Benchmark studies using GPU accelaration for large-scale molecular dynamics simulations of the Arp2/3 complex

Paraskevi Gkeka, George Patargias, Zoe Cournia

Biomedical Research Foundation of the Academy of Athens,
4 Soranou Ephessiou, 11527 Athens, Greece

Abstract

The present report summarizes the benchmark study on GPU acceleration performed during the GPU CyTera Workshop that took place in the Cyprus Institute, Nicosia, 10-12 December 2012 for Cy-Tera project number lsprob10541. Two systems were used for the benchmarking: 1) A system with the Arp2/3-VCA complex in explicit water and 2) a system with the Arp2/3-2VCA-2actin complex in explicit water. The total size of the systems was approximately 350,000 and 600,000 atoms, respectively. According to the literature, runtimes of several hundreds of nanoseconds are required to identify putative VCA binding sites on Arp2/3 and to elucidate the structural details of Arp2/3 activation, which are the main goals of this project. Based on the size of the Arp2/3-VCA and Arp2/3-2VCA-2actin systems the specific project would greatly benefit by the use of a large cluster and GPU acceleration, as confirmed by the results presented herein.

1 Background and significance

Actin filaments play an important role in the cell as they act as force-generating polymer motors, structural scaffolds, and tracks for motor protein and aid important cellular processes such as cell migration, vesicle trafficking and lamellar motility. [1] Precise regulation of actin filament growth is crucial in order to enable the cell to change its shape and stiffness rapidly in response to external stimuli. [2] One of the three classes of proteins that initiate new filament polymerization is the Arp2/3 complex, which includes two actin-related proteins (ARP2 and ARP3) and five novel protein subunits (ARPC1, ARPC2, APRC3, APRC4, ARPC5). Before the formation of an actin filament branch, the complex is initially found in an open, "inactive" state, where Arp2 and Arp3 are found 30 Å apart. When an activating factor binds the complex, a large conformational change takes place to form the closed, "active" state, where Arp2 and Arp3 are brought together and initiate the actin filament nucleation. [1, 3, 4, 5]. The mechanism of this conformational change that causes the Arp2/3 complex activation, which is an essential step for actin filament nucleation, is still under debate.

The comparison of the crystal structure representing the inactive state with EM reconstructions of the Arp2/3 complex bound at the mother filament actin junction, suggested that Arp2/3 activation is accompanied by a large conformational change.[6] A number of studies have attempted to shed light to the structural nature of Arp2/3 activation. [3, 7, 8] However, a number of questions remain open:

- Does the conformational change include rearrangement of only Arp2, Arp3 or, both?
- Do the other subunits (ARPC1, ARPC3-5) of the Arp2/3 complex take part in the motions required for the conformational change and, if so, in what way?
- Is there an intermediate state, as suggested by EM studies [6] between the open and closed conformation of Arp2/3 complex?

- Is there one or more pathways characterizing the transition from the open to the closed state and is it possible to discriminate them in terms of energy differences?
- What is the mode of action of ATP and WASP as ligands that favor the closed state of the Arp2/3 complex?

Molecular dynamics (MD) simulations is a powerful technique that has been extensively employed to explain structure-function relationships in biomolecules at atomistic level [9, 10, 11] and within complex chemical environments [12]. Through equilibrium or biased MD simulations of large proteins using state-of-the-art computer clusters with GPU accelaration one can reach the timescales needed to explore the questions described above.

2 Materials and methods

2.1 Hardware

The benchmark study was performed on the CyTera cluster located in the Cyprus Institute during the LinkSCEEM/Cy-Tera GPU training CyTera is a hybrid CPU/GPGPU Linux cluster. The facility cosists of 98 CPU only compute nodes and a further 18 computes nodes with 2 GPUs each (see GPU specifications below).

Compute Node specifications
CyTera is an IBM dx360 M3 cluster with dual Xeon X 5650 Westmere hexacore 2.67 GHz CPUs and the following specifications:

Node Type	RAM	Disk	Other
compute	48GB DDR3 @1333MHz	146GB	NA
gpucompute	48GB DDR3 @1333MHz	146GB	2 x Nvidia M2070 (Tesla)

Interconnect (MPI Message Passing and Storage Access Network)
Mellanox 4xQDR Infiniband (40Gbps)

Storage
300TB GPFS

2.2 Software

For our benchmarks, we used GROMACS version 4.6 and NAMD version 2.9. Different compilations of GROMACS were tested with and without GPUs, with threads, and mpi-enabled.

