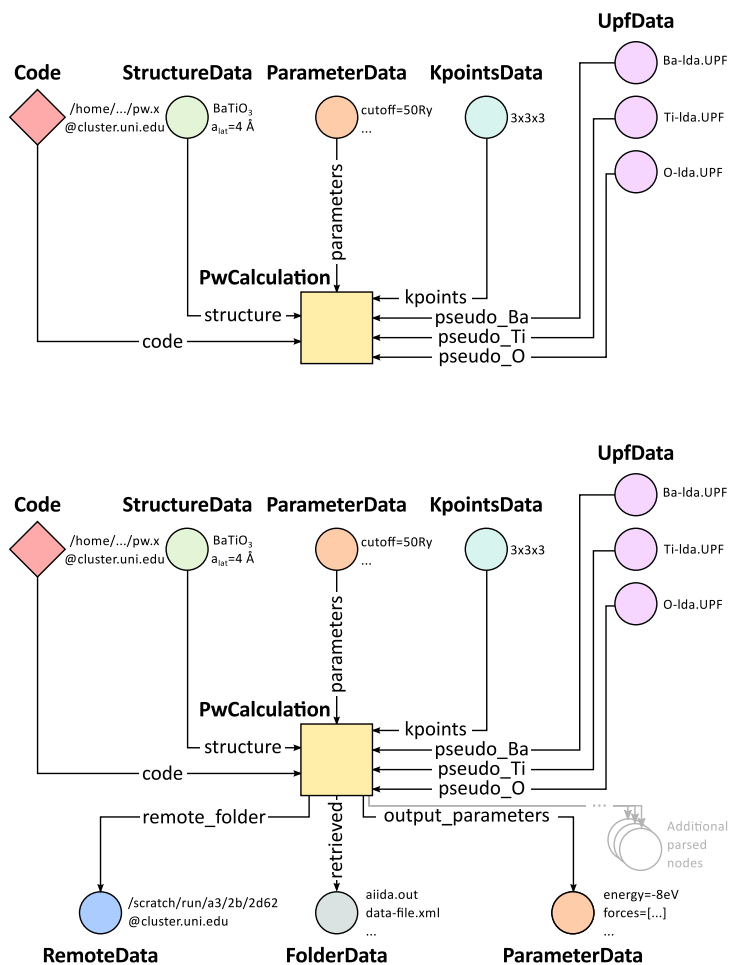


Figure 1 – (a, top) Graph with all inputs (data, circles; and code, diamond) to the Quantum Espresso calculation (square) that you will create in Sec. 4 of this tutorial. **(b, bottom)** Same as (a), but also with the outputs that the daemon will create and connect automatically. The **RemoteData** node is created during submission and can be thought as a symbolic link to the remote folder in which the calculation runs on the cluster. The other nodes are created when the calculation has finished, after retrieval and parsing. The node with linkname “retrieved” contains the raw output files stored in the AiiDA repository; all other nodes are added by the parser. Additional nodes (symbolized in gray) can be added by the parser (e.g., an output **StructureData** if you performed a relaxation calculation, a **TrajectoryData** for molecular dynamics, ...).

You can create a similar graph for any calculation node by using the utility `verdi graph generate <pk>`. For example, before you obtained information (in text form) for `UUID = ce81c420`. To visualize similar information in graph(ical) form, run (replacing `<pk>` with the PK of the node):

```
verdi graph generate <pk>
```

This command will create the file `<pk>.dot` that can be rendered by means of the utility `dot`. If you now type

```
dot -Tpdf -o <pk>.pdf <pk>.dot
```

you will create a pdf file `<pk>.pdf`. You can open this file on the Amazon machine by using `evince` or, if you feel that the ssh connection is too slow, copy it via `scp` to your local machine. To do so, if you are using Linux/Mac OS X, you can type in your *local* machine:

```
scp aiiadatutorial:<path_with_the_graph_pdf> <local_folder>
```

and then open the file. Alternatively, you can use graphical software to achieve the same, for instance: WinSCP on Windows, Cyberduck on the Mac, or the “Connect to server” option in the main menu after clicking on the desktop for Ubuntu.

Spend some time to familiarize yourself with the graph structure. Choose the root node (highlighted in blue) and trace back the parent calculation which produced the structure used as an input. This is an example of a Quantum ESPRESSO `pw.x` calculation, where the input structure was actually obtained as the output of a previous calculation. We will now inspect the different elements of this graph.

2.3 Inspecting the nodes of a graph

ParameterData and Calculations

Now, let us have a closer look at some of the nodes appearing in the graph. Choose the node of the type `ParameterData` with input link name `parameters` (to double check, it should have `UUID d1bbe1ea`) and type in the terminal:

```
verdi data parameter show <pk>
```

A `ParameterData` contains a dictionary (i.e., key–value pairs), stored in the database in a format ready to be queried (we will learn how to run queries later on in this tutorial). The command above will print the content dictionary, containing the parameters used to define the input file for the calculation. You can compare the dictionary with the content of the raw input file to Quantum ESPRESSO (that was generated by AiiDA) via the command

```
verdi calculation inputcat <pk>
```

where you substitute the pk of the calculation node. Check the consistency of the parameters written in the input file and those stored in the `ParameterData` node. Even if you don’t know the meaning of the input flags of a Quantum ESPRESSO calculation, you should be able to see how the input dictionary has been converted to Fortran namelists.

The previous command just printed the content of the “default” input file `aiida.in`. To see a list of all the files used to run a calculation (input file, submission script, etc.) instead type

```
verdi calculation inputls <pk>
```

(Adding a `--color` flag allows you to easily distinguish files from folders by a different coloring).

Once you know the name of the file you want to visualize, you can call the `verdi calculation inputcat` command specifying the path. For instance, to see the submission script, you can do:

```
verdi calculation inputcat <pk> -p _aiidasubmit.sh
```

StructureData

Now let us focus on `StructureData` objects, representing a crystal structure. We can consider for instance the input structure to the calculation we were considering before (it should have UUID prefix `3a4b1270`). Such objects can be inspected interactively by means of an atomic viewer such as the one provided by `ase`. AiiDA however supports several other viewers such as `xcrysden`, `jmol`, and `vmd`. Type in the terminal

```
verdi data structure show --format ase <pk>
```

to show the selected structure (it will take a few seconds to appear, and you can rotate the view with the right mouse button—if you receive some errors, make sure you started your SSH connection with the `-X` or `-Y` flag).

Alternatively, especially if showing them interactively is too slow over SSH, you can export the content of a structure node in various popular formats such as `xyz` or `xsf`. This is achieved by typing in the terminal

```
verdi data structure export --format xsf <pk> > <pk>.xsf
```

You can open the generated `xsf` file and observe the cell and the coordinates. Then, you can then copy `<pk>.xsf` from the Amazon machine to your local one and then visualize it, e.g. with `xcrysden` (if you have it installed):

```
xcrysden --xsf <pk>.xsf
```

Codes and computers

Let us focus now on the nodes of type `code`. A code represents (in the database) the actual executable used to run the calculation. Find the pk of such a node in the graph and type

```
verdi code show <pk>
```

The command prints information on the plugin used to interface the code to AiiDA, the remote machine on which the code is executed, the path of its executable, etc. To show a list of all available codes type

```
verdi code list
```

If you want to show all codes, including hidden ones and those created by other users, use `verdi code list -a -A`. Now, among the entries of the output you should also find the code just shown.

Similarly, the list of computers on which AiiDA can submit calculations is accessible by means of the command

```
verdi computer list -a
```

(`-a` shows all computers, also the one imported in your database but that you did not configure, i.e., to which you don't have access). Details about each computer can be obtained by the command

```
verdi computer show <COMPUTERNAME>
```

Now you have the tools to answer the question:

What is the scheduler installed on the computer where the calculations of the graph have run?

Calculation results

The results of a calculation can be accessed directly from the calculation node. Type in the terminal

```
verdi calculation res <pk>
```

which will print the output dictionary of the “scalar” results parsed by AiiDA at the end of the calculation. Note that this is actually a shortcut for

```
verdi data parameter show <pk2>
```

where `pk2` refers to the `ParameterData` node attached as an output of the calculation node, with link name `output_parameters`.

By looking at the output of the command, what is the Fermi energy of the calculation with UUID prefix `ce81c420`?

Similarly to what you did for the calculation inputs, you can access the output files via the commands

```
verdi calculation outputls <pk>
```

and

```
verdi calculation outputcat <pk>
```

Use the latter to verify that the Fermi energy that you have found in the last step has been extracted correctly from the output file (Hint: filter the lines containing the string “Fermi”, e.g. using `grep`, to isolate the relevant lines).

The results of calculations are stored in two ways: `ParameterData` objects are stored in the database, which makes querying them very convenient, whereas `ArrayData` objects are stored on the disk. Once more, use the command `verdi data array show <pk>` to know the Fermi energy obtained from calculation with UUID prefix `ce81c420` (you need to use, this time, the PK of the output `ArrayData` of the calculation, with link name `output_trajectory_array`). As you might have realized the difference now is that the whole series of values of the Fermi energy calculated after each relax/vc-relax step are stored. (The choice of what to store in `ParameterData` and `ArrayData` nodes is made by the parser of `pw.x` implemented in the `aiida-quantumespresso` plugin.)

(Optional section) Comments

AiiDA offers the possibility to attach comments to a calculation node, in order to be able to remember more easily its details. Node with UUID prefix `ce81c420` has no comment already defined, but you can add a very instructive one by typing in the terminal

```
verdi comment add -c "vc-relax of a BaTiO3 done with QE pw.x" <pk>
```

Now, if you ask for a list of all comments associated to that calculation by typing

```
verdi comment show <pk>
```

the comment that you just added will appear together with some useful information such as its creator and creation date. We let you play with the other options of `verdi comment` command to learn how to update or remove comments.

2.4 AiiDA groups of calculations

In AiiDA, calculations (and more generally nodes) can be organized in groups, which are particularly useful to assign a set of calculations or data to a common project. This allows you to have quick access to a whole set of calculations with no need for tedious browsing of the database or writing complex scripts for retrieving the desired nodes. Type in the terminal

```
verdi group list
```

to show a list of the groups that already exist in the database. Choose the PK of the group named `tutorial_pbesol` and look at the calculations that it contains by typing

```
verdi group show <pk>
```

In this case, we have used the name of the group to organize calculations according to the pseudopotential that has been used to perform them. Among the rows printed by the last command you will be able to find the calculation we have been inspecting until now.

If, instead, you want to know all the groups to which a specific node belongs, you can run

```
verdi group list --node <pk>
```

3 Using the verdi shell and familiarizing with AiiDA objects

In this section we will use an interactive ipython environment with all the basic AiiDA classes already loaded. We propose two realizations of such a tool. The first consist of a special ipython shell where all the AiiDA classes, methods and functions are accessible. Type in the terminal

```
verdi shell
```

For all the everyday AiiDA-based operations, i.e. creating, querying and using AiiDA objects, the `verdi shell` is probably the best tool. In this case, we suggest that you use two terminals, one for the `verdi shell` and one to execute bash commands.

