

How to launch NAMD

- [NAMD on the clusters](#)
- [Charm++](#)
- [Running NAMD](#)
 - [Example job script](#)
 - [Potential errors](#)

NAMD on the clusters

In order to use NAMD one must first load a compiler, MPI implementation and linear algebra library. NAMD is currently only built for the Intel stack so this implies:

```
$ module load intel intel-mpi intel-mkl
$ module load namd
$ module list
Currently Loaded Modules:
1) intel/17.0.2 2) intel-mpi/2017.2.174 3) intel-mkl/2017.2.174 4) namd/2.12
```

NAMD is parallelised using the Charm++ framework which is somewhat different to the more commonly used MPI and OpenMP models. For this reason we briefly describe Charm++ and its options.

Charm++

"Charm++ is a C++-based parallel programming system, founded on the migratable-objects programming model, and supported by a novel and powerful adaptive runtime system." (<http://charm.cs.illinois.edu/manuals/html/charm++/1.html>)

Charm++ is available as a module (`module load intel intel-mpi charm`). Note that there is no need to load the Charm++ module to run NAMD

Charm can be configured in many different ways and on the SCITAS clusters it is built using the SMP mode. From the documentation:

SMP mode in Charm++ spawns one OS process per logical node. Within this process there are two types of threads:

- *Worker Threads that have objects mapped to them and execute entry methods*
- *Communication Thread that sends and receives data (depending on the network layer)*

Charm++ always spawns one communication thread per process when using SMP mode and as many worker threads as the user specifies (see the options below). In general, the worker threads produce messages and hand them to the communication thread, which receives messages and schedules them on worker threads.

The options for Charm++ itself also apply to NAMD and the possible options are described at <http://charm.cs.illinois.edu/manuals/html/charm++/C.html>

As we use SMP mode the most important option is the number of worker threads per compute node. This needs to be given explicitly as Charm++ cannot deduce it from the run-time environment.

```
++ppn
Number of PEs (or worker threads) per logical node (OS process). This option should be specified even when
using platform specific launchers (e.g., aprun, ibrun).
```

This value should be one less than the number of CPU cores per node so that one is left free for the communications process.

Running NAMD

Due to the SMP mode of Charm++ we need to specify 1 task per node and assign all the available CPUs to this task. The number depends on the cluster so you will need one script per processor type.

Cluster	CPUs per node	Notes
Deneb phase 1	16	E5v2
Deneb phase 2	24	E5v3
Fidis	28	E5v4 and s6g1

Example job script

We assume that you already know how to prepare input for NAMD - if not please contact your local specialist.

A sample script that asks for 2 nodes on Fidis and sets the ++ppn option based on the number of CPU cores you have asked for is given below:

```
#!/bin/bash
#SBATCH --nodes=2
#SBATCH --ntasks-per-node=1
#SBATCH --cpus-per-task=28

module purge
module load intel intel-mpi intel-mkl
module load namd

let "workers = $SLURM_CPUS_ON_NODE - 1"

srun namd2 ++ppn=${workers} myinput.conf
```

This should then give an output starting with:

```
Charm++> Running on MPI version: 3.1
Charm++> level of thread support used: MPI_THREAD_FUNNELED (desired: MPI_THREAD_FUNNELED)
Charm++> Running in SMP mode: numNodes 2, 27 worker threads per process
```

Potential errors

Please check that there are no warning messages given at the start which are often due to not correctly setting the number of tasks per node to be 1. An example of a "bad" run is as follows with the warning being "**Warning: the number of SMP threads (256) is greater than the number of physical cores (24)**"

```
Charm++> Running on 2 unique compute nodes (24-way SMP).
Charm++> cpu topology info is gathered in 1.715 seconds.
Charm++> Warning: the number of SMP threads (256) is greater than the number of physical cores (24), so threads
will sleep while idling. Use +CmiSpinOnIdle or +CmiSleepOnIdle to control this directly.
Info: NAMD 2.12 for linux-x86_64
Info:
Info: Please visit http://www.ks.uiuc.edu/Research/namd/
Info: for updates, documentation, and support information.
Info:
Info: Please cite Phillips et al., J. Comp. Chem. 26:1781-1802 (2005)
Info: in all publications reporting results obtained with NAMD.
Info:
Info: Based on Charm++/Converse 60701 for charm-6.7.1-f2vzlvdlhxoy73gtjvrbdd2ohzk73s
Info: Built Sun Aug 20 17:08:56 CEST 2017 by scitasbuild on r01-node49
Info: 1 NAMD 2.12 linux-x86_64 480 r01-node01 user
Info: Running on 480 processors, 32 nodes, 2 physical nodes.
```