2.3 System setup

In the present benchmark study, we used two test systems one with Arp2/3-2VCA (312,158 atoms for GROMACS and 291,360 atoms for NAMD) and another with Arp2/3-2VCA-2actin (546,263 atoms) both in explicit water. The systems were prepared by Dr. George Patargias.

2.4 Simulation parameters

The Molecular Dynamics simulations were carried out using the GROMACS 4.6 and NAMD 2.9 MD packages. The CHARMM22 force field (includes phi, psi cross term map (CMAP) correction) was used to model all protein interactions for NAMD and the AMBER99SB-ILDN was used for GROMACS. The TIP3P model was used for water in both cases. The simulations were performed at constant pressure, temperature, and number of particles (NPT ensemble). The temperature was kept constant at 310 K. The nonbonded potential energy functions were cut off and shifted at 14 Å, with forces smoothly decaying between 10 and 14 Å. The particle-mesh Ewald method (PME) was employed to calculate long-range electrostatic interactions. The simulations were run using a 2-fs integration time step and a total number of 50,000 timesteps. The output coordinates, velocities and energies were saved every 100 steps. More details about the simulation parameters can be found in the Appendix.

3 Benchmarks and scaling

3.1 Strong scalability

The two systems in the present project contained approximately $300,000$ and $550,000$ interaction centers over a domain of 16 nm $\times$ 13 nm $\times$ 15 nm for the small system and 19 nm $\times$ 18 nm $\times$ 16 nm for the large system. We used PME for to resolve the long range electrostatics with a short range cutoff of 1.4 nm and a timestep of 2 fs along with the SETTLE algorithm to maintain the water bonds rigid.

The strong scalability can be measured as:

$$\eta_{\text{efficiency}} = \mathrm{N}^{\text{ref}}_{\text{CPUs}}\mathrm{T}(\mathrm{N}^{\text{ref}}_{\text{CPUs}})/(\mathrm{N}_{\text{CPUs}}\mathrm{T}(\mathrm{N}_{\text{CPUs}}))$$

where T is the average computation time of one time step and $N^{ref}_{CPUS} = 2$ for our scaling tests. Moreover, the speed-up of the calculations relative a reference system has been calculated.

4 GROMACS 4.6

In Figure 1, we present the speed up of our calculation as a result of GPU accelaration with GROMACS 4.6. The results show that in both systems a 10-fold speed-up is apparent upon introduction of two GPUs in a two-CPU calculation. The effect of GPU accelaration seems to dimish upon usage of more than eight CPUs. The same effect can be observed in Figure 2, where the speed-up over increasing number of CPUs and a constant number of GPUs equal to two is used. In an effort to optimize the CPU-GPU ratio, we ended up with a scheme of two CPUs and two GPUs per node. In Figure 3, we present the speed-up of increasing the number of nodes in a calculation and thus increasing the

number of CPUs/GPUs by two each time. The ns/day and s/step for these calculations are presented in Figure 4.

We should also mention that, as discussed in GROMACS official web page (*http* : *//www.gromacs.org*), for optimal performance on multi-socket servers, groups of OpenMP threads belonging to an MPI process/ thread-MPI should run on the same CPU/socket and this requires that the number of processes is a multiple of the number of CPUs/sockets in the respective machine and the number of cores per CPU is divisible by the number of threads per process. Launching M MPI processes with N OpenMP threads each:

mpirun -np M mdrun_mpi -ntomp N

Thus, based on the above, on a dual 6-core machine N=6, M=2 or N=3, M=4 should run more efficiently than N=4 M=3. In our case then, an N=6 was used with the -ntomp flag.