The second option is based on `jupyter` notebooks and is probably most suitable to the purposes of our tutorial. Go to the browser where you have opened `jupyter` and click `New` → `Python 2` (top right corner). This will open an ipython notebook based on cells where you can type portions of python code. The code will not be executed until you press `Shift+Enter` from within a cell. Type in the first cell

```
%aiida
```

and execute it. This will set exactly the same environment as the `verdi shell`. The notebook will be automatically saved upon any modification and when you think you are done, you can export your notebook in many formats by going to `File` → `Download as`. We suggest you to have a look to the drop-down menus `Insert` and `Cell` where you will find the main commands to manage the cells of your notebook. **The verdi shell and the jupyter notebook are completely equivalent. Use either according to your personal preference.**

Note: you will still need sometimes to type command-line instructions in `bash` in the first terminal you opened today. To differentiate these from the commands to be typed in the `verdi shell`, the latter will be marked in this document by a vertical line on the left, like:

```
| some verdi shell command
```

while command-line instructions in `bash` to be typed on a terminal will be encapsulated between horizontal lines:

```
some bash command
```

Alternatively, to avoid changing terminal, you can execute `bash` commands within the `verdi shell` or the notebook adding an exclamation mark before the command itself

```
| !some bash command
```

3.1 Loading a node

Most AiiDA objects are represented by nodes, identified in the database by its pk number (an integer). You can access a node using the following command in the shell:

```
| node = load_node(PK)
```

Load a node using one of the calculation pks visible in the graph you displayed in the previous section of the tutorial. Then get the energy of the calculation with the command

```
| node.res.energy
```

You can also type

```
| node.res.
```

and then press `TAB` to see all the possible output results of the calculation.

3.2 Loading different kinds of nodes

3.2.1 Pseudopotentials

From the graph displayed in Section 2.2, find the pk of the barium pseudopotential file (LDA). Load it and verify that it describes barium. Type

```
| upf = load_node(PK)
| upf.element
```

All methods of `UpfData` are accessible by typing `upf.` and then pressing `TAB`.

3.2.2 k-points

A set of k-points in the Brillouin zone is represented by an instance of the `KpointsData` class. Choose one from the graph of Section 2.2, load it as `kpoints` and inspect its content:

```
| kpoints.get_kpoints_mesh()
```

Then get the full (explicit) list of k-points belonging to this mesh using

```
| kpoints.get_kpoints_mesh(print_list=True)
```

If you incurred in a `AttributeError`, it means that the `kpoints` instance does not represent a regular mesh but rather a list of k-points defined by their crystal coordinates (typically used when plotting a band structure). In this case, get the list of k-points coordinates using

```
| kpoints.get_kpoints()
```

If you prefer Cartesian (rather than crystal) coordinates, type

```
| kpoints.get_kpoints(cartesian=True)
```

For later use in this tutorial, let us try now to create a `kpoints` instance, to describe a regular $2 \times 2 \times 2$ mesh of k-points, centered at the Gamma point (i.e. without offset). This can be done with the following commands:

```
| from aiida.orm.data.array.kpoints import KpointsData
| kpoints = KpointsData()
| kpoints_mesh = 2
| kpoints.set_kpoints_mesh([kpoints_mesh,kpoints_mesh,kpoints_mesh])
| kpoints.store()
```

The import performed in the first line is however unpractical as it requires to remember the exact location of the module containing the `KpointsData` class. Instead, it is easier to use the `DataFactory` function instead of an explicit import.

```
| KpointsData = DataFactory("array.kpoints")
```

This function loads the appropriate class defined in a string (here `array.kpoints`).¹ Therefore, `KpointsData` is not a class instance, but the `kpoints` class itself!

3.2.3 Parameters

Nested dictionaries with individual parameters, as well as lists and arrays, are represented in AiiDA with `ParameterData` objects. Get the PK and load the input parameters of a calculation in the graph of Section 2.2. Then display its content by typing

```
| params.get_dict()
```

where `params` is the `ParameterData` node you loaded. Modify the dictionary content so that the wave-function cutoff is now set to 20 Ry. Note that you cannot modify an object already stored in the database. To save the modification, you must create a new `ParameterData` object. Similarly to what discussed before, first load the `ParameterData` class by typing

```
| ParameterData = DataFactory('parameter')
```

Then an instance of the class (i.e. the parameter object that we want to create) is created and initialized by the command

```
| new_params = ParameterData(dict=YOUR_DICT)
```

where `YOUR_DICT` is the modified dictionary. Note that the parameter object is not yet stored in the database. In fact, if you simply type `new_params` in the verdi shell, you will be prompted with a string notifying you the “unstored” status. To save an entry in the database corresponding to the `new_params` object, you need to type a last command in the verdi shell:

```
| new_params.store()
```

¹The string provided to the `DataFactory` encodes both the location and the name of the required class according to some specific rules.

3.2.4 Structures

Find a structure in the graph of Section 2.2 and load it. Display its chemical formula, atomic positions and species using

```
| structure.get_formula()
| structure.sites
```

where `structure` is the structure you loaded. If you are familiar with ASE and PYMATGEN, you can convert this structure to those formats by typing

```
| structure.get_ase()
| structure.get_pymatgen()
```

Let's try now to define a new structure to study, specifically a silicon crystal. In the `verdi shell`, define a cubic unit cell as a 3×3 matrix, with lattice parameter $a_{lat} = 5.4 \text{ \AA}$:

```
| alat = 5.4
| the_cell = [[alat/2, alat/2, 0.], [alat/2, 0., alat/2], [0., alat/2, alat/2]]
```

Note: Default units for crystal structure cell and coordinates in AiiDA are \AA .

Structures in AiiDA are instances of `StructureData` class: load it in the `verdi shell`

```
| StructureData = DataFactory("structure")
```

Now, initialize the class instance (i.e. is the structure we want to study) by the command

```
| structure = StructureData(cell=the_cell)
```

which sets the cubic cell defined before. From now on, you can access the cell with the command

```
| structure.cell
```

Finally, append each of the 2 atoms of the cell command. You can do it using commands like

```
| structure.append_atom(position=(alat/4., alat/4., alat/4.), symbols="Si")
```

for the first 'Si' atom. Repeat it for the other atomic site (0, 0, 0). You can access and inspect² the structure sites with the command

```
| structure.sites
```

If you make a mistake, start over from `structure = StructureData(cell=the_cell)`, or equivalently use `structure.clear_kinds()` to remove all kinds (atomic species) and sites. Alternatively, AiiDA structures can also be converted directly from ASE [2] structures using³

```
| from ase.lattice.spacegroup import crystal
| ase_structure = crystal('Si', [(0,0,0)], spacegroup=227,
|                     cellpar=[alat, alat, alat, 90, 90, 90], primitive_cell=True)
| structure=StructureData(ase=ase_structure)
```

Now you can store the new structure object in the database with the command:

```
| structure.store()
```

Finally, we can also import the silicon structure from an external (online) repository such as the Crystallography Open Database [6]:

```
| from aida.tools.dbimporters.plugins.cod import CodDbImporter
| importer = CodDbImporter()
| for entry in importer.query(formula='Si', spacegroup='F d -3 m'):
|     structure = entry.get_aida_structure()
|     print "Formula", structure.get_formula()
|     print "Unit cell volume: ", structure.get_cell_volume()
```

In that case two duplicate structures are found for Si.

²if you set the structure incorrectly, for example with overlapping atoms, it is very likely that any DFT code will fail!

³We purposefully do not provide advanced commands for crystal structure manipulation in AiiDA, because python packages that accomplish such tasks already exist (such as ASE or pymatgen).

3.3 Accessing inputs and outputs

Load again the calculation node used in Section 3.1:

```
| calc = load_node(PK)
```

Then type

```
| calc.inp.
```

and press TAB: you will see all the link names between the calculation and its input nodes. You can use a specific linkname to access the corresponding input node, e.g.:

```
| calc.inp.structure
```

You can use the `inp` method multiple times in order to browse the graph. For instance, if the input structure node that you just accessed is the output of another calculation, you can access the latter by typing

```
| calc2 = calc.inp.structure.inp.output_structure
```

Here `calc2` is the `PwCalculation` that produced the structure used as an input for `calc`.

Similarly, if you type:

```
| calc2.out.
```

and then TAB, you will list all output link names of the calculation. One of them leads to the structure that was the input of `calc` we loaded previously:

```
| calc2.out.output_structure
```

Note that links have a single name, that was assigned by the calculation that used the corresponding input or produced the corresponding output, as illustrated in Fig. 1.

For a more programmatic approach, you can get a list of the inputs and outputs of a node, say `calc`, with the methods

```
| calc.get_inputs()
| calc.get_outputs()
```

Alternatively, you can get a dictionary where the keys are the link names and the values are the linked objects, with the methods

```
| calc.get_inputs_dict()
| calc.get_outputs_dict()
```

Note: You will sometime see entries in the dictionary with names like `output_kpoints_3511`. These exist because standard python dictionaries require unique key names while link labels may not be unique. Therefore, we use the link label plus the PK separated by underscores.

3.4 Pseudopotential families

Pseudopotentials in AiiDA are grouped in “families” that contain one single pseudo per element. We will see how to work with UPF pseudopotentials (the format used by Quantum ESPRESSO and some other codes).

Download and untar the SSSP [5] pseudopotentials via the commands:

```
wget https://archive.materialscloud.org/file/2018.0001/v1/SSSP_efficiency_pseudos.tar.gz
tar -zxvf SSSP_efficiency_pseudos.tar.gz
```

Then you can upload the whole set of pseudopotentials to AiiDA by to the following `verdi` command:

```
verdi data upf uploadfamily SSSP_efficiency_pseudos 'SSSP' 'SSSP pseudopotential library'
```

In the command above, `SSSP_efficiency_pseudos` is the folder containing the pseudopotentials, `'SSSP'` is the name given to the family and the last argument is its description.

Finally, you can list all the pseudo families present in the database with

```
verdi data upf listfamilies
```

4 Submit, monitor and debug calculations

The goal of this section is to understand how to create new data in AiiDA. We will launch a total energy calculation and check its results. We will introduce intentionally some common mistakes along the process of defining and submitting a calculation and we will explain you how to recognize and correct them. While this debugging is done here ‘manually’, workflows (that we will learn later in this tutorial) can automate this procedure considerably. For computing the DFT energy of the silicon crystal (with a PBE functional) we will use Quantum ESPRESSO [1], in particular the PWscf code (`pw.x`). Besides the AiiDA-core package, a number of plugins exist for many different codes. These are listed in the [AiiDA plugin registry](#)⁴. In particular, the “aiida-quantumespresso” plugin (already installed in your machine) provides a very extensive set of plugins, covering most (if not all) the functionalities of the underlying codes.

