RUMD User Manual [Version 3.6]

The RUMD development team: Thomas B. Schrøder, Nicholas P. Bailey,
Trond Ingebrigtsen, Jesper Schmidt Larsen, Claire Lemarchand,
Lasse Bøhling, Heine Larsen, Lorenzo Costigliola, Ulf R. Pedersen

© May 2019

Contents

1   Preliminaries and general information

This document is meant to be a more complete documentation of the functionality available to a user of RUMD than is available in the tutorial, without describing the internal workings in detail. It is assumed that the reader has studied the tutorial; there is little point in duplicating what is explained there. The overall aim is that having studied this manual, the user should have fairly complete knowledge of what RUMD can do, though not how it does it.

1.1 Availability and latest version

The latest version of this manual, and of the RUMD source code, are available at the RUMD website `http:rumd.org`. The version of a given installation of RUMD can be found in Python after importing the rumd module (see below) via `rumd.GetVersion()`. If using RUMD compiled with source code checked out from the subversion repository then the revision number can be found via `rumd.Get_SVN_Revision()`.

1.2 RUMD paths and modules

When RUMD is correctly installed, the relevant paths should be already in Python's search path and the import statements should work. When compiling from source without installing, the relevant paths which need to be included in Python's search path (via `import sys` and `sys.path.append(...)`) are `<path-to-source>/lib` and `<path-to-source>/Python`. Note this changed in version 3.1; before that there were three paths that should be included.

From a software point of view, RUMD consists of two Python modules[1]. The first is called `rumd`, is implemented in C++ (which in turn makes calls to device functions which run on the GPU) and represents the bulk of the code. The second is called `rumd.Simulation`, which is a pure Python module implementing the `Simulation` class. Both modules need to be imported near the start of a user script, for example as follows

```
import rumd
from rumd.Simulation import Simulation
```

After this the `Simulation` class is available without prefix and classes in the `rumd` module can be accessed by prefixing with `rumd.`, for example to create a Lennard-Jones pair potential object:

```
pot = rumd.Pot_LJ_12_6(...)
```

One can also use the `import *` device to import all names from `rumd` into the current name space, or `from rumd import Pot_I` say, to import some names. This can be convenient but it is not considered good Python practice because of potential name clashes and the idea that it should be clear to someone reading the code which module a function or class comes from. For Python access to the analysis tools it is necessary to import the tools module:

```
import rumd.Tools
```

1.3 Basic RUMD scripts

A minimal RUMD script consists of four steps. (1) Create a `Simulation` object from a configuration file. (2) Create a potential object and attach it to the simulation object. (3) Create an integrator object and attach it to the simulation object. (4) Call `Run(nSteps)` on the simulation object. See the tutorial for examples.

1.4 Acquiring start-configurations

The program `rumd_init_conf` can be used to create a configuration file named `start.xyz.gz`. Type

```
rumd_init_conf --lattice=fcc --cells=6,6,6
```

or

```
rumd_init_conf -lfcc -c6,6,6
```

to make a face centered cubic (fcc) lattice with $6{\times}6{\times}6$ unit cells. Type

---

[1] Three if the `rm` module is included, which provides a Python interface to the analysis tools.

```
rumd_init_conf -c10 -N800,200 -r1.2 -m -o KobAndersen.xyz.gz
```

to make a configuration for the binary Kob-Andersen system on a (default) simple cubic (sc) lattice.

Below is a list of options for `rumd_init_conf`

rm Print information on using the program.

rm Hide program output.

rm Set lattice type. Default: sc. Allowed values (n="particles in unit cell"):

    sc Simple cubic (n=1).

    rp Randomly packed (n=1).

    bcc Body centered cubic (n=2).

    fcc Face centered cubic (n=4).

    hcp Hexagonal close packed (n=4).

    hex Hexagonal layers in xy-planes (n=2).

    dc Diamond cubic lattice (n=8).

    NaCl Rock salt lattice. (n=2x4).

    CsCl Caesium Chloride lattice (n=2x1).

rm Set number of unit cells. Default: 5.

rm Reset particle types in lattice. Note: The sum of particles must be < or = lattice sites.

rm Set masses of types. Default: 1.

rm Make positions of particles on lattice sites random.

rm Set number density.

rm Change length of box vectors (skip resetting when <0).

rm Do not allow distances shorter than distance given.

rm Temperature of random velocity vectors.

rm Seed for pseudo random numbers.

rm *.xyz or *.xyz.gz output file.