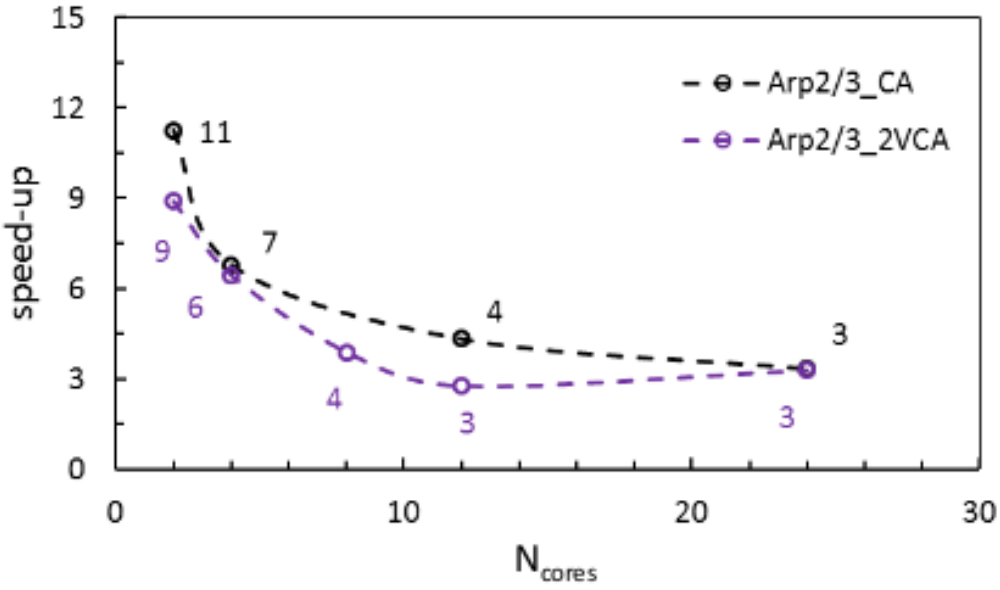

Figure 1: Speed-up between single precision simulations with and without GPUs over increasing number of CPUs on one node. Black line corresponds to the Arp2/3 CA system and purple line corresponds to the Arp2/3 2VCA system. The tests with the GPUs were done on the CyTera GPU nodes whereas the rest of the benchmarks were run on the CPU nodes, using GROMACS 4.6 with and without GPU-enabled respectively. The best speed-up is achieved for both systems with 2CPUs/2GPUs per node.

4.0.1 Single vs Double precision

Scaling studies were also performed without GPU accelation to compare between single and double precision as well as the use of threads or mpi. The strong scalability results are presented in Figures 5 and 6. From the results, we can see that all the systems scale very well, however double precision calculations are always two times slower than single precision (Figures 7 and 8).

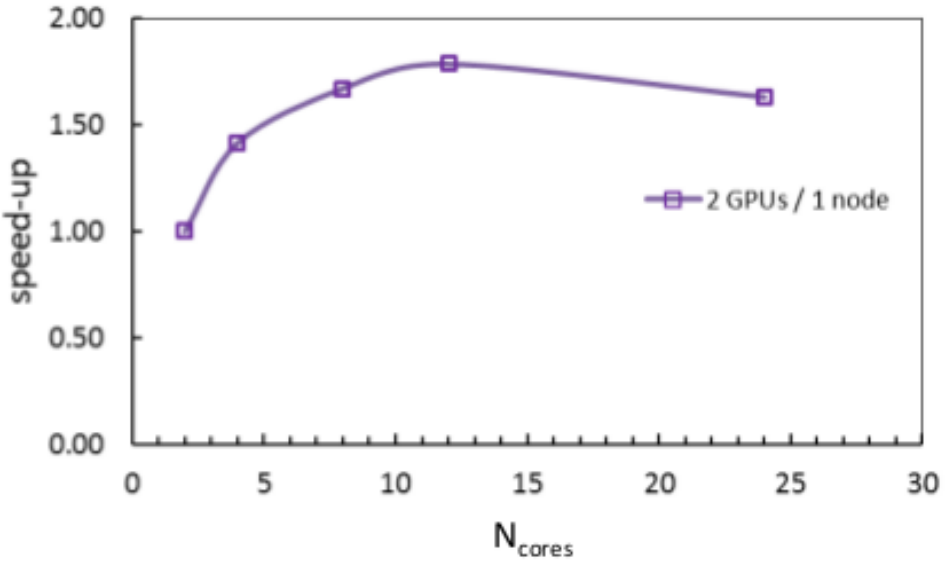

Figure 2: Speed-up over increasing number of CPUs on one node carrying two GPUs. The benmarks were run for the Arp2/3 2VCA system on the CyTera GPU nodes.