4.1 The AiiDA daemon

First of all, check that the AiiDA daemon is actually running. The AiiDA daemon is a program running all the time in the background, checking if new calculations appear and need to be submitted to the scheduler¹. The daemon also takes care of all the necessary operations before the calculation submission², and after the calculation has completed on the cluster.³ Type in the terminal

```
verdi daemon status
```

If the daemon is running, the output should look like

```
Profile: default
Daemon is running as PID 1650 since 2018-05-16 16:26:04
Active workers [1]:
  PID      MEM %    CPU %  started
  ----  -
  1653    8.225      0  2018-05-16 16:26:04
Use verdi daemon [incr | decr] [num] to increase / decrease the amount of workers
```

If this is not the case, type in the terminal

```
verdi daemon start
```

to start the daemon.

4.2 Creating a new calculation

To launch a calculation, you will need to interact with AiiDA mainly in the `verdi shell`. We strongly suggest you to first try the commands in the shell, and then copy them in a script “test_pw.py” using a text editor. This will be very useful for later execution of a similar series of commands.

The best way to run python scripts using AiiDA functionalities is to run them in a terminal by means of the command

```
verdi run <scriptname>
```

Every calculation sent to a cluster is linked to a code, which describes the executable file to be used. Therefore, first load the suitable code:

```
| code = Code.get_from_string(<codename>)
```

Here `Code` is the general AiiDA class handling all possible codes, and `code` is a class instance tagged as `<codename>` (see the first part of the tutorial for listing all codes installed in your AiiDA machine). You might also want to list only the codes that define a default calculation plugin for the `pw.x` code of Quantum ESPRESSO. You can do this with the following command:

⁴<https://aiidateam.github.io/aiida-registry/>

¹i.e. first queued on the cluster, and then run when the queuing system allows it

²creating a directory where to run and copying there the input files

³retrieving and copying some of the output files in the database repository, parsing one or several output files, attaching some data to the calculation and storing them in the database.

```
verdi code list -p quantumespresso.pw
```

Pick the correct codename, that might look like, e.g. `qe-pw-6.2.1@localhost`.

Once run, AiiDA calculations are instances of the class `Calculation`, more precisely of one of its subclasses, each corresponding to a code specific plugin (for example, the `PWscf` plugin). You have already seen `Calculation` classes in the previous sections.

However, to create a new calculation, rather than manually creating a new class, the suggested way is to use a `Builder`, that helps in setting the various calculation inputs and parameters, and provides TAB-completion.

To obtain a new builder, we can use the `get_builder` method of the `code` object:

```
| builder = code.get_builder()
```

This returns a builder that helps in setting up the inputs for the `PwCalculation` class (associated to the `quantumespresso.pw` plugin, i.e. the default plugin for the code you chose before).

As the first step, you can assign a (short) label or a (long) description to the calculation that you are going to create, that you might find convenient in the future. This can be achieved with:

```
| builder.label = "PW test"
| builder.description = "My first AiiDA calc with Quantum ESPRESSO on BaTiO3"
```

This information will be saved in the database for later query or inspection. Note that you can press TAB after writing `builder.` to see all available inputs.

Now you have to specify the number of machines (a.k.a. cluster nodes) you are going to run on and the maximum time allowed for the calculation. These general calculation options, that are independent of the code or plugin, but rather mainly passed later to the scheduler that handles the queue, are all grouped under “`builder.options`”:

```
| builder.options.resources = {'num_machines': 1}
| builder.options.max_wallclock_seconds = 30*60
```

Just like the normal inputs, these builder options are also TAB-completed. Type “`builder.options.`” and hit the TAB button to see the list of available options.

4.2.1 Preparation of inputs

Quantum ESPRESSO requires an input file containing Fortran namelists and variables, plus some cards sections (the documentation is available [online](http://www.quantum-espresso.org/wp-content/uploads/Doc/INPUT_PW.html)⁵). The Quantum ESPRESSO plugin of AiiDA requires quite a few nodes in input, which are documented [online](http://aiida-core.readthedocs.io/en/latest/plugins/quantumespresso/pw.html)⁶. Here we will instruct our calculation with a minimal configuration for computing the energy of silicon. We need:

1. Pseudopotentials
2. a structure
3. the k-points
4. the input parameters

We leave the parameters as the last thing to setup and start with structure, k-points, and pseudopotentials.

Use what you learned in the previous section and define these two kinds of objects in this script. Define in particular a silicon structure and a $2 \times 2 \times 2$ mesh of k-points. Notice that if you just copy and paste the code that you executed previously, you will create duplicated information in the database (i.e. every time you will execute the script, you will create another `StructureData`, another `KpointsData`, ...). In fact, you already have the opportunity to re-use an already existing structure.^a Use therefore a combination of the bash command `verdi data structure list` and of the shell command `load_node()` to get an object representing the structure created earlier.

^aHowever, to avoid duplication of `KpointsData`, you should first learn how to query the database, therefore we will ignore this duplication issue for now.

⁵http://www.quantum-espresso.org/wp-content/uploads/Doc/INPUT_PW.html

⁶<http://aiida-core.readthedocs.io/en/latest/plugins/quantumespresso/pw.html>

4.2.2 Attaching the input information to the calculation

So far we have defined (or loaded) some of the input data, but we haven't instructed the calculation to use them. To do this, let's just set the appropriate attributes of the builder (we assume here that you created the structure and k-points AiiDA nodes before and called them `structure` and `kpoints`, respectively):

```
builder.structure = structure
builder.kpoints = kpoints
```

Note that you can set in the builder both stored and unstored nodes. AiiDA will take care of storing the unstored nodes upon submission. Otherwise, if you decide not to submit, nothing will be stored in the database.

Moreover, PWscf also needs information on the pseudopotentials, specified by UpfData objects. This is set by storing a dictionary in “builder.pseudo”, with keys being the kind names, and value being the UpfData pseudopotential nodes. To simplify the task of choosing pseudopotentials, we can however use a helper function that automatically returns this dictionary picking the pseudopotentials from a given UPF family.

You can list the preconfigured families from the command line:

```
verdi data upf listfamilies
```

Pick the one you configured earlier or one of the SSSP families that we provide, and link it to the calculation using the command:

```
from aiida.orm.data.upf import get_pseudos_from_structure
builder.pseudo = get_pseudos_from_structure(structure, '<PSEUDO_FAMILY_NAME>')
```

4.2.3 Preparing and debugging input parameters

The last thing we miss is a set of parameters (i.e. cutoffs, convergence thresholds, etc...) to launch the Quantum ESPRESSO calculation. This part requires acquaintance with Quantum ESPRESSO and, very often, this is the part to tune when a calculation shows a problem. Let's therefore use this part of the tutorial to learn how to debug problems, and **let's introduce errors intentionally**. Note also that some of the problems we will investigate appear the first times you launch calculations and can be systematically avoided by using workflows.

Let's define a set of input parameters for Quantum ESPRESSO, preparing a dictionary of the form:

```
parameters_dict = {
    'CONTROL': {
        'calculation': 'scf',
        'tstress': True,
        'tprnfor': True,
    },
    'SYSTEM': {
        'ecutwfc': 30.,
        'ecutrho': 200.,
        'mickeymouse': 240.,
    },
    'ELECTRONS': {
        'conv_thr': 1.e-8,
    },
}
```

This dictionary is almost a valid input for the Quantum ESPRESSO plugin, except for an invalid key called “mickeymouse”. When Quantum ESPRESSO receives an unrecognized key (even when you misspell one) its behavior is to stop almost immediately. By default, the AiiDA plugin will not validate your input and simply pass it over. Therefore let's pass this dictionary to the calculation and observe this unsuccessful behavior.

As done before, load the ParameterData class

```
ParameterData = DataFactory("parameter")
```

and create an instance of the class containing all the input parameters you just defined

```
| parameters = ParameterData(dict=parameters_dict)
```

Finally, set the parameters in the builder

```
| builder.parameters = parameters
```

4.2.4 Simulate submission

At this stage, you have recreated in memory (it's not yet stored in the database) the input of the graph shown in Fig. 1a, whereas the outputs will be created later by the daemon.

In order to check how AiiDA creates the actual input files for the calculation, we can simulate the submission process with the (otherwise optional) command

```
| builder.submit_test()
```

This creates a folder of the form `submit_test/[date]-0000[x]` in the current directory. Check (in your second terminal) the input file `aiida.in` within this folder, comparing it with the content of the input data nodes you created earlier, and that the 'pseudo' folder contains the needed pseudopotentials. You can also check the submission script `_aiidasubmit.sh` (the scheduler that is installed on the machine is Torque, so AiiDA creates the files with the proper format for this scheduler). Note: you cannot correct the input file from the "submit_test" folder: you have to correct the script and re-execute it; the files created by `submit_test()` are only for final inspection.

4.2.5 Storing and submitting the calculation

Up to now the calculation `calc` is kept in memory and not in the database. We will now submit it, that will implicitly create a `PwCalculation` class, store it in the database, store also all its inputs parameters, k-points, structure, and properly link them. To submit it, run

```
| from aiida.work.launch import submit
| calc = submit(builder)
```

`calc` will now be the stored `PwCalculation`, already submitted to the daemon. The calculation has now a "database primary key" or `pk` (an integer ID) to the calculation (typing `calc.pk` will print this number). Moreover, it also gets a universally-unique ID (UUID), visible with `calc.uuid` that does not change even upon sharing the data with collaborators (while the `pk` will change in that case).

Now that the calculation is stored, you can also attach any additional attributes of your choice, which are called "extra" and defined in as key-value pairs. For example, you can add an extra attribute called `element`, with value `Si` through

```
| calc.set_extra("element", "Si")
```

You will see later the advantage of doing so for querying.

In the meantime, as soon as you submitted your calculation, the daemon picked it up and started to perform all the operations to do the actual submission, going through input file generation, submission to the queue, waiting for it to run and finish, retrieving the output files, parsing them, storing them in the database and setting the state of the calculation to `Finished`.

N.B. If the daemon is not running the calculation will remain in the `NEW` state until when you start it.

4.2.6 Checking the status of the calculation

You can check the calculation status from the command line:

```
| verdi calculation list
```

Note that `verdi` commands can be slow in this tutorial when the calculation is running (because you just have one CPU which is also used by the `PWscf` calculation).

By now, it is possible that the calculation you submitted has already finished, and therefore that you don't see any calculation in the output. In fact, by default, the command only prints calculations that are still being handled by the daemon, i.e. those with a state that is not `FINISHED` yet⁷.

To see also (your) calculations that have finished (and limit those only to the one created in the past day), use instead

```
verdi calculation list -a -p1
```

as explained in the first section.

