

CRYSTAL17 - Parallel implementation (Replicated data)

Overview

The basic parallel versions of CRYSTAL modules, denoted as **Pcrystal** and **Pproperties**, uses a replicated data algorithm. Each host runs the same code and performs a number of independent *tasks* which are distributed at run time. One host is chosen as the master. The master host spawns the program onto other hosts (slaves) and operates dynamical load balancing of the task execution via a shared atomic counter. During integral generation a task is defined as the calculation of a block of integrals. Thus each node computes a number of integrals which are stored to its local disk. During a SCF cycle, a partial Hamiltonian matrix (**F**) is built on each node from those integrals which have been stored locally. The matrices are then passed between nodes so that each has a complete copy. The diagonalization of **F** at each k-point is treated as an independent task which is distributed. After diagonalization the eigenvalues are communicated to all nodes.

This strategy is comparatively easy to implement and successful on architectures where each node has access to fast disk storage and sufficient memory to run a complete copy of CRYSTAL. Low speed communication hardware (such as Ethernet) is usually sufficient. Performance depends critically on the system considered.

The integral generation step is performed efficiently when the number of integrals to be generated is much larger than the number of nodes. This condition is satisfied in most applications. Machines with up to 64 nodes have been used effectively on large cases. In the SCF process the construction of **F** is also efficient. Diagonalization of **F** is performed efficiently if the number of k-points is much larger than the number of nodes. This condition is usually not satisfied for large systems and thus diagonalization may be the most costly phase.

The parallel version of CRYSTAL requires a mechanism for initiating processes on remote machines and a library of routines to provide inter-process communication. There are many implementations of this functionality available and CRYSTAL has been modified to take advantage of the MPI message-passing library.

MPI Parallel version of CRYSTAL17

Running the MPI parallel version of CRYSTAL under Linux

The CRYSTAL17 parallel executables for Linux (*Pcrystal*, *Pproperties*) are based on **OpenMPI** implementation of the MPI message-passing library and have been generated with the following features:

1. Fortran compiler: Intel Fortran Compiler XE 17.0
2. MPI libraries:
 - **Openmpi-2.1.0** (see <http://www.open-mpi.org/>)
3. Processor communication: TCP/IP (ch_p4)
4. Processor connection: ssh

The CRYSTAL17 parallel version is supposed to run on homogeneous workstation networks, Beowulf cluster and individual workstations.

To run the MPI parallel version of CRYSTAL17 under Linux special attention must be paid to set the proper environment:

1. OpenMPI must be installed according to the adopted processor connection remote shell (either rsh or ssh) The `mpirun` load module may then be used to initiate parallel execution of CRYSTAL17 from the master host.
2. each node must allow access via a remote shell, either rsh or ssh, to the master host. Note that the CRYSTAL17 parallel executable to be used will depend on the adopted remote shell.

In the following procedure, we refer to **Pcrystal**, but the same holds for **Pproperties**. Before a parallel job can be submitted:

1. there must be a consistent set of CRYSTAL17 parallel modules (**Pcrystal**) available on each node (e.g. through a NFS filesystem)
2. the CRYSTAL input deck must be provided on each node of the cluster in a file named INPUT.

Workstation clusters require each process in a parallel job be started individually. The procedure to run CRYSTAL17 can then be summarized as:

- 1) create a temporary directory on each node (workstation)
 - 2) either copy **Pcrystal** in the temporary directory of each node (workstation) or make **Pcrystal** available to each node through a NFS filesystem
 - 3) copy the CRYSTAL17 input deck as INPUT in the temporary directory of each node (workstation)
- 4) OpenMPI – prepare a file with the list of nodes (workstations) to be used in the parallel run. The file is usually indicated as `machines.arch`, where `arch` is the architecture of the system (e.g. LINUX) and it can be located in the working directory or in the directory. The format is one hostname per line, with either `hostname` or `hostname:n`, where `n` is the number of processors in a cluster of symmetric multiprocessors. The `hostname` should be the same as the result from the command "`hostname`".

A sample file for a cluster of 6 nodes with a processor each will look like:

```
#
node9
node10
node11
node12
node13
node14
```

You can change this file to contain the machines that you want to use to run MPI jobs on.

- connect to the node which will be the master host (e.g. `node9`).
- move to the temporary directory of the master host and run `mpirun` as:

```
mpirun -np nprocs -machinefile machines.arch Pcrystal
```

This will run **Pcrystal** on the first `nprocs` processors in the `machines.arch`, located in the working directory. According to the list of nodes above, if `nprocs=4`, the program will run on: `node9`, `node10`, `node11` and `node12`.

– For more details please refer to the *Open MPI Software Documentation*

The output file will be displayed on the standard error. Use common Unix commands for redirecting `stderr` to a file.

Note that scripts are available to run Pcrystal and Pproperties (see: <http://www.crystal.unito.it/utills/utills17.zip>)