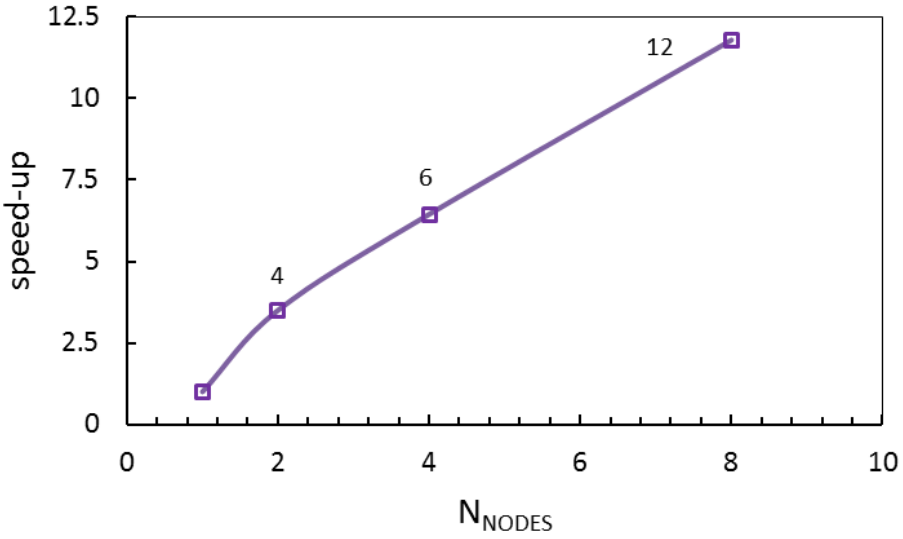

Figure 3: Speed-up over increasing number of nodes. The benmarks were run for the Arp2/3 2VCA system on the CyTera GPU nodes.

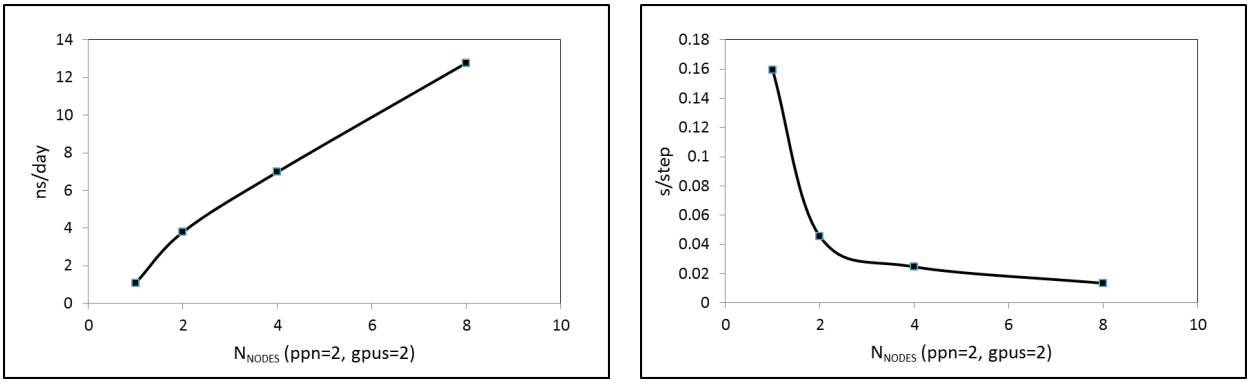

(a) (b)

Figure 4: 4(a) ns/day and 4(b) s/step using a 2CPUs:2GPUs scheme and increasing number of nodes.

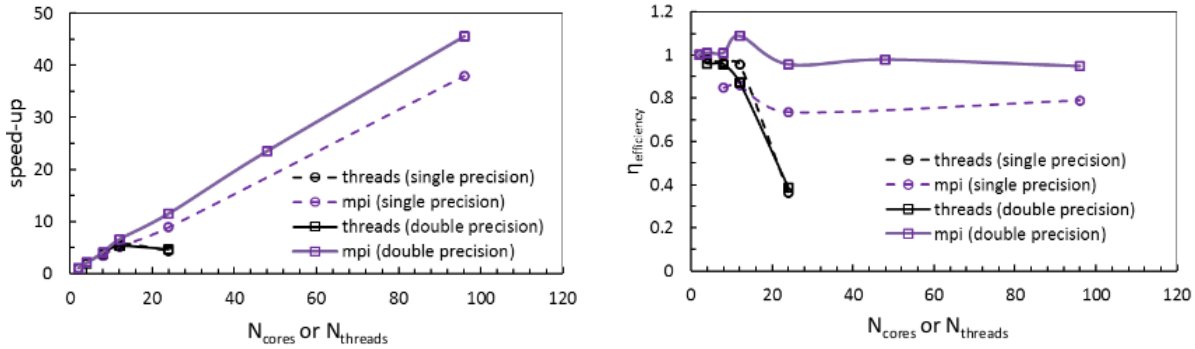

(a) (b)

Figure 5: (a) speed-up and (b) efficiency over increasing number of CPUs for the Arp 2/3 VCA system without GPUs.