To inspect the list of input files generated by the AiiDA (this can be done even when the calculation did not finish yet), type

```
verdi calculation inputls <pk_number> -c
```

with `pk_number` the pk number of your calculation. This will show the contents of the input directory (`-c` prints directories in colour). Then you can also check the content of the actual input file with

```
verdi calculation inputcat <pk_number> | less
```

4.3 Troubleshooting

After all this work the calculation should end up in a `FAILED` Job state (last column of `verdi calculation list`), and correspondingly the error code near the “Finished” status of the State should be non-zero (400 for `FAILED` calculations). This was expected, since we used an invalid key in the input parameters. Situations like this happen (probably often...) in real life, so we built in AiiDA the tools to traceback the problem source and correct it.

A first way to proceed is the manual inspection of the output file of `PWscf`. You can visualize it with:

```
verdi calculation outputcat <pk_number> | less
```

This can be a good primer for problem inspection. For something more compact, you can also try to inspect the calculation log (from AiiDA):

```
verdi calculation logshow <pk_number>
```

If the calculation has encountered a mistake, this log shows a handful of warnings coming from various processes, such as the daemon, the parser of the output or the scheduler on the cluster. In production runs, errors will mostly come from an unexpected termination of the `PWscf` calculation. The most programmatic way to handle these errors is to inspect the warnings key by loading the calculation object, say `calc`, and the using the following method:

```
| calc.res.warnings
```

This will print a list of strings reporting errors experienced during the execution, that can be easily read in python (and thus addressed programmatically), but are also reported in the calculation log. With any of these three methods you can understand that the problem is something like an ‘invalid input key’, which is exactly what we did.

Let's use a parameters dictionary that actually works. Modify the script `test_pw.py` script modifying the parameter dictionary as

```
parameters_dict = {
    "CONTROL": {"calculation": "scf",
               },
    "SYSTEM": {"ecutwfc": 30.,
              "ecutrho": 200.,
```

⁷For JobCalculations (i.e., calculations that are submitted to a remote computer through a scheduler) there is an additional “Job state” (last column of the output of `verdi calculation list`) that can either be `FINISHED` if all went well, or one of the possible failure states (`FAILED`, `PARSINGFAILED`, `SUBMISSIONFAILED`, `RETRIEVALFAILED`). These states are represented as a Finished state (third column of `verdi calculation list`, with a zero/non-zero error code depending if they finished/did not finish correctly). This latter state is more general than just JobCalculations and also applies to workflows, as we will see later in the tutorial.


```

        },
        "ELECTRONS": {"conv_thr": 1.e-6,
                       }
    }

```

If you launch the modified script by typing

```
verdi run test_pw.py
```

you should now be able to see a calculation reaching successfully the FINISHED state. Now you can access the results as you have seen earlier. For example, note down the pk of the calculation so that you can load it in the `verdi shell` and check the total energy with the commands:

```
calc=load_node(<pk>)
calc.res.energy
```

5 Queries in AiiDA: The QueryBuilder

This part of the tutorial is provided only in interactive mode through a Jupyter notebook, which you will be able to run in your browser. To accomplish this we first need to start the Jupyter server, if you didn't do it already at the very beginning of the tutorial. First make sure you are connected to the virtual machine with local forwarding enabled, as described in section ?? . Then, on the virtual machine, first make sure you are in the `aiida` virtual environment:

```
workon aiida
```

If the virtual environment is successfully loaded, your prompt should be prefixed with `(aiida)`. To finally launch the Jupyter server, execute the following commands:

```
cd ~/examples/aiida-demos/tutorial/
jupyter notebook --no-browser
```

If all went well, you should now be able to open up a browser on your local machine and point it to the following address `http://localhost:8888/?token=2a3ba3...` (replace the token with the one printed on output by the previous command). This should now show you a directory navigator. To open the notebook, click on `querybuilder` and then select the file `tutorial.ipynb`. Note that there is also a `solution.ipynb`, which is a copy of the same notebook, but which contains the solutions to all the exercises. You can use this version at your own discretion if you get stuck at some point (but we suggest that you try not to look at it at first).

6 AiiDA Workflows: workfunctions, workchains

The aim of the last part of this tutorial is to introduce the concept of workflows in AiiDA.

In this section, we will ask you to:

1. Understand how to keep the provenance when running small python scripts to convert one data object into another (postprocessing, preparation of inputs, ...)
2. Understand how to represent simple python functions in the AiiDA database
3. Learn how to write a simple workflow in AiiDA (without and with remote calculation submission)
4. Learn how to write a workflow with checkpoints: this means that, even if your workflows requires external calculations to start, them and their dependences are managed through the daemon. While you are waiting for the calculations to complete, you can stop and even shutdown the computer in which AiiDA is running. When you restart, the workflow will continue from where it was.

5. (optional) Go a bit deeper in the syntax of workflows with checkpoints (WorkChain), e.g. implementing a convergence workflow using `while` loops.

A note: this is probably the most “complex” part of the tutorial. We suggest that you try to understand the underlying logic behind the scripts, without focusing too much on the details of the workflows implementation or the syntax. If you want, you can then focus more on the technicalities in a second reading.

6.1 Introduction

The ultimate aim of this section is to create a workflow to calculate the equation of state of silicon. This is a very common task for an *ab initio* researcher. An equation of state consists in calculating the total energy E as a function of the unit cell volume V . The minimal energy is reached at the equilibrium volume V^* . Equivalently, the equilibrium is defined by a vanishing pressure $p = -dE/dV$. In the vicinity of the minimum, the functional form of the equation of state can be approximated by a parabola. Such an approximation greatly simplifies the calculation of the bulk modulus, that is proportional to the second derivative of the energy d^2E/dV^2 (a more advanced treatment requires fitting the curve with, e.g., the Birch–Murnaghan expression).

The process of calculating an equation of state puts together several operations. First, we need to define and store in the AiiDA database the basic structure of, e.g., bulk Si. Next, one has to define several structures with different lattice parameters. Those structures must be connected between them in the database, in order to ensure that their provenance is recorded. In other words, we want to be sure that in the future we will know that if we find a bunch of rescaled structures in the database, they all descend from the same one. How to link two nodes in the database in a easy way is the subject of Sec. 6.2.

In the following sections, the newly created structures will then serve as an input for total energy calculations performed, in this tutorial, with Quantum ESPRESSO. This task is very similar to what you have done in the previous part of the tutorial. Finally, you will fit the resulting energies as a function of volume to get the bulk modulus. As the EOS task is very common, we will show how to automate its computation with workflows, and how to deal with both serial and parallel (i.e., independent) execution of multiple tasks. Finally, we will show how to introduce more complex logic in your workflows such as loops and conditional statements (Sec. D), with an example on a convergence loop to find iteratively the minimum of an EOS.

6.2 Workfunctions: a way to generalize provenance in AiiDA

Imagine to have a function that takes as input a string of the name of a chemical element and generates the corresponding bulk structure as a `StructureData` object. The function might look like this (you will find this function in the folder `/home/aiida/tutorial_scripts/create_rescale.py` on your virtual machine):

```
def create_diamond_fcc(element):
    """
    Workfunction to create the crystal structure of a given element.
    For simplicity, only Si and Ge are valid elements.
    :param element: The element to create the structure with.
    :return: The structure.
    """
    import numpy as np
    elem_alat= {
        "Si": 5.431, # Angstrom
        "Ge": 5.658,
    }

    # Validate input element
    symbol = str(element)
    if symbol not in elem_alat.keys():
        raise ValueError("Valid elements are only Si and Ge")

    # Create cel starting having lattice parameter alat corresponding to the element
    alat = elem_alat[symbol]
```

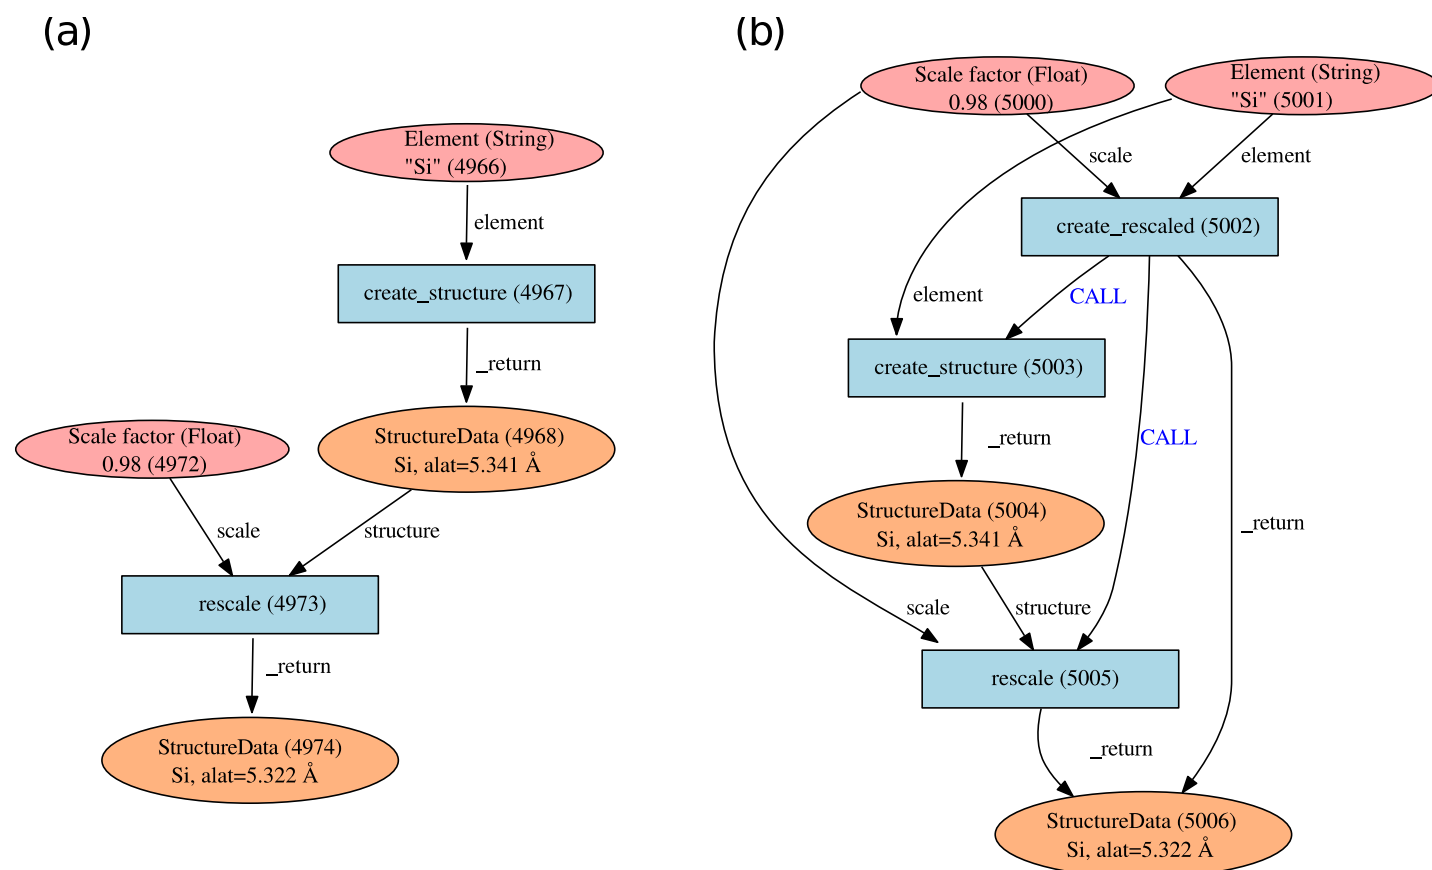


Figure 2 – Typical graphs created by using a workfunction. (a) The workfunction “create_structure” takes a **Str** object as input and returns a single **StructureData** object which is used as input for the workfunction “rescale” together with a **Float** object. This latter workfunction returns another **StructureData** object, defining a crystal having the rescaled lattice constant. (b) Graph generated by nesting workfunctions. A wrapper workfunction “create_rescaled” calls serially “create_structure” and “rescale”. This relationship is stored via “CALL” links.

```
the_cell = np.array([[0., 0.5, 0.5],
                    [0.5, 0., 0.5],
                    [0.5, 0.5, 0.]]) * alat

# Create a structure data object
StructureData = DataFactory("structure")
structure = StructureData(cell=the_cell)
structure.append_atom(position=(0., 0., 0.), symbols=str(element))
structure.append_atom(position=(0.25*alat, 0.25*alat, 0.25*alat),
                      symbols=str(element))

return structure
```