In addition, some configurations are supplied with the source code in the directory Conf. Configurations for single component liquids are in the sub-directory `Conf/SingleComponentLiquid`, for example `LJ_N1000_rho0.80_T0.80.xyz` is a configuration with 1000 particles and density 0.80 which has been at least partly equilibrated at temperature 0.80 using the Lennard-Jones potential. One should not rely on these configurations being properly equilibrated however. In particular there is no problem using them for a different potential, temperature and density than those specified in the filename. The directory `Conf/Mol` contains configurations for molecular systems, along with the associated topology files (`.top`).

1.5 Units

Input and output of RUMD is done in reduced units. E.g. temperatures are not given in Kelvin, Celsius or Fahrenheit but in units of ε/k_B. The table below gives conversion from reduced units to two choices of unit systems.

Table 1: Examples of unit systems

		Molar units	Shifted SI units
Length	σ	Å	nm
Energy	ε	kJ/mol	zJ
Mass	m	u	yg
Boltzmann's const.	k_B	0.0083145 kJ/(mol K)	0.013806 zJ/K
Coulomb's const.	k_e	1389.3 Å kJ/(mol e^2)	230.71 nm zJ/e^2
Time	$\sigma\sqrt{m/\varepsilon}$	0.1 ps	ps
Temperature	ε/k_B	120.27 K	72.432 K
Velocity	$\sqrt{\varepsilon/m}$	km/s	km/s
Force	ε/σ	16.605 pN	pN
Pressure	ε/σ^3	1.6605 GPa	GPa
Charge	$\sqrt{\varepsilon\sigma/k_e}$	0.026829 e	0.065837 e
Mass density	m/σ^3	1.6605 g/ml	g/ml
Diffusion	$\sigma\sqrt{\varepsilon/m}$	0.001 cm^2/s	0.01 cm^2/s

Values in standard SI units: $u = 1.66053892 \times 10^{-27}$ kg; $e = 1.602176565 \times 10^{-19}$ C; Å $= 10^{-10}$ m; 1 mol $= 6.02214129 \times 10^{23}$; $k_B = 1.3806488 \times 10^{-23}$ J/K; $k_e = \frac{1}{4\pi\varepsilon_0} = 8.9875517873681764 \times 10^9$ J m/C^2. SI prefixes: yocto, y$=10^{-24}$; zepto, z$=10^{-21}$; femto, f$=10^{-15}$; pico, p$=10^{-12}$; nano, n$=10^{-9}$; centi, c$=10^{-2}$; kilo, k$=10^3$; mega, M$=10^6$; giga, G$=10^9$. Note the use of non-SI units for densities: g/ml$=1000$ kg/m^3. Values in the table are truncated to 5 significant digits.

1.6 Precision for floating-point arithmetic

Like many GPU-based simulation codes, most of the RUMD kernels use single-precision floating point arithmetic. Code on the host (the CPU) generally uses double precision; this is important particularly when summing up the particle contributions to the total potential or kinetic energy. The effect of single precision can be seen in the drift of conserved quantities due to round-off error. These include the total momentum, which therefore needs to be periodically reset to zero, and the total energy in NVE simulations. Even in double precision drift will be observed in long runs; the simplest fix is to use a thermostat (NVT). A potentially more serious limitation is the convergence of energy minimization to so-called "inherent states", although we have not investigated how severe this is in practice.

1.7 Maximum system sizes

In version 3.0 a new neighbor-list array structure is used. For systems larger than around 10000 particles space is allocated for a number of neighbors estimated from the global density. If this turns out to be insufficient the neighbor-list is re-allocated, increasing the allocated space by a factor of 2. For a simple Lennard-Jones system with a cutoff of 2.5σ and a skin of 0.5 σ the memory requirements are of order 800 bytes per particle of which 600 are for the neighbor-list (assuming a density of 1 σ^{-3}). Explicit, though limited, testing on a GTX 780 Ti (with 3072 MB RAM) is consistent with this, allowing simulation of systems up to $N \sim 3.5 \times 10^6$. Using a cutoff of 1.5 σ (still with skin of 0.5 σ) requires roughly a factor of two less space, and explicit testing gives a limit $N \sim 7.3 \times 10^6$. In these tests 1000 time steps were simulated on a Lennard-Jones fcc crystal with temperature 1.5. The autotuner was not used. The autotuner makes copies of the system so requires more memory, although a separate neighbor-list is not created. Using more than one pair potential will have a significant effect because each pair potential object has its own neighbor-list. Molecular systems also use more memory because they require an exclusion list which has the same structure as the neighbor-list and a hard-coded size of 100 excluded neighbors per particle, which is comparable to the neighbor-list size, leading to a size limit of order half that for

atomic systems. On GPUs with compute capability less than 3.0, another limit is relevant, namely the number of thread blocks cannot exceed 65535; with 32 particles per block this limit is 2097120 particles.