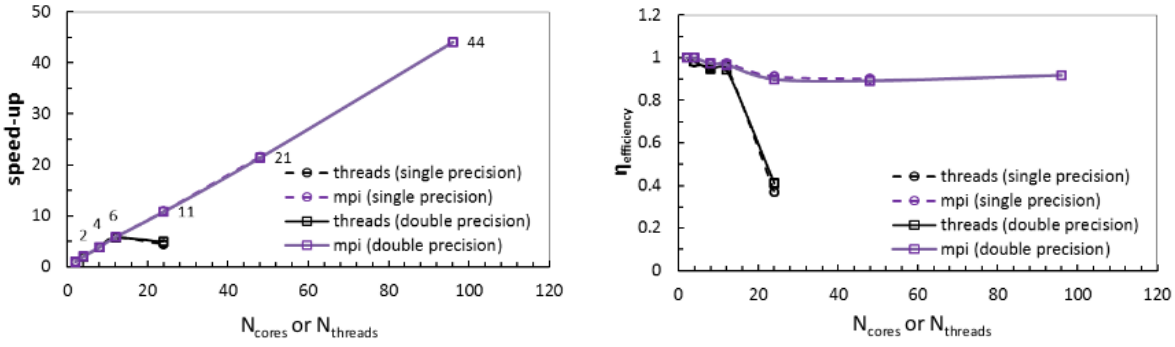

(a) (b)

Figure 6: (a) speed-up and (b) efficiency over increasing number of CPUs for the Arp 2/3-2VCA-actin system without GPUs.

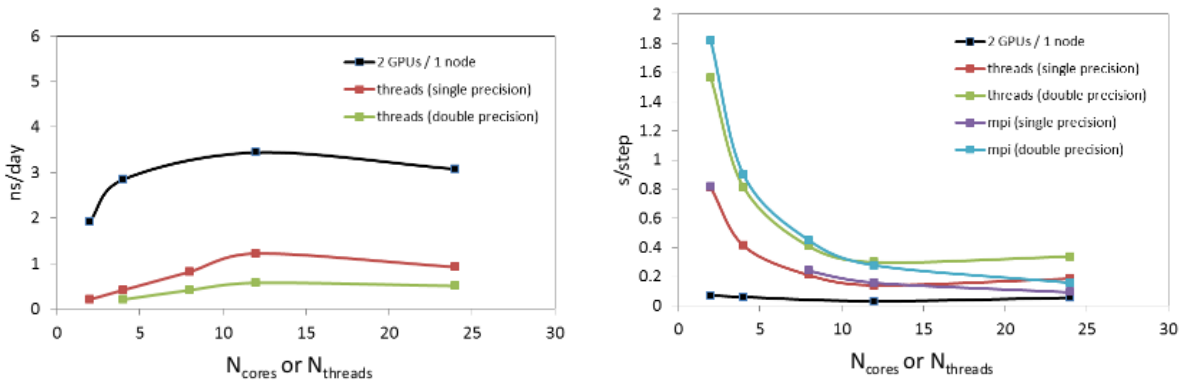

(a) (b)

Figure 7: Performance of GROMACS 4.6 with GPUs, single precision (without GPUs), and double precision (without GPUs) versions for Arp 2/3-VCA.

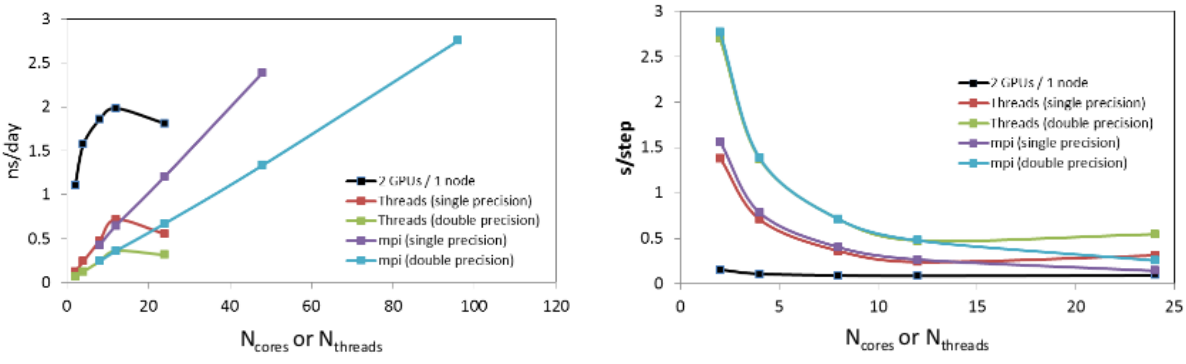

(a) (b)

Figure 8: Performance of GROMACS 4.6 with GPUs, single precision (without GPUs), and double precision (without GPUs) versions for Arp 2/3-2VCA-actin.