For the equation of state you need another function that takes as input a **StructureData** object and a rescaling factor, and returns a **StructureData** object with the rescaled lattice parameter (you will find this function in the same file `create_rescale.py` on your virtual machine):

```
def rescale(structure, scale):
    """
    Workfunction to rescale a structure

    :param structure: An AiiDA structure to rescale
    :param scale: The scale factor (for the lattice constant)
    :return: The rescaled structure
    """
    the_ase = structure.get_ase()
```

```

new_ase = the_ase.copy()
new_ase.set_cell(the_ase.get_cell() * float(scale), scale_atoms=True)
new_structure = DataFactory('structure')(ase=new_ase)
return new_structure

```

In order to generate the rescaled starting structures, say for five different lattice parameters you would combine the two functions. Enter the following commands in the `verdi shell` from the `tutorial_scripts` folder.

```

from create_rescale import create_diamond_fcc, rescale

s0 = create_diamond_fcc("Si")
rescaled_structures = [rescale(s0, factor) for factor
                        in (0.98, 0.99, 1.0, 1.1, 1.2)]

```

and store them in the database:

```

s0.store()
for struct in rescaled_structures:
    struct.store()

```

Run the commands above to store all the structures.

As expected, all the structures that you have created are not linked in any manner as you can verify via the `get_inputs()/get_outputs()` methods of the `StructureData` class. Instead, you would like these objects to be connected as sketched in Fig. 2a. Now that you are familiar with AiiDA, you know that the way to connect two data nodes is through a calculation. In order to “wrap” python functions and automate the generation of the needed links, in AiiDA we provide you with what we call “workfunctions”. A normal function can be converted to a workfunction by using the `@workfunction` decorator⁸ that takes care of storing the execution as a calculation and adding the links between the input and output data nodes.

In our case, what you need to do is to modify the two functions as follows (note that we import `workfunction` as `wf` to be shorter, but this is not required). You can do it in the file `create_rescale.py`:

```

# Add this import
from aiida.work import workfunction as wf

# Add decorators
@wf
def create_diamond_fcc(element):
    ...
    ...

@wf
def rescale(structure, scale):
    ...
    ...

```

Important: when you use workfunctions, you have to make sure that their input and output are actually `Data` nodes, so that they can be stored in the database. AiiDA objects such as `StructureData`, `ParameterData`, etc. carry around information about their provenance as stored in the database. This is why we must use the special database-storable types `Float`, `Str`, etc. as shown in the snippet below.

Try now to run the following script:

⁸In simple (or even simplified) words, a decorator is a function that modifies the behavior of another function. In python, a function can be decorated by adding a line of the form `@decorating_function_name` on the line just before the `def` line of the decorated function. If you want to know more, there are many online resources explaining python decorators.

```
from aiida.orm.data.base import Float, Str
from create_rescale import create_diamond_fcc, rescale

s0 = create_diamond_fcc(Str("Si"))
rescaled_structures = [rescale(s0,Float(factor)) for factor in (0.98, 0.99, 1.0, 1.1, 1.2)]
```

and check now that the output of `s0` as well as the input of the rescaled structures point to an intermediate `ProcessCalculation` node, representing the execution of the workfunction, see Fig. 2.

For instance, you can check that the output links of `s0` are the five `rescale` calculations:

```
s0.get_outputs()
```

which outputs

```
[<FunctionCalculation: uuid: 01b0b137-974c-4d80-974f-ea4978b12019 (pk: 4970)>,
 <FunctionCalculation: uuid: 1af5ead6-0ae0-42a7-969c-1f0e88300f4a (pk: 4973)>,
 <FunctionCalculation: uuid: 22dee9d5-0382-48a3-9319-e800506946f1 (pk: 4976)>,
 <FunctionCalculation: uuid: dc4c93b7-3e7a-4f51-8d44-cf15c5707ddb (pk: 4979)>,
 <FunctionCalculation: uuid: f5b4e9f2-0d50-4b3d-a76b-7a12d232ea54 (pk: 4982)>]
```

and the inputs of each `ProcessCalculation` (“rescale”) are obtained with:

```
for s in s0.get_outputs():
    print s.get_inputs()
```

that will return

```
[0.98, <StructureData: uuid: 9b76b5fa-2908-4f88-a4fb-7a9aa343a1f3 (pk: 4968)>]
[0.99, <StructureData: uuid: 9b76b5fa-2908-4f88-a4fb-7a9aa343a1f3 (pk: 4968)>]
[1.0, <StructureData: uuid: 9b76b5fa-2908-4f88-a4fb-7a9aa343a1f3 (pk: 4968)>]
[1.1, <StructureData: uuid: 9b76b5fa-2908-4f88-a4fb-7a9aa343a1f3 (pk: 4968)>]
[1.2, <StructureData: uuid: 9b76b5fa-2908-4f88-a4fb-7a9aa343a1f3 (pk: 4968)>]
```

6.2.1 Workfunction nesting

One key advantage of workfunctions is that they can be nested, namely, a workfunction can invoke workfunctions inside its definition, and this “call” relationship will also be automatically recorded in the database. As an example, let us combine the two previously defined workfunctions by means of a wrapper workfunction called “create_rescaled” that takes as input the element and the rescale factor.

Type in your shell (or modify the functions defined in `create_rescale.py` and then run):

```
@wf
def create_rescaled(element, scale):
    """
    Workfunction to create and immediately rescale
    a crystal structure of a given element.
    """
    s0 = create_diamond_fcc(element)
    return rescale(s0,scale)
```

and create an already rescaled structure by typing

```
| s1 = create_rescaled(element=Str("Si"), scale=Float(0.98))
```

Now inspect the input links of `s1`:

```
| In [6]: s1.get_inputs()
| Out[6]:
| [<FunctionCalculation: uuid: a672317b-3091-4135-9d84-12c2fff34bfe (pk: 5005)>,
|  <FunctionCalculation: uuid: a672317b-3091-4135-9d84-12c2fff34bfe (pk: 5005)>,
|  <FunctionCalculation: uuid: f64f4a70-70ff-4551-ba4d-c186328d8bd6 (pk: 5002)>]
```

The object `s1` has three incoming links, corresponding to *two* different calculations as input (in this case, pks 5002 and 5005). These correspond to the calculations “create_rescaled” and “rescale” as shown in Fig. 2b. It is normal that calculation 5005 has two links, don’t worry about that⁹. To see the “call” link, inspect now the outputs of the calculation appearing only once in the list. Write down its `<pk>` (in general, it will be different from 5002), then in the shell load the corresponding node and inspect the outputs:

```
| In [12]: p1 = load_node(<pk>)
| In [13]: p1.get_outputs_dict()
```

You should be able to identify the two “children” calculations as well as the final structure (you will see the calculations linked via CALL links: these are calculation-to-calculation links representing the fact that `create_rescaled` called two sub-workfunctions). The graphical representation of what you have in the database should match Fig. 2b.

6.3 Run a simple workflow

Let us now use the workfunctions that we have just created to build a simple workflow to calculate the equation of state of silicon. We will consider five different values of the lattice parameter obtained rescaling the experimental minimum, $a = 5.431 \text{ \AA}$, by a factor in $[0.96, 0.98, 1.0, 1.02, 1.04]$. We will write a simple script that runs a series of five calculations and at the end returns the volume and the total energy corresponding to each value of the lattice parameter. For your convenience, besides the functions that you have written so far in the file `create_rescale.py`, we provide you with some other utilities to get the correct pseudopotential and to generate a pw input file, in the module `common_wf.py` which has been put in the `tutorial_scripts` folder.

We have already created the following script named `simple_sync_workflow.py`, which you are free to look at but please go through the lines carefully and make sure you understand them. If you decide to create your own new script, make sure to also place it in the folder `tutorial_scripts`, otherwise the imports won’t work.

Besides the functions in the local folder

```
| from create_rescale import create_diamond_fcc, rescale
| from common_wf import generate_scf_input_params
```

you need to import few further AiiDA classes and functions:

```
| from aiida.work import run, Process
| from aiida.work import workfunction as wf
| from aiida.orm.data.base import Str, Float
| from aiida.orm import CalculationFactory, DataFactory
```

⁹If you are curious: the two links have the same label, but are of different *link_type*: one is a **create** link, that keeps track of the calculation that actually generated the node. Instead the other one is of type **return**, stating that the workfunction, beside creating that node, also returned it as an output. Calculation 5002 instead only returned the node but it did not generate it, therefore there is only one link between it and the final `StructureData`.