1.8 On the use of the word "block" in RUMD

For historical reasons the word "block" is used in two different senses in RUMD. This should not cause too much confusion, since the contexts are quite different. Within the CUDA kernels, "block" refers to a thread-block in the usual CUDA sense. More relevant for the user is the sense of dividing the simulation run into (say) 100 blocks, such that there is one output file of a given category for each block.

1.9 How to cite RUMD

If you use RUMD in published work, please cite it by giving the RUMD homepage `http://rumd.org`. A paper describing the code in detail is available in SciPost Physics. A BibTeX snippet for this paper is:

```
@Article{Bailey-SP-2017-RUMD,
author    = {Nicholas Bailey and Trond Ingebrigtsen and Jesper Schmidt Hansen
            and Arno Veldhorst and Lasse B{\o}hling and Claire Lemarchand and Andreas Olsen
            and Andreas Bacher and Lorenzo Costigliola and Ulf Pedersen and Heine Larsen
            and Jeppe Dyre and Thomas Schr{\o}der},
title     = {{RUMD}: A general purpose molecular dynamics package optimized to utilize {GPU}
            hardware down to a few thousand particles},
journal   = {{SciPost} Physics},
year      = {2017},
volume    = {3},
number    = {6},
pages     = {038},
doi       = {10.21468/scipostphys.3.6.038},
publisher = {Stichting {SciPost}},
timestamp = {2020-07-16},
}
```

2 Pair potentials

An overview of the available potentials is given in table 2 along with the arguments needed for each pair of types.

2.1 Number of types

Previously there was a hard-coded limit to the maximum number of types, set to 16 by default. In RUMD 3.0 this limit was removed, because parameters for potentials are now stored differently.

2.2 Cut-offs

For almost all pair potentials the constructor takes a required argument `cutoff_method` which is one of `rumd.ShiftedPotential`, `rumd.ShiftedForce` or `rumd.NoShift`. The traditional choice is shifted-potential, but in our group we use shifted-force cutoffs more and more. Note that neither the cut-off method nor the cut-off distance may be specified for Dzugutov potential, which has an exponential cut-off built into the potential itself.

Table 2: List of pair potentials implemented in RUMD.