4.1 NAMD

NAMD was tested only for GPU accelaration over different number of nodes. The scaling was poor and the results are presented in Figure 4.1.

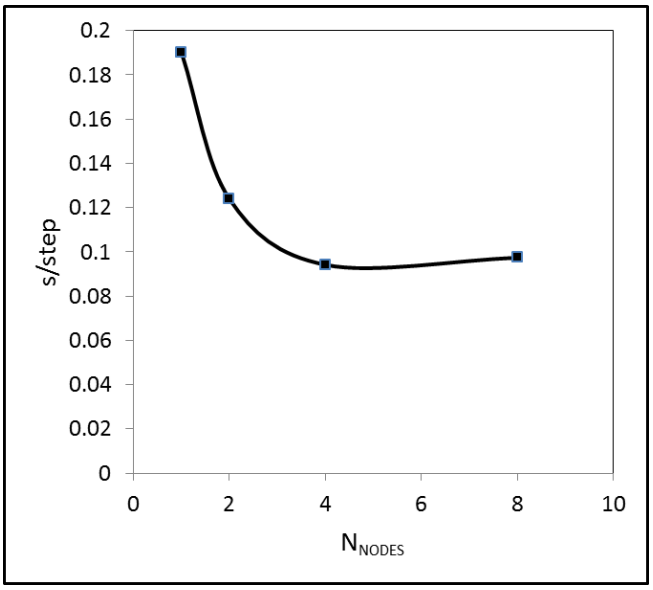

Figure 9: s/step for the Arp 2/3-VCA simulations using NAMD 2.9.

5 Conclusions

The benchmark studies presented herein show an impressive speed-up upon GPU accelaration in the case of GROMACS 4.6. The optimal usage for the under study systems was a 2CPUs/2GPUs per node scheme with an impressive performance of approximately 12 ns/day in eight nodes for a system of more than 0.5 M atoms. Without GPUs, both systems scaled excellent up to 96 cores for the mpi (double and single precision). For the Arp2/3-2VCA-actin system, in 48 cores (without GPU acceleration), 2.4 ns/day were achieved with GROMACS 4.6 mpi single precision, whereas 1.3 ns/day were achieved with GROMACS 4.6 mpi doublem precision. Results from the comparison between GROMACS 4.9 single and double presicion show the same scalability between the two versions, but with double precision being two times slower than single. A thread-enabled versus mpi-enabled version was also tested for both systems under investigation and showed minor differences, apart from the case of 24 threads per node that seemed to not scale well. NAMD 2.9 performed poorly in the present benchmark study.