The only imported function that deserves an explanation is `run`. For the time being, you just need to know that it is a function that needs to be used to execute a new workflow. The actual body of the script is the following. We suggest that you first have a careful look at it before running it.

```
# Load the calculation class 'PwCalculation' using its entry point 'quantumespresso.pw'
PwCalculation = CalculationFactory('quantumespresso.pw')

scale_facs = (0.96, 0.98, 1.0, 1.02, 1.04)
labels = ["c1", "c2", "c3", "c4", "c5"]

@wf
def run_eos_wf(codename, pseudo_family, element):
    print "Workfunction node identifiers: {}".format(Process.current().calc)
    s0 = create_diamond_fcc(Str(element))

    calcs = {}
    for label, factor in zip(labels, scale_facs):
        s = rescale(s0, Float(factor))
        inputs = generate_scf_input_params(s, str(codename), Str(pseudo_family))
        print "Running a scf for {} with scale factor {}".format(element, factor)
        result = run(PwCalculation, **inputs)
        print "RESULT: {}".format(result)
        calcs[label] = get_info(result)

    eos = []
    for label in labels:
        eos.append(calcs[label])

    # Return information to plot the EOS
    ParameterData = DataFactory("parameter")
    retdict = {
        'initial_structure': s0,
        'result': ParameterData(dict={'eos_data': eos})
    }

    return retdict
```

If you look into the previous snippets of code, you will notice that the way we submit a QE calculation is slightly different from what you have seen in the first part of the tutorial. The following:

```
| result = run(PwCalculation, **inputs)
```

runs in the current python session (without the daemon), waits for its completion and returns the output in the user-defined variable `result`. The latter is a dictionary whose values are the output nodes generated by the calculation, with the link labels as keys. For example, once the calculation is finished, in order to access the total energy, we need to access the `ParameterData` node which is linked via the “output_parameters” link (see again Fig. 1 of Day 1 Tutorial, to see inputs and outputs of a Quantum ESPRESSO calculation). Once the right node is retrieved as `result['output_parameters']`, we need to get the `energy` attribute. The global operation is achieved by the command

```
| result['output_parameters'].dict.energy
```

As you see, the function `run_eos_wf` has been decorated as a workfunction to keep track of the provenance. Finally, in order to get the <pk> associated to the workfunction (and print on the screen for our later reference), we have used the following command to get the node corresponding to the `ProcessCalculation`:

```
| from aiida.work import Process
| print Process.current().calc
```

To run the workflow it suffices to call the function `run_eos_wf` in a python script providing the required input parameters. For simplicity, we have included few lines at the end of the script that invoke the function with a static choice of parameters:

```
def run_eos(codename='pw-5.1@localhost', pseudo_family='GBRV_lda', element="Si"):
    return run_eos_wf(Str(codename), Str(pseudo_family), Str(element))

if __name__ == '__main__':
    run_eos()
```

Run the workflow by running the following command from the `tutorial_scripts` directory:

```
verdi run simple_sync_workflow.py
```

and write down the `<pk>` of the `ProcessCalculation` printed on screen at execution.

The command above locks the shell until the full workflow has completed (we will see in a moment how to avoid this). While the calculation is running, you can use (in a different shell) the command `verdi work list` to show ongoing and finished workfunctions. You can “grep” for the `<pk>` you are interested in. Additionally, you can use the command `verdi work status <pk>` to show the tree of the sub-workfunctions called by the root workfunction with a given `<pk>`.

Wait for the calculation to finish, then call the function `plot_eos(<pk>)` that we provided in the file `common_wf.py` to plot the equation of state and fit it with a Birch–Murnaghan equation.

6.4 Run multiple calculations

You should have noticed that the calculations for different lattice parameters are executed serially, although they might perfectly be executed in parallel because their inputs and outputs are not connected in any way. In the language of workflows, these calculations are executed in a synchronous (or blocking) way, whereas we would like to have them running *asynchronously* (i.e., in a non-blocking way, to run them in parallel). One way to achieve this is to submit the calculation to the daemon using the `submit` function. Make a copy of the script `simple_sync_workflow.py` that we worked on in the previous section and name it `simple_submit_workflow.py`. To make the new script work asynchronously, simply change the following subset of lines:

```
from aiida.work import run
[...]
for label, factor in zip(labels, scale_facs):
    [...]
    result = run(PwCalculation, **inputs)
    calcs[label] = get_info(result)
[...]
eos = []
for label in labels:
    eos.append(calcs[label])
```

replacing them with

```
from aiida.work import submit
from time import sleep
[...]
for label, factor in zip(labels, scale_facs):
    [...]
    calcs[label] = submit(PwCalculation, **inputs)
[...]
# Wait for the calculations to finish
```



```

    for calc in calcs.values():
        while not calc.is_finished:
            sleep(1)

eos = []
for label in labels:
    eos.append(get_info(calcs[label].get_outputs_dict()))

```

The main differences are:

- `run` is replaced by `submit`
- The return value of `submit` is not a dictionary describing the outputs of the calculation, but it is the calculation node for that submission.
- Each calculation starts in the background and calculation nodes are added to the `calc` dictionary.
- At the end of the loop, when all calculations have been launched with `submit`, another loop is used to wait for all calculations to finish before gathering the results as the final step.

In the next section we will show you another way to achieve this, which has the added bonus that it introduces checkpoints in the workfunction, from which the calculation can be resumed should it be interrupted.

After applying the modifications, run the script. You will see that all calculations start at the same time, without waiting for the previous ones to finish.

If in the meantime you run `verdi work status <pk>`, all five calculations are already shown as output. Also, if you run `verdi calculation list`, you will see how the calculations are submitted to the scheduler.

6.5 Workchains, or how not to get lost if your computer shuts down or crashes

The simple workflows that we have used so far have been launched by a python script that needs to be running for the whole time of the execution, namely the time in which the calculations are submitted, and the actual time needed by Quantum ESPRESSO to perform the calculation and the time taken to retrieve the results. If you had killed the main python process during this time, the workflow would not have terminated correctly. Perhaps you have kill the calculation and you experienced the unpleasant consequences: intermediate calculation results are potentially lost and it is extremely difficult to restart a workflow from the exact place where it stopped.

In order to overcome this limitation, in AiiDA we have implemented a way to insert checkpoints, where the main code defining a workflow can be stopped (you can even shut down the machine on which AiiDA is running!). We call these workfunctions with checkpoints “workchains” because, as you will see, they basically amount to splitting a workfunction in a chain of steps. Each step is then ran by the daemon, in a way similar to the remote calculations.

The basic rules that allow you to convert your workfunction-based script to a workchain-based one are listed in Table 1, which focus on the code used to perform the calculation of an equation of state. The modifications needed are put side-to-side to allow for a direct comparison. In the following, when referencing a specific part of the code we will refer to the line number appearing in Table 1.

- Instead of using decorated functions you need to define a class, inheriting from a prototype class called `WorkChain` that is provided by AiiDA (line 4)
- Within your class you need to implement a `define` classmethod that always takes `cls` and `spec` as inputs. (lines 6–7). Here you specify the main information on the workchain, in particular:
 - the *inputs* that the workchain expects. This is obtained by means of the `spec.input()` method, which provides as the key feature the automatic validation of the input types via the `valid_type` argument (lines 8–10). The same holds true for outputs, as you can use the `spec.output()` method to state what output types are expected to be returned by the workchain. Both `spec.input()` and `spec.output()` methods are optional, and if not specified, the workchain will accept any set of inputs and will not perform any check on the outputs, as long as the values are database storable AiiDA types.

Table 1: Side-to-side comparison of the EOS workflow using standard workfunctions (left panel) or “workchains” (right panel).

Workfunctions		Workchains	
1	from aiida.work import submit, Process	1	from aiida.work.workchain import WorkChain, ToContext
2	#	2	#
3		3	
4		4	class EquationOfStates(WorkChain):
5		5	@classmethod
6		6	def define(cls, spec):
7		7	super (EquationOfStates, cls).define(spec)
8		8	spec. input ('element', valid_type=Str)
9		9	spec. input ('code', valid_type=Str)
10		10	spec. input ('pseudo-family', valid_type=Str)
11		11	spec.outline(
12		12	cls.run-pw,
13		13	cls.return-results,
14		14)
15		15	
16	@wf	16	
17	def run-eos-wf(codename, pseudo-family, element):	17	def run-pw(self):
18	# ...	18	# ...
19	s0 = create-diamond-fcc(Str(element))	19	self.ctx.s0 = create-diamond-fcc(Str(self.inputs.element))
20		20	
21		21	
22	calcs = {}	22	calcs = {}
23	for label, factor in zip(labels, scale-facs):	23	for label, factor in zip(labels, scale-facs):
24	s = rescale(s0, Float(factor))	24	s = rescale(self.ctx.s0, Float(factor))
25	inputs = generate-scf-input-params(25	inputs = generate-scf-input-params(
26	s, str (codename), pseudo-family)	26	s, str (self.inputs.code), self.inputs.pseudo-family)
27	# ...	27	# ...
28	calcs[label] = submit(PwCalculation, **inputs)	28	future = self.submit(PwCalculation, **inputs)
29		29	calcs[label] = future
30		30	
31	<i># Wait for the calculations to finish</i>	31	<i># Ask the workflow to continue when the results are ready</i>
32	for calc in calcs.values():	32	<i># and store them in the context</i>
33	while not calc.is-finished:	33	return ToContext(**calcs)
34	sleep(1)	34	
35		35	
36	eos = []	36	def return-results(self):
37	for label in labels:	37	eos = []
38	eos.append(get-info(calcs[label].get-outputs-dict()))	38	for label in labels:
39		39	eos.append(get-info(self.ctx[label].get-outputs-dict()))
40	<i>#Return information to plot the EOS</i>	40	<i># Return information to plot the EOS</i>
41	ParameterData = DataFactory("parameter")	41	ParameterData = DataFactory('parameter')
42	retdict = {	42	retdict = {
43	'initial-structure': s0,	43	'initial-structure': self.ctx.s0,
44	'result': ParameterData(dict ={'eos-data': eos})	44	'result': ParameterData(dict ={'eos-data': eos})
45	}	45	}
46	return retdict	46	for link-name, node in retdict.iteritems():
47		47	self.out(link-name, node)

- the `outline` consisting in a list of “steps” that you want to run, put in the right sequence (lines 11–14). This is obtained by means of the method `spec.outline()` which takes as input the steps. *Note*: in this example we just split the main execution in two sequential steps, that is, first `run_pw` then `return_results`. However, more complex logic is allowed, as will be explained in the Sec. D.
- You need to split your main code into methods, with the names you specified before into the outline (`run_pw` and `return_results` in this example, lines 17 and 35). Where exactly should you split the code? Well, the splitting points should be put where you would normally block the execution of the script for collecting results in a standard workfunction, namely whenever you call the method `.result()`. Each method should accept only one parameter, `self`, e.g. `def step_name(self)`.
- You will notice that the methods reference the attribute `ctx` through `self.ctx`, which is called the *context* and is inherited from the base class `WorkChain`. A python function or workfunction normally just stores variables in the local scope of the function. For instance, in the example of the subsection 6.3, you stored the `calc_results` in the `eos` list, that was a local variable. In workchains, instead, to preserve variables between different steps, you need to store them in a special dictionary called *context*. As explained above, the context variable `ctx` is inherited from the base class `WorkChain`, and at each step method you just need to update its content. AiiDA will take care of saving the context somewhere between workflow steps (on disk, in the database, ..., depending on how AiiDA was configured). For your convenience, you can also access the value of a context variable as `self.ctx.varname` instead of `self.ctx['varname']` (see e.g. lines 19, 24, 38, 43).
- Any submission within the workflow should not call the normal `run` or `submit` functions, but `self.submit` to which you have to pass the `Process` class, and a dictionary of inputs (line 28).
- The submission in line 28, returns a future and not the actual calculation, because at that point in time we have only just launched the calculation to the daemon and it is not yet completed. Therefore it literally is a “future” result. Yet we still need to add these futures to the context, so that in the next step of the workchain, when the calculations are in fact completed, we can access them and continue the work. To do this, we can use the `ToContext` class. This class takes a dictionary, where the values are the futures and the keys will be the names under which the corresponding calculations will be made available in the context when they are done. See line 33 how the `ToContext` object is created and returned from the step. By doing this, the workchain will implicitly wait for the results of all the futures you have specified, and then call the next step *only when all futures have completed*.
- *Return values*: While in a normal workfunction you attach output nodes to the `FunctionCalculation` by invoking the `return` statement, in a workchain you need to call `self.out(link_name, node)` for each node you want to return (line 46-47). Of course, if you have already prepared a dictionary of outputs, you can just use the following syntax:

```
|self.out_many(retdict) # Keys are link names, value the nodes
```

The advantage of this different syntax is that you can start emitting output nodes already in the middle of the execution, and not necessarily at the very end as it happens for normal functions (`return` is always the last instruction executed in a function). Also, note that once you have called `self.out(link_name, node)` on a given `link_name`, you can no longer call `self.out()` on the same `link_name`: this will raise an exception.