Name of Potential	Definition $v(r_{ij})$	Parameters/Constants	SetParams	Name in code
Lennard–Jones (12,6)	$4\cdot\epsilon_{ij}\left\{\left(\frac{\sigma_{ij}}{r_{ij}}\right)^{12}-\left(\frac{\sigma_{ij}}{r_{ij}}\right)^{6}\right\}$	$\epsilon_{ij},\sigma_{ij}$	$(i,j,\sigma_{ij},\epsilon_{ij},R_{cut})$	`Pot_LJ_12_6`
Lennard–Jones (m,n)	$\frac{\epsilon_{ij}}{(m-n)}\left\{n\cdot\left(\frac{\sigma_{ij}}{r_{ij}}\right)^{m}-m\cdot\left(\frac{\sigma_{ij}}{r_{ij}}\right)^{n}\right\}$	$\epsilon_{ij},\sigma_{ij},m,n$	$(i,j,\sigma_{ij},\epsilon_{ij},R_{cut})$	`Pot_gLJ_m_n`
Lennard–Jones (12,6)	$\epsilon_{ij}\left\{\left(\frac{\sigma_{ij}}{r_{ij}}\right)^{12}-2\cdot\left(\frac{\sigma_{ij}}{r_{ij}}\right)^{6}\right\}$	$\epsilon_{ij},\sigma_{ij}$	$(i,j,\sigma_{ij},\epsilon_{ij},R_{cut})$	`Pot_gLJ_m12_n6`
Lennard–Jones Gaussian	$\epsilon_{ij}\left\{\left(\frac{\sigma_{ij}}{r_{ij}}\right)^{12}-2\left(\frac{\sigma_{ij}}{r_{ij}}\right)^{6}-\epsilon_0\cdot\exp\left[-\left(\frac{r-r_0}{2\cdot\sigma_0}\right)^{2}\right]\right\}$	$\epsilon_{ij},\sigma_{ij},\epsilon_0,\sigma_0,r_0$	$(i,j,\sigma_{ij},\epsilon_{ij},R_{cut},\epsilon_0,\sigma_0,r_0)$	`Pot_LJ_12_6_Gauss`
Gaussian core	$\epsilon_{ij}\cdot\exp\left[\left(\frac{-r_{ij}}{\sigma_{ij}}\right)^{2}\right]$	$\epsilon_{ij},\sigma_{ij}$	$(i,j,\sigma_{ij},\epsilon_{ij},R_{cut})$	`Pot_Gauss`
Buckingham (exp,six)	$\epsilon_{ij}\left\{\frac{6}{(\alpha-6)}\exp\left[\alpha\left(1-\frac{r_{ij}}{r_m}\right)\right]-\frac{\alpha}{(\alpha-6)}\left(\frac{r_m}{r_{ij}}\right)^{6}\right\}$	ϵ_{ij},r_m,α	$(i,j,r_m,\epsilon_{ij},R_{cut})$	`Pot_Buckingham`
Inverse Power Law, n = 12	$\epsilon_{ij}\left(\frac{\sigma_{ij}}{r_{ij}}\right)^{12}$	$\epsilon_{ij},\sigma_{ij}$	$(i,j,\sigma_{ij},\epsilon_{ij},R_{cut})$	`Pot_IPL_12`
Inverse Power Law, n = 18	$\epsilon_{ij}\left(\frac{\sigma_{ij}}{r_{ij}}\right)^{18}$	$\epsilon_{ij},\sigma_{ij}$	$(i,j,\sigma_{ij},\epsilon_{ij},R_{cut})$	`Pot_IPL_18`
Inverse Power Law, arbitrary n	$\epsilon_{ij}\left(\frac{\sigma_{ij}}{r_{ij}}\right)^{n}$	$\epsilon_{ij},\sigma_{ij},n$	$(i,j,\sigma_{ij},\epsilon_{ij},R_{cut},n)$	`Pot_IPL_n`
IPL6/LN	$\epsilon_{ij}\left(\frac{\sigma_{ij}}{r_{ij}}\right)^{6}\ln\left(\frac{r_{ij}}{\sigma_{ij}}\right)^{-1}$	$\epsilon_{ij},\sigma_{ij}$	$(i,j,\sigma_{ij},\epsilon_{ij},R_{cut})$	`Pot_IPL6_LN`
Dzugutov	$A\cdot\left(r^{-n}-B\right)exp\left(\frac{c}{r_{ij}-a}\right)+B\cdot exp\left(\frac{d}{r_{ij}-b}\right)$	A,B,a,b,c,d	(i,j,A,B,a,b,c,d)	`Pot_Dzugutov`
Shifted-force Coulomb	ϵ_{ij}/r_{ij} [+force-shifting term]	ϵ_{ij}	$(i,j,\epsilon_{ij},R_{cut})$	`Pot_ShiftedForceCoulomb`

2.3 Attaching potential objects to the simulation object

For a simple system with a single pair potential, one typically uses `rumd.AddPotential(pot)`, where `pot` is the potential object. The old method `SetPotential` has been deprecated in Version 3.5. The difference was that `SetPotential` also cleared any previously added potentials. If that is necessary, one can simply supply the argument `clear=True` to `AddPotential`. Potentials for intra-molecular forces (bond, bond-angles, and dihedral angles) were previously added via special methods but as of Version 3.5 they work like any other potential; one creates an object, calls `SetParams`, and then `AddPotential`. These potentials will be described in Section 5.

2.4 General methods for pair potentials

Constructor As mentioned above, almost all pair potentials take a constructor argument specifying the cutoff method. Some also take additional arguments, which describe aspects of the potential which are fixed for all types. An example is the general inverse power law potential `Pot_IPL_n`:

```
pot = rumd.Pot_IPL_n(13.6, rumd.ShiftedPotential)
```

Another example is the generalized Lennard-Jones potential `Pot_gLJ_m_n` which takes two additional arguments, the exponents m and n.