6 Appendix

6.1 GROMACS

```
; VARIOUS PREPROCESSING OPTIONS
; Preprocessor information: use cpp syntax.
; e.g.: -I/home/joe/doe -I/home/mary/roe
include                  = -I../top
; e.g.: -DPOSRES -DFLEXIBLE (note these variable names are case sensitive)
define                   = -DPOSRES

; RUN CONTROL PARAMETERS
integrator               = md
; Start time and timestep in ps
tinit                    = 0.0
dt                       = 0.002
nsteps                   = 50000
; For exact run continuation or redoing part of a run
init_step                = 0
; Part index is updated automatically on checkpointing (keeps files separate)
simulation_part          = 1
; mode for center of mass motion removal
comm-mode                = Linear
; number of steps for center of mass motion removal
nstcomm                  = 10
; group(s) for center of mass motion removal
comm-grps                =

; ENERGY MINIMIZATION OPTIONS
; Force tolerance and initial step-size
emtol                    = 0.01
emstep                   = 0.01
; Max number of iterations in relax_shells
niter                    = 0
; Step size (ps^2) for minimization of flexible constraints
fcstep                   = 0
; Frequency of steepest descents steps when doing CG
nstcgsteep               = 1000
nbfgscorr                = 10

; OUTPUT CONTROL OPTIONS
; Output frequency for coords (x), velocities (v) and forces (f)
nstxout                  = 5000
nstvout                  = 5000
nstfout                  = 0
; Output frequency for energies to log file and energy file
nstlog                   = 1000
```

```
nstcalcenergy            = -1
nstenergy                = 1000
; Output frequency and precision for .xtc file
nstxtcout                = 1000
xtc_precision            = 1000
; This selects the subset of atoms for the .xtc file. You can
; select multiple groups. By default all atoms will be written.
xtc-grps                 =
; Selection of energy groups
energygrps               = Protein Water_and_ions

; NEIGHBORSEARCHING PARAMETERS
; nblist update frequency
nstlist                  = 10
; ns algorithm (simple or grid)
ns_type                  = grid
; Periodic boundary conditions: xyz, no, xy
pbc                      = xyz
periodic_molecules       = no
; nblist cut-off
rlist                    = 1.4
; long-range cut-off for switched potentials
rlistlong                = 1.4

; OPTIONS FOR ELECTROSTATICS AND VDW
; Method for doing electrostatics
coulombtype              = PME
rcoulomb-switch          = 0
rcoulomb                 = 1.4
; Relative dielectric constant for the medium and the reaction field
epsilon_r                = 1.0
epsilon_rf               = 1
; Method for doing Van der Waals
vdw-type                 = switch
; cut-off lengths
rvdw-switch              = 0.8
rvdw                     = 1.4
; Apply long range dispersion corrections for Energy and Pressure
DispCorr                 = No
; Extension of the potential lookup tables beyond the cut-off
table-extension          = 1
; Seperate tables between energy group pairs
energygrp_table          =
; Spacing for the PME/PPPM FFT grid
fourierspacing           = 0.14
; FFT grid size, when a value is 0 fourierspacing will be used
fourier_nx               = 0
```

```
fourier_ny               = 0
fourier_nz               = 0
; EWALD/PME/PPPM parameters
pme_order                = 4
ewald_rtol               = 1e-05
ewald_geometry           = 3d
epsilon_surface          = 0
optimize_fft             = yes


; OPTIONS FOR WEAK COUPLING ALGORITHMS
; Temperature coupling
tcoupl                   = Berendsen
nsttcouple               = -1
nh-chain-length          = 1
; Groups to couple separately
tc-grps                  = Protein Water_and_ions
; Time constant (ps) and reference temperature (K)
tau_t                    = 0.5 0.5
ref_t                    = 310 310
; Pressure coupling
Pcoupl                   = Berendsen
Pcoupltype               = isotropic
nstpcouple               = -1
; Time constant (ps), compressibility (1/bar) and reference P (bar)
tau_p                    = 1
compressibility          = 4.5e-5
ref_p                    = 1.01325
; Scaling of reference coordinates, No, All or COM
refcoord_scaling         = No
; Random seed for Andersen thermostat
andersen_seed            = 815131

; SIMULATED ANNEALING
; Type of annealing for each temperature group (no/single/periodic)
annealing                = no
; Number of time points to use for specifying annealing in each group
annealing_npoints        =
; List of times at the annealing points for each group
annealing_time           =
; Temp. at each annealing point, for each group.
annealing_temp           =

; GENERATE VELOCITIES FOR STARTUP RUN
gen_vel                  = yes
gen_temp                 = 300
gen_seed                 = 173529
```

```
; OPTIONS FOR BONDS
constraints              = all-bonds
; Type of constraint algorithm
constraint-algorithm     = Lincs
; Do not constrain the start configuration
continuation             = no
; Use successive overrelaxation to reduce the number of shake iterations
Shake-SOR                = no
; Relative tolerance of shake
shake-tol                = 1e-04
; Highest order in the expansion of the constraint coupling matrix
lincs-order              = 4
; Number of iterations in the final step of LINCS. 1 is fine for
; normal simulations, but use 2 to conserve energy in NVE runs.
; For energy minimization with constraints it should be 4 to 8.
lincs-iter               = 1
; Lincs will write a warning to the stderr if in one step a bond
; rotates over more degrees than
lincs-warnangle          = 30
; Convert harmonic bonds to morse potentials
morse                    = no
```