Inspect the example in the table that compares the two versions of workfunctions to understand in detail the different syntaxes.

Finally, the workflow has to be run. For this you have to use the function `run` passing as arguments the `EquationOfStates` class and the inputs as key-value arguments. For example, you can execute

```
|run(EquationOfStates, element=Str('Si'), code=Str('qe-pw-6.2.1@localhost'),
    pseudo_family=Str('GBRV_lda'))
```

While the workflow is running, you can check (in a different terminal) what is happening to the calculations using `verdi calculation list`. You will see that after a few seconds the calculations are all submitted to the scheduler and can potentially run at the same time.

Note: You will see warnings that say ‘Exception trying to save checkpoint, this means you will not be able to restart in case of a crash until the next successful checkpoint’, these are generated by the `PwCalculation` which is unable to save a checkpoint because it is not in a so called ‘importable path’. Simply put this means that if AiiDA were to try and reload the class it wouldn’t know which file to find it in.

To get around this you could simply put the workflow in a different file that is in the ‘PYTHONPATH’ and then launch it by importing it in your launch file, this way AiiDA knows where to find it next time it loads the checkpoint.

As an additional exercise (optional), instead of running the main workflow (`EquationOfStates`), try to submit it. Note that the file where the `WorkChain` is defined will need to be globally importable (so the daemon knows how to load it) and you need to launch it (with `submit`) from a different python file. The easiest way to achieve this is typically to embed the workflow inside a python package.

Note: As good practice, you should try to keep the steps as short as possible in term of execution time. The reason is that the daemon can be stopped and restarted only between execution steps and not if a step is in the middle of a long execution.

Finally, as an optional exercise if you have time, you can jump to the [Appendix D](#), which shows how to introduce more complex logic into your `WorkChains` (if conditionals, while loops etc.). The exercise will show how to realize a convergence loop to obtain the minimum-volume structure in a EOS using the Newton’s algorithm.

7 WorkChains, a real-world example: computing a band structure for a simple crystal structure

Note: *If you still have enough time, you might want to check first [Appendix D](#) before continuing with this section.*

As a final demonstration of the power of `WorkChains` in AiiDA, we want to give a demonstration of a `WorkChain` that we have written that will take a structure as its only input and will compute its band structure. All of the steps that would normally have to be done manually by the researcher, choosing appropriate pseudopotentials, energy cutoffs, k-points meshes, high-symmetry k-point paths and performing the various calculation steps, are performed automatically by the `WorkChain`.

The demonstration of the workflow will be performed in a Jupyter notebook. To run it, follow the instructions that were given for the `querybuilder` notebook in [section 5](#). The only difference is that instead of selecting the notebook in the `querybuilder` directory, go to `pw/bandstructure` instead and choose the `bandstructure.ipynb` notebook. There you will find some example structures that are loaded from COD, through the importer integrated within AiiDA. Note that the required time to calculate the bandstructure for these example structures ranges from 3 minutes to almost half an hour, given that the virtual machine is running on a single core with minimal computational power. It is not necessary to run these examples as it may take too long to complete. For reference, the expected output band structures are plotted in [Fig. 3](#).

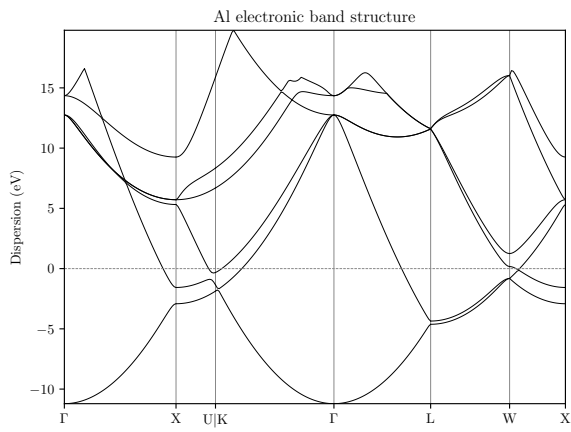
Appendices

The following appendices consist of optional exercises, and are mentioned in earlier parts of the tutorial. Go through them only if you have time.

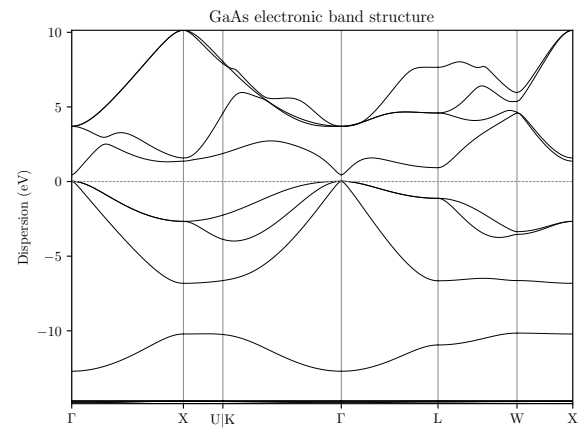
A Calculation input validation

This appendix shows additional ways to debug possible errors with QE, how to use a useful tool that we included in AiiDA to validate the input to Quantum ESPRESSO (and possibly suggest the correct name to misspelled keywords)

There are various reasons why you might end up providing a wrong input to a Quantum ESPRESSO calculation. Let’s check for example this input dictionary, where we inserted two mistakes:



(a) Al



(b) GaAs

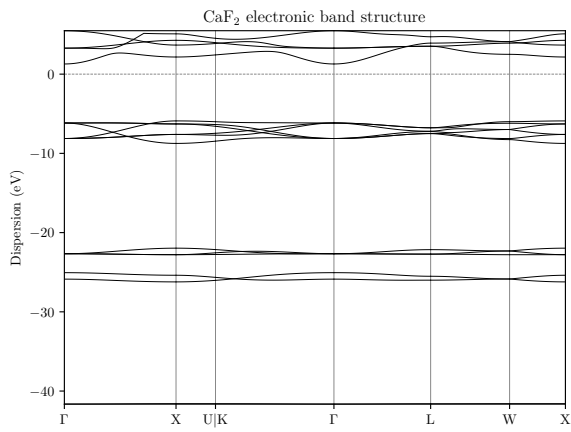

(c) CaF₂(d) *h*-BN

Figure 3 – Electronic band structures of four different crystal structures computed with AiiDA’s PwBandsWorkChain

```
parameters_dict = {
    "CTRL": {
        "calculation": "scf",
        "restart_mode": "from_scratch",
    },
    "SYSTEM": {
        "nat": 2,
        "ecutwfc": 30.,
        "ecutrho": 200.,
    },
    "ELECTRONS": {
        "conv_thr": 1.e-6,
    }
}
```

The two mistakes in the dictionary are the following. First, we wrote a wrong namelist name ('CTRL' instead of 'CONTROL'). Second, we inserted the number of atoms explicitly: while that is how the number of atoms is specified in Quantum ESPRESSO, in AiiDA this key is reserved by the system: in fact this information is already contained in the StructureData. Modify the script with this ParameterData. In this case, we use a tool (the input validator that we provide in AiiDA) to check the input file before submitting. Therefore, the behavior will be slightly different from the previous example: the plugin will check for some keys and will refuse to submit the calculation. This kind of mistakes is revealed by either

- Submitting a calculation. You will see the calculation ending up in the SUBMISSIONFAILED status. Check therefore the logs to recognize the source of the error.
- Submitting a test. You will not be able to successfully create the test and the traceback will guide you to the problem.

Over the time this kind of trivial mistakes can be annoying, but they can be avoided with a utility function that checks the “grammar” of the input parameters. In your script, after you defined the parameters_dict, you can validate it with the command (note that you need to pass also the input crystal structure, `s`, to allow the validator to perform all needed checks):

```
PwCalculation = CalculationFactory('quantumespresso.pw')
validated_dict = PwCalculation.input_helper(parameters_dict, structure=s)
parameters = ParameterData(dict=validated_dict)
```

The `input_helper` method will check for the correctness of the input parameters. If misspelling or incorrect keys are detected, the method raises an exception, which stops the script before submitting the calculation and thus allows for a more effective debugging.

With this utility, you can also provide a list of keys, without providing the namelists (useful if you don't remember where to put some variables in the input file). Hence, you can provide to `input_helper()` a dictionary like this one:

```
parameters_dict = {
    'calculation': 'scf',
    'tstress': True,
    'tprnfor': True,
    'ecutwfc': 30.,
    'ecutrho': 200.,
    'conv_thr': 1.e-6,
}
validated_dict = PwCalculation.input_helper(
    parameters_dict, structure=s, flat_mode=True)
```

If you print the `validated_dict`, it will look like:

```
{
    "CONTROL": {
        "calculation": "scf",
        "tstress": True,
        "tprnfor": True,
    },
    "SYSTEM": {
        "ecutwfc": 30.,
        "ecutrho": 200.,
    },
    "ELECTRONS": {
        "conv_thr": 1.e-6,
    }
}
```

B Restarting calculations

Up to now, we have only presented cases in which we were passing wrong input parameters to the calculations, which required us to modify the input scripts and relaunch calculations from scratch. There are several other scenarios in which, more generally, we need to restart calculations from the last step that they have executed. For example when we run molecular dynamics, we might want to add more time steps than we initially thought, or as another example you might want to refine the relaxation of a structure with tighter parameters.

In this section, you will learn how to restart and/or modify a calculation that has run previously. As an example, let us first submit a total energy calculation using a parameters dictionary of the form:

```
parameters_dict = {
    'CONTROL': {
        'calculation': 'scf',
        'tstress': True,
        'tprnfor': True,
    },
    'SYSTEM': {
        'ecutwfc': 30.,
        'ecutrho': 200.,
    },
    'ELECTRONS': {
        'conv_thr': 1.e-14,
        'electron_maxstep': 3,
    },
}
```

and submit the calculation with this input. In this case, we set a very low number of self consistent iterations (3), too small to be able to reach the desired accuracy of 10^{-14} : therefore the calculation will not reach a complete end and will be flagged in a FAILED state. However, there is no mistake in the parameter dictionary.