SetParams(...) Set the parameters for a given pair of types (described in the tutorial). The exact arguments depend on the pair potential. To find out what they are within python, type `help(rumd.Pot_LJ_12_6.SetParams)` (for example). Typically the first two arguments are the two types `i` and `j` involved in a given interaction. New in version 2.1 is that Newton's third law is assumed by default, so only it is no longer necessary to call `SetParams` both for types 0 and 1 and then 1 and 0, say.

rm Set the string which identifies a potential (in for example the energies file if the potential is being computed in addition to the actual simulation potential, or is one term of the latter).

rm Get the ID string for this potential.

rm If `v` is false, messages to standard output are suppressed.

SetNbListSkin(skin) Set the skin for the neighbor list (the width of the extra shell of neighbors which allow some time steps to be taken before re-building).

GetNbListSkin() Get the skin for the neighbor list.

rm Pass `False` to disable symmetric interactions between unlike types by default. Pass `True` to enable (default: `True`).

rm Return the potential energy contribution from this pair potential. This will cause a kernel launch; it does not use the results from a previous call to `CalcF`.

rm Write the pair potentials and pair forces (for all type-combinations) to a file (whose name consists of `potentials_` plus the ID string, with a `.dat` extension). A simulation box is passed so that the appropriate range of r values is known. It is more convenient, however, to call `WritePotentials` method in `Simulation`, which will call it on all attached pair potentials and automatically pass the simulation box.

3 Other potentials

3.1 Coulomb

The Coulomb potential is on the face of it a pair potential, but requires special consideration because its long range means that generally a straightforward implementation as a pair potential is not appropriate both for performance reasons and due to issues related to periodic boundary conditions. Typically an Ewald-based method is used, such as particle-mesh Ewald (PME). We plan to implement PME but have implemented a direct pair-sum using a shifted force cut-off, which has been shown to give reasonable accuracy and performance in bulk ionic systems with a cutoff of order 5-6 inter-particle spacings [Hansen *et al.*, J. Phys. Chem. B, 116, 5378, 2012]. This implementation is the pair potential `Pot_ShiftedForceCoulomb`. The user cannot choose the cutoff method here: Shifted-force is enforced. The cut-off itself is chosen by the user; note that because of the large number of neighbors the memory use is significant.

3.2 Many-body potential for metals: Effective Medium Theory

The effective medium theory approach in metallic interactions has been developed by researchers at the Technical University of Denmark. The potential is described in "A semi-empirical effective medium theory for metals and alloys", K. W. Jacobsen et al, Surface Science, Vol. 366, pp394-402 (1996). Parameters for six metals are given there using eV as the energy unit and the bohr radius as the length unit. The following snippet shows how to set up a potential for pure gold using the parameters from the 1996 paper but with lengths converted into Å.

```
bohr   = 0.5291772
E0     = -3.80
s0     = 3.00 * bohr
V0     = 2.321
eta2   = 1.674 / bohr
kappa  = 2.873 / bohr
Lambda = 2.182 / bohr
n0     = 0.00703 / (bohr*bohr*bohr)


pot = rumd.EMT_Potential()
# set parameters for particle type 0 to those for Au
pot.SetParams(0, E0, s0, V0, eta2, kappa, Lambda, n0)
```

Note that alloys can be simulated without any new parameters.

3.3 TetheredGroup

This potential allows you you identify a group of particles which are to be "tethered" to a set of so-called "virtual lattice sites" by harmonic springs. The particles are currently identified by type: A list of types is passed to the constructor, and all particles of any of those types are considered part of the group[2]. The spring constant is the same for all particles in the group, and the virtual lattice sites can be displaced using the Move function (typically inside a method to be called using the "run-time action" mechanism (see section 10.3; this can be used to implement constant-speed boundary-driven shearing, for example.

```
potWall1 = TetheredGroup(solidAtomTypes=[0], springConstant=500.0)
sim.AddPotential(potWall1)
```

---

[2]This typically means that extra types have to be introduced which are otherwise identical to existing types, just for identifying sets of particles for tethering. A more general particle selection mechanism will be implemented in a future version.

```
potWall1.SetDirection(2) # z direction (default is 0 ie x)
potWall1.Move(0.1) # move the virtual lattice sites a common amount
```

4 Integrators

This section provides details on the available integrators listed in the RUMD tutorial. We start with some common methods available on all integrators (listing only those of interest to the user).

4.1 Common methods for integrators

GetTimeStep Return the current time step

SetTimeStep Set the time step

GetKineticEnergy Return the total kinetic energy (i.e. not per particle). In Leap-Frog-type integrators, this requires either calculating a half time step to get the velocities synchronous with the positions, or a full time step to get the "forward" velocities, and taking an average of the kinetc energy based on the backward and forward velocities. For the SLLOD integrator this is not required as the kinetic energy is a conserved quantity.