6.2 NAMD

```
# Molecular system
structure ionized.psf
coordinates equil2_out.pdb
bincoordinates   equil2_out.coor
binvelocities    equil2_out.vel
extendedSystem   equil2_out.xsc

set outputname equil3b_out ;

#~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~#
firsttimestep    0
set temperature  310

#~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~#
# Input
paraTypeCharmm    on
parameters        par_all27_prot_lipid.prm
#temperature       $temperature

#~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~#
# Force-Field Parameters
```

```
exclude        scaled1-4
1-4scaling     1.0
cutoff         12.
margin         10.0
switching      on
switchdist     10.
pairlistdist   13.5

#~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~#
# Integrator Parameters
timestep             2.0
rigidBonds           all
useSettle            on
nonbondedFreq        1
fullElectFrequency   2
stepspercycle        10

#~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~#
#Harmonic constraints
constraints        on
consexp            2
conskcol           B
consref            equil2_out.pdb
conskfile          equil2_out.pdb
constraintScaling  0.5

#~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~#
# Constant Temperature Control
langevin            on
langevinDamping     5 # damping coefficient in ps^{-1}
langevinTemp        310
langevinHydrogen    off

wrapWater   on
wrapAll     on

#~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~#
# PME (for full-system periodic electrostatics)
#use numbers with small integer factor: 2,3,5
PME            yes
PMEGridSizeX    162
PMEGridSizeY    162
PMEGridSizeZ    162

#~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~#
# Constant Pressure Control
langevinPiston        on
```

```
langevinPistonTarget  1.01325
langevinPistonPeriod  1000.
langevinPistonDecay   500. # damping time constant in fs
langevinPistonTemp    310

#~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~#
# Output - Write every 2ps
outputName      $outputname
binaryoutput    on
restartfreq     1000
dcdfreq         1000
xstFreq         1000
outputEnergies  1000
outputPressure  1000

#~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~#
# Molecular Dynamics - Equilibrate the system for 3ns
reinitvels $temperature
run        50000
```

References

[1] E. D. Goley and M. D. Welch, "The arp2/3 complex: an actin nucleator comes of age," *Mol. Cell Biol.*, vol. 7, pp. 713–726, 2006.

[2] B. Alberts, A. Johnson, J. Lewis, M. Raff, K. Roberts, and P. Walter, *Molecular Biology of the Cell*. Garland Science, fifth edition ed., 2008.

[3] R. C. Robinson, K. Turbedsky, D. Kaiser, J. B. Marchand, H. Henry, S. Choe, and T. D. Pollard, "Crystal structure of arp2/3 complex," *Science*, vol. 294, pp. 1679–1684, 2001.

[4] C. C. Beltzner and T. D. Pollard, "Pathway of actin filament branch formation by arp2/3 complex," *J. Biol. Chem.*, vol. 283, pp. 7135–7144, 2008.

[5] R. E. Mahaffy and T. D. Pollard, "Kinetics of the formation and dissociation of actin filament branches mediated by arp2/3 complex.," *Biophys. J.*, vol. 91, pp. 3519–3528, 2006.

[6] A. Rodal, O. Sokolova, D. B. Robins, K. M. Daugherty, S. Hippenmeyer, H. Riezman, N. Grigorieff, and B. L. Goode, "Conformational changes in the arp2/3 complex leading to actin nucleation.," *Nat. Struct. Biol.*, vol. 12, no. 1, pp. 26–31, 2005.

[7] P. Dallhaimer and T. Pollard, "Molecular dynamics simulations of arp2/3 complex activation.," *Biophys. J.*, vol. 99, pp. 2568–2576, 2010.

[8] J. Pfaendtner and G. A. Voth, "Molecular dynamics simulation and coarse-grained analysis of the arp2/3 complex.," *Biophys. J.*, vol. 95, pp. 5324–5333, 2008.

[9] Z. Cournia, G. Ullmann, and J. Smith, "Differential effects of cholesterol, ergosterol and lanosterol on a dipalmitoyl phosphatidylcholine membrane: A molecular dynamics simulation study.," *J. Phys. Chem. B*, vol. 111, no. 7, pp. 1786–1801, 2007.

[10] G. Patargias, H. Martay, and W. B. Fischer, "Reconstructing potentials of mean force from short steered molecular dynamics simulations of vpu from hiv-1.," *J. Biomol. Struct. Dyn.*, vol. 1, pp. 1–12, 2009.

[11] G. Patargias, T. Barke, A. Watts, and W. B. Fischer, "Model generation of viral channel forming 2b protein bundles from polio and coxsackie virus.," *Mol. Mem. Biol.*, vol. 26, no. 5, pp. 1–13, 2009.

[12] G. Patargias, P. Bond, D. D. Deol, and M. Sansom, "Molecular dynamics simulations of glpf in a micelle vs. in a bilayer: Conformational dynamics of a membrane protein as a function of environment.," *J. Phys. Chem. B*, vol. 109, no. 1, pp. 575–582, 2005.