



## 3DIX

# Calculation of X-ray absorption spectra

The 3DIX project produces 3D Imaging with X-rays. The objective is to make 3D images (volumes) of nanoscale objects to study their characteristics on a nanoscale. These studies can then be applied to a wide variety of objects and scientific fields such as chemical studies, life sciences, structure of materials etc.

## Goal

Experiments at the ESRF include the collection of data containing X-ray absorption spectra. The interpretation of this data is of fundamental interest for the understanding of the underlying phenomena, but also for applications in material science. The FDMNES programme (developed at the CNRS in Grenoble, [www.neel.cnrs.fr/fdmnes](http://www.neel.cnrs.fr/fdmnes)) allows you to calculate this spectra, using a variety of theoretical methods.

The calculations done by FDMNES require very little data transfer and storage: a total of a few GByte for the executable and the input and output data. They are, however, often very compute intensive, involving 50 - 100 cores for several days. They may also require large amounts of memory, 8 GB / core is common. Running such computations therefore uses considerable computing resources, but this load varies considerably over time depending on the experiments done and the research interests of the scientists. **Instead of building upon the computing facilities in our institute for a rarely occurring maximum demand, the 3DIX project wants to offload such peak loads into the cloud.**

## Preconditions

Tasks of the cloud access interface:

1. individual user accounts and user groups
2. budget allocation and usage monitoring with a cut-off per group
3. startup and shutdown of cloud clusters
4. file transfer between local and cloud clusters
5. login and job startup on cloud cluster

Tasks 1 & 2 can be done by GUI or CLI. For 3-5, develop scripts with minimal number of parameters

The hardware must be able to support:

6. up to 20 jobs submitted simultaneously, 20 - 200 cores / job and up to 8 GB memory / job
7. availability of GPUs is needed for later use case (not present use case)

To distribute the calculations over the nodes, the required minimum is:

8. support for MPI (ideally OpenMPI), OpenMP and hybrid parallel processing
9. possibility of under-subscription of nodes for memory-intensive jobs
10. suspending / checkpointing / restarting of jobs

Task 9 & 10 are typically handled by a job scheduler.

## The Challenge

The principal challenges that the job is required to run are:

- » the functionality and user-friendliness of the cloud access interface;
- » the availability of hard- and software resources to run compute-intensive jobs on many cores.

The cloud access interface must be structured in a way that an every-day user can create a compute cluster on the cloud, transfer the programmes and input data to it, run the job and get the result back without having to ask for help from the IT staff.

The hardware must be able to support computing loads as mentioned for the FDMNES case above. The software available on the cloud must allow the user to run a multi-node job without having to worry about the details of how the calculations are distributed over the nodes.

## From the user perspective

In order to serve a useful purpose for the ESRF, cloud computing must not tie up too many resources of support staff. Our users are normally scientists with rather limited IT knowledge. It must be possible in less than half a day to give them a general introduction in the use of the cloud computing facilities, create an account and provide them with a few sets of scripts that cover typical use cases. From there, the users are expected to act on their own, modify the examples to fit their needs and contact IT staff only in the case of problems.

## Benefits and impact

The community using FDMNES only involves about 10 scientists, but FDMNES is just one example of theoretical calculations that can be done to understand the experimental results at the ESRF. Once we have established a procedure to run FDMNES on the cloud, it should be easy to adapt this for other calculations that have the same requirements: small data sets transfer, but long (several days) and memory intensive (10 - 20 GB / core) calculations running on a large (up to several hundreds) number of cores. The expected benefit is a faster turnaround for those calculations on the one hand, but it should also lessen the load on our local compute resources and thus make those more easily available for the data processing that we need or want to do on our own compute cluster.

## Procurer sponsoring the use case: ESRF

The ESRF is an international research institute funded by 21 partner countries. Its purpose is fundamental and applied research with synchrotron radiation, which is an extremely focussed and intense form of X-rays. This offers scientists the opportunity to explore materials and living matter in a multitude of fields, ranging from chemistry and materials physics to archaeology and from structural biology, health and life sciences to environmental sciences, information science and nanotechnologies. These experiments are done at 43 experimental stations, each specialised for a certain type of research.