4.2 IntegratorNVE, IntegratorNVT

`IntegratorNVE` realizes Newtonian NVE dynamics via the Leap-frog algorithm whereas `IntegratorNVT` realizes Nose-Hoover NVT dynamics [S. Nose, J. Chem. Phys. 81, 511 (1984)]. The implementation of the Nose-Hoover NVT dynamics is detailed in [S. Toxvaerd, Mol. Phys. 72, 159 (1991)]. In this implementation, the Leap-frog algorithm naturally appears by setting the thermostat variable to zero. The former two integrators are thus based on the same common C++ class `IntegratorNVT` where `IntegratorNVE` is simply a wrapper for this class with fixed parameters. Examples of use may be found in the tutorial but are also listed below. Nose-Hoover NVT dynamics can be run with

```
# Nose-Hoover NVT dynamics with time step dt = 0.0025 and temperature T = 1.0
itg = rumd.IntegratorNVT(timeStep=0.0025, targetTemperature=1.0, thermostatRelaxationTime=0.2)
```

The third constructor argument is the relaxation time of the thermostat, which can also be set using

```
itg.SetRelaxationTime(0.2)
```

A linear temperature ramp can also be associated with the NVT dynamics via

```
itg.SetTemperatureChangePerTimeStep(0.01)
```

This causes `targetTemperature` to be changed every time step with $\Delta T = 0.01$ until the simulation is complete. Alternatively, Newtonian NVE dynamics can be run with

```
# Newtonian NVE dynamics with time step dt = 0.0025.
itg = rumd.IntegratorNVE(timeStep=0.0025)
```

Note: No correction of the energy E is performed, and round-off errors will eventually cause the energy to drift over a long simulation. One should always estimate whether this drift is significant for the phenomena under investigation.

4.3   IntegratorNPTAtomic

`IntegratorNPTAtomic` realizes isothermal-isobaric dynamics via Leap-frog algorithm coupled with a thermostat and a barostat. In the current version of RUMD, NPT dynamics is not implemented for molecules. Details on the equations used to implement an isothermal-isobaric ensemble can be found in [Martyna *et al*, J.Chem.Phys. 101, 5, (1994)]. The following is an example of how to use the NPTAtomic integrator.

```
# NPT dynamics with time step dt = 0.0025
# T = 2.0, thermostat relaxation time = 0.4, p = 5.0, barostat relaxation time = 10.0
itg=IntegratorNPTAtomic(timeStep=0.0025, targetTemperature=2.0,
thermostatRelaxationTime=0.4, targetPressure=5.0, barostatRelaxationTime=10.0)
```

Instantaneous volume is not a default output. It is possible to print instantaneous volume in the energy files using `rm`:

```
# print instantaneous volume in the energy files.
sim.SetOutputMetaData("energies", volume=True)
```

4.4   IntegratorNPTLangevin

This integrator realizes the NPT ensemble using different dynamics. The Langevin equation is used both for particles and for the box dynamics. The method is based on N. Grønbech-Jensen and Oded Farago, J. Chem. Phys. 141, 194108 (2014). It has been reformulated as a leap-frog-type algorithm, as outlined in Grønbech-Jensen and Farago, Comput. Phys. Commun. 185, 524-527 (2014) (the latter covers the NVT case, but the equations for the box degree of freedom parallel those for the particles). A slight difference from their version is to define the random forces as integrals over the offset-timesteps (from $t_{n-1/2}$ to $t_{n+1/2}$) instead of taking the mean of two successive non-offset integrals, which is what comes out when straightforwardly reformulating to a leap-frog version.

The user obviously must specify time step, target temperature and target pressure. In addition a friction α must be specified for the particles, and a friction $\tilde{\alpha}$ and inertial constant Q must be specified for the box dynamics. The inertial constant is like a mass, although it has different units. See below for advice on the choice of

By calling SetBarostatMode with arguments rumd.Off, rumd.Iso or rumd.Ansio, the user may disable barostatting (giving thus an NVT ensemble), or isotropic barostatting (default) or anisotropic barostatting. For the latter the box size fluctuates only in one direction, by default the z-direction. The direction may be set by calling `SetBoxFlucCoord`.

```
itg = rumd.IntegratorNPTLangevin(timeStep=dt, targetTemperature=T, friction=0.1,
targetPressure=0.1, barostatFriction=1.e-3, barostatMass=1.e-2);
# make the barostatting anisotropic, specifically in the x-direction
itg.SetBarostatMode(rumd.Aniso)
itg.SetBoxFlucCoord(0)
```