Now, create a new script file, where you will try to restart and correct the input dictionary. We first load the calculation that has just failed (let's call it `c1`)

```
old_calc = load_node(PK)
```

(take care of using the correct PK). Then, create a new Builder `builder` which is set to reuse all inputs from the previous step, with a few adaptations to the input parameters that might be needed by the code to properly deal with restarts.

```
from aiida_quantumespresso.utils.restart import create_restart_pw
builder = create_restart_pw(
    old_calc,
    use_output_structure=False,
```

```
restart_from_beginning=False,
force_restart=True)
```

The flag usage (most of them are optional) is:

- `use_output_structure`: if `True` and `old_calc` has an output structure, the new calculation will use it as input;
- `restart_from_beginning`: if `False` the new calculation will start from the charge density of `old_calc`, it will start from the beginning otherwise;
- `force_restart`: if `True`, the new calculation will be created even if `old_calc` is not in a `FINISHED` job state.

Since this calculation has exactly the same parameters of before, we have to modify the input parameters and increase `electron_maxstep` to a larger value. To this aim, let's load the dictionary of values and change it

```
old_parameters = builder.parameters
parameters_dict = old_parameters.get_dict()
parameters_dict['ELECTRONS']['electron_maxstep'] = 100
```

Note that you cannot modify the `old_parameters` object: it has been used by calculation `c1` and is saved in the database; hence a modification would break the provenance. We have to create a new `ParameterData` and pass it to `c2`:

```
ParameterData = DataFactory('parameter')
new_parameters = ParameterData(dict=parameters_dict)
builder.parameters = new_parameters
```

Now you can launch the new calculation

```
from aiida.work.run import submit
new_calc = submit(builder)
print new_calc.pk
```

that this time can proceed until the end and return converged total energy. Using the restart method, the script is much shorter than the one needed to launch a new one from scratch: you didn't need to define pseudopotentials, structures and k-points, which are the same as before. You can indeed inspect the new calculation to check that now it actually completed successfully.

C Queries in AiiDA - Optional examples and exercises

Optional exercises on relationships (Task 3)

Hint for the exercises:

- You can have projections on properties of more than one entity in your query. You just have to add the *project* key (specifying the list of properties that you want to project) along with the corresponding entity when you append it.

Exercises:

Try to write the following queries:

- Find all descendants of a `StructureData` with a certain uuid. Print both the `StructureData` and the descendant.
- Find all the `FolderData` created by a specific user.

Optional exercises on attributes and extras (Task 4)

Hint for the exercises:

- You can easily order or limit the number of the results by using the *order_by()* and *limit()* methods of the QueryBuilder. For example, we can order all our job calculation by their *id* and limit the result to the first 10 as follows:

```
qb = QueryBuilder()
qb.append(JobCalculation, tag='calc')
qb.order_by({'calc':'id'})
qb.limit(10)
qb.all()
```

Exercises:

- Write a code snippet that informs you how many pseudopotentials you have for each element.
- Smearing contribution to the total energy for calculations:
 - Write a query that returns the smearing contribution to the energy stored in some instances of ParameterData.
 - Extend the previous query to also get the input structures.
 - Which structures have a smearing contribution to the energy smaller or equal to -0.02?

Summarizing what we learned by now - An example

At this point you should be able to do queries with projections, joins and filters. Moreover, you saw how to apply filters and projections even on attributes and extras. Let's discover the full power of the QueryBuilder with a complex graph query that allows you to project various properties from different nodes and apply different filters and joins.

Imagine that you would like to get the smearing energy for all the calculations that have finished and have a Sn_2O_3 as input. Moreover, besides from the smearing energy, you would like to print the units of this energy and the formula of the structure that was given to the calculation. The graphical representation of this query can be seen in Figure 4 and the actual query follows:

```
qb = QueryBuilder()
qb.append(
    StructureData,
    project=["extras.formula"],
    filters={"extras.formula":"Sn2O3"},
    tag="structure"
)
qb.append(
    Calculation,
    tag="calculation",
    output_of="structure"
)
qb.append(
    ParameterData,
    tag="results",
    filters={"attributes.energy_smearing":{"<=":-0.0001}},
    project=[
        "attributes.energy_smearing",
        "attributes.energy_smearing_units",
```

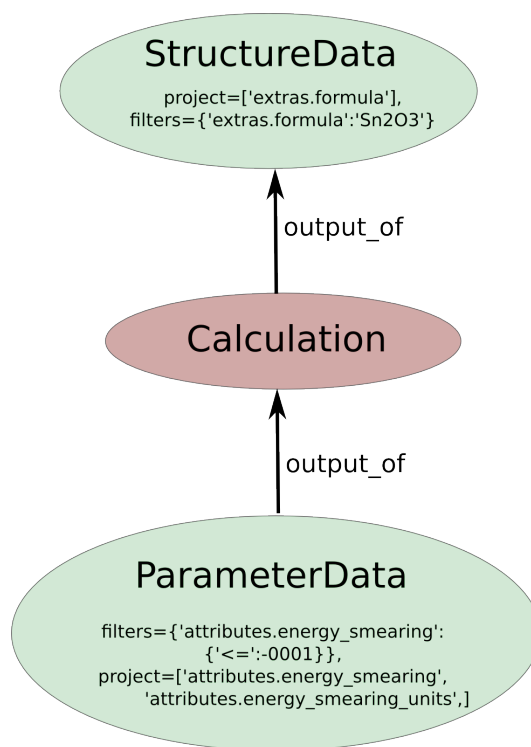


Figure 4 – Complex graph query.

```

    ],
    output_of="calculation"
)
qb.all()

```

D More complex logic in workflows: while loops and conditional statements

In the previous sections, you have been introduced to WorkChains, and the reason for using them over “standard” workfunctions (i.e., functions decorated with `@wf`).

However, in the example of Sec. 6.5, the `spec.outline` was quite simple, with a “static” sequence of two steps. Most often, however, you need dynamic workflows, where you need to decide at runtime whether to continue to compute or not (e.g. in a convergence loop, where you need to stop if convergence has been achieved). To support this scenario, the `spec.outline` can support logic: *while* loops and *if/elif/else* blocks. The simplest way to explain it is to show an example:

```

from aiida.work.workflow import if_, while_

spec.outline(
    cls.s1,
    if_(cls.isA)(
        cls.s2
    ).elif_(cls.isB)(
        cls.s3
    ).else_(
        cls.s4
    ),
    cls.s5,
    while_(cls.condition)(
        cls.s6
    ),
)

```

that would *roughly* correspond, in a python syntax, to:

```
s1()
if isA():
    s2()
elif isB():
    s3()
else:
    s4()
s5()
while condition():
    s6()
```

The only constraint is that condition functions (in the example above `isA`, `isB` and `condition`) must be class methods that returns `True` or `False` depending on whether the condition is met or not.

A suggestion on how to write new workchains: Use the outline to help you in designing the logic. First create the spec outline writing, almost if you were explaining it in words, what you expect the workflow to do. Then, define one by one the methods. For example, we have prepared a simple workfunction to optimize the lattice parameter of silicon efficiently using a Newton’s algorithm on the energy derivative, i.e. the pressure $p = -dE/dV$. You can find it the code at `tutorial_scripts/pressure_convergence.py`. The outline looks like this:

```
spec.outline(
    cls.init,
    cls.put_step0_in_ctx,
    cls.move_next_step,
    while_(cls.not_converged)(
        cls.move_next_step,
    ),
    cls.report
)
```

This outline already roughly explains the algorithm: after an initialization (`init`) and putting the first step (number zero) in the ctx (`put_step0_in_ctx`), a function to move to the next step is called (`move_next_step`). This is iterated while a given convergence criterion is not met (`not_converged`). Finally, some reporting is done, including returning some output nodes (`report`).

If you are interested in the details of the algorithm, you can inspect the file. The main ideas are described here:

init Generate a `pw.x` calculation for the input structure (with volume V), and one for a structure where the volume is $V + 4\text{\AA}^3$ (just to get a closeby volume). Store the results in the context as `r0` and `r1`

put_step0_in_ctx Store in the context V , $E(V)$ and dE/dV for the first calculation `r0`

move_next_step This is the most important function. Calculate V , $E(V)$ and dE/dV for `r1`. Also, estimate d^2E/dV^2 from the finite difference of the first derivative of `r0` and `r1` (helper functions to achieve this are provided). Get the a , b and c coefficients of a parabolic fit $E = aV^2 + bV + c$ and estimated the expected minimum of the EOS function as the minimum of the fit $V_0 = -b/2a$. Finally, replace `r0` with `r1` in the context (i.e., get rid of the oldest point) and launch a new pw calculation at volume V_0 , that will be stored in the context replacing `r1`. In this way, at the next iteration `r0` and `r1` will contain the latest two simulations. Finally, at each step some relevant information (coefficients a , b and c , volumes, energies, energy derivatives, ...) are stored in a list called `steps`. This whole list is stored in the context because it provides quantities to be preserved between different workfunction steps.

not_converged Return `True` if convergence has not been achieved yet. Convergence is achieved if the difference in volume between the two latest simulations is smaller than a given threshold (`volume_tolerance`).

report This is the final step. Mainly, we return the output nodes: `steps` with the list of results at each step, and `structure` with the final converged structure.

The results returned in `steps` can be used to represent the evolution of the minimisation algorithm. A possible way to visualize it is presented in Fig. 5, obtained with an initial lattice constant of $a_{\text{lat}} = 5.2\text{\AA}$.




Figure 5 – Example of results of the convergence algorithm presented in Sec. D. The bottom plot is a zoom near the minimum. The dots represent the (volume,energy) points obtained from Quantum ESPRESSO, and the numbers indicate at which iteration they were obtained. The parabolas represent the parabolic fits used in the algorithm; the minimum of the parabola is represented with a small cross, in correspondence of the vertical lines, used as the volume for the following step.

References

- [1] P. Giannozzi et al., J.Phys. Cond. Matt. 29, 465901 (2017).
- [2] S. R. Bahn and K. W. Jacobsen, Comput. Sci. Eng., 4, 56-66 (2002).
- [3] S. Ping Ong et al., Comput. Mater. Sci. 68, 314-319 (2013).
- [4] K.F. Garrity, J.W. Bennett, K.M. Rabe and D. Vanderbilt, Comput. Mater. Sci. 81, 446 (2014).
- [5] G. Prandini, A. Marrazzo, I. E. Castelli, N. Mounet, N. Marzari, A Standard Solid State Pseudopotentials (SSSP) library optimized for accuracy and efficiency (Version 1.0, data download), Materials Cloud Archive (2018), [doi:10.24435/materialscloud:2018.0001/v1](https://doi.org/10.24435/materialscloud:2018.0001/v1).
- [6] Crystallographic Open Database (COD), <http://www.crystallography.net/cod/>.