To see the volume fluctuations in the energies files, add the following line:

```
sim.SetOutputMetaData(''energies'', volume=True)
```

There is an important difference between the ensemble generated by this integrator and that generated by IntegratorNPTAtomic. Momentum is not conserved by the dynamics. This means (1) the SetMomentumToZero function, which is called to curb drift in momentum due to round-off error, does nothing; and (2) the ensemble allows non-zero values of the total momentum (although centered around zero due to the friction). In fact this is a more accurate NPT ensemble, whereas the one generated by the other integrators has fixed total momentum

so is technically a different ensemble. A consequence is that the number of degrees of freedom when estimating the temperature is 3N, rather than 3N-3, which is otherwise assumed in rumd. This difference is not handled automatically, so the user must include a line in their script to set the number of degrees of freedom correctly:

```
sim.sample.SetNumberOfDOFs(sim.GetNumberOfParticles()*3)
```

4.4.1 Choosing friction parameters and box inertia

The user has a certain freedom when choosing the atomic friction parameter. The time scale for a particle velocity to decay is m/α where m is the particle mass. A large value of α gives a fast time scale and strong coupling to the heat bath, and distorts single-particle dynamics substantially, but can help equilibrate the system quickly. For minimal distortion of particle dynamics a low value of α can be chosen in which case longer equilibration times could be required. (If the system has a long relaxation time anyway, e.g. a viscous liquid, then this is not an issue).

Next consider the box inertia Q. First note that Q has dimensions mass/length^4 or density/length. By linearizing the box equation of motion around the desired pressure (with corresponding volume V_0) and omitting the stochastic terms we get an approximate equation for a basic "breathing" mode involving volume fluctuations ΔV

$$Q\frac{d^2(\Delta V)}{dt^2} = \frac{K_s}{V_0}\Delta V \tag{1}$$

where K_s is the adiabatic bulk modulus. This has a frequency

$$\omega_{ad} = \sqrt{K_s/(QV_0)} = 2\pi/\tau_{ad}. \tag{2}$$

A choice consistent with that made by the authors in their example is that τ_{ad} should be of the order of the time for sound to travel one inter-particle spacing a,

$$\tau_{ad} = a/c_s = a\sqrt{\rho_m/K_s} \tag{3}$$

where c_s is the sound velocity and ρ_m is the mass density. Combining with Eq. (2) gives the following equation for Q which does not involve K_s:

$$Q = \frac{a^2}{V_0}\frac{\rho_m}{(2\pi)^2} \tag{4}$$

For a Lennard-Jones system with density of order 1 and around 1000 particles this gives $Q \sim 3\times10^{-5}$, as used by Grønbech-Jensen and Farago. Note that the inertial parameter here must be chosen inversely proportional to the system volume.

For box friction parameter $\tilde{\alpha}$, a very small value could lead to underdamped box dynamics, and thus long lived oscillations of the volume, which are not desireable. A natural choice would be one which implies critical damping of the box dynamics, that is the damping time scale matches that of the adiabatic breathing mode, or more precisely,

$$\frac{\tilde{\alpha}}{Q} = 2\omega_{ad} \tag{5}$$

Given the choice of Q above this implies that

$$\tilde{\alpha} = \frac{a}{\pi V_0}\sqrt{\rho_m K_s} \tag{6}$$

Note that here we do need to know the adiabatic bulk modulus, at least roughly. Choosing value ~ 50 for a Lennard-Jones system, with $V_0 \sim 1000$ and $a \sim \rho_m \sim 1$ gives $\tilde{\alpha} \sim 2 \times 10^{-3}$. This does not quite match the

value chosen by Grønbech-Jensen and Farago, who used a value 10^{-4} or a factor of 20 smaller, implying an underdamped dynamics (with a quality factor of 20), although the effective damping was probably smaller due to coupling with the particle dynamics. A recommendation would thus be to use Eq. (6) as an upper bound for $\tilde{\alpha}$ and experiment with reducing it by a factor of order 10.

4.5   IntegratorNVU

This is a new molecular dynamics that instead of conserving the energy E (as in Newtonian dynamics) conserves the total potential energy U. NVU dynamics traces out a geodesic curve on the so-called constant-potential-energy hypersurface Ω, given by ($\mathbf{R} \equiv (\mathbf{r}_1, ..., \mathbf{r}_N)$)

$$\Omega = \{\mathbf{R} \in R^{3N} \mid U(\mathbf{R}) = U_0\}, \tag{7}$$

in which U_0 is a constant. More details on the dynamics and its implementation can be found in T. S. Ingebrigtsen *et al.*, J. Chem. Phys. 135, 104101 (2011); ibid. 135, 104102 (2011); ibid. 137, 244101 (2012). Examples of use are given below

```
# NVU dynamics with (mass-weighted) step size in 3N-dimensional
# configuration space dispLength = 0.01 and potentialEnergyPerParticle = -1.0
itg = rumd.IntegratorNVU(dispLength=0.01, potentialEnergyPerParticle=-1.0)
```

NVU dynamics gives equivalent results to NVE dynamics for most properties, and `potentialEnergyPerParticle` could, for instance, be chosen as the average potential energy of a given state point. There is no drift of the potential energy as this is explicitly corrected by the implementation. `dispLength` (also known as l_0) is similar to the time step Δt of NVE dynamics and cannot be chosen arbitrarily large.

4.6   IntegratorSLLOD

This integrator uses the SLLOD equations of motion are used to simulate simple shear flow. The flow is in the x-direction with a gradient in the y-direction; that is, the further you go in the (positive) y-direction, the faster the streaming velocity (in the positive x-direction). As is typically done we use an isokinetic thermostat, which keeps the kinetic energy constant. Our implementation is based on Pan et al., "Operator splitting algorithm for isokinetic SLLOD molecular dynamics", J. Chem. Phys., 122, 094114 (2005), and conserves the kinetic energy to a high precision. Note that while RUMD in general is coded in single precision, parts of SLLOD are done in double precision. These include the kinetic energy-like quantities which are summed for the isokinetic thermostat. In addition the incrementing of the box-shift is done in double precision on the host in order to allow small strain rates.

   This integrator requires a simulation box that implements a special kind of periodic boundary conditions, so-called Lees-Edwards boundary conditions, implemented using `LeesEdwardsSimulationBox`. Here is the Python code for running a SLLOD simulation:

```
sim = Simulation(...)

# create a Lees-Edwards box based on the existing box
# (can be omitted if using a configuration that came from a
# SLLOD simulation)
le = rumd.LeesEdwardsSimulationBox(sim.sample.GetSimulationBox())
sim.SetSimulationBox(le)

# strain rates below 1.e-6 may not be stable for small box sizes (the box-shift
# per time step must not be swallowed by round-off error)
```

```
itg = rumd.IntegratorSLLOD(timeStep=0.004, strainRate=1.e-5)
sim.SetIntegrator(itg)

# T is the desired temperature; N the number of particles
# -3 due to momentum conservation
# (perhaps should be -4 because of the kinetic energy constraint)
itg.SetKineticEnergy(0.5*T*(3*N-3.))

# xy is the relevant component of the stress (from version 3.0 includes kinetic contribution by default)
# for stress-strain curves, determining viscosity, etc.
sim.SetOutputMetaData("energies", stress_xy=True)

sim.Run(...)
```

Note that the analysis tools that analyze configurations (such as `rumd_rdf` and `rumd_msd`) correctly handle the Lees-Edwards boundary conditions. It is important to realize, however, that the `rumd_msd` only considers displacements transverse to the shearing (x) direction.

4.7   IntegratorMolecularSLLOD

The above integrator, `IntegratorSLLOD`, is for use on atomic systems. It can be used with molecular systems, but there are some issues which arise with such systems, in particular to do with the thermostatting mechanism. To avoid these issues there is a separate integrator class, `IntegratorMolecularSLLOD`, which integrates the molecular formulation of the SLLOD equations of motion. The essential difference is that (1) the streaming velocity used to adjust the changes in position and momentum is evaluated at the center of mass of the molecule containing that atom, and (2) the thermostat constrains the kinetic energy associated with the molecular centers of mass, leaving rotations and other degrees of freedom unthermostatted. The algorithm is a straightforward generalization of the one of Pan et al that was used for the atomic case. Use of `IntegratorMolecularSLLOD` is exactly as for `IntegratorSLLOD`; note that it will give an error if no molecules have been defined (section 5).

4.8   IntegratorIHS

Potential energy minimization is performed with `IntegratorIHS` via Newtonian dynamics (Leap-frog algorithm). A path is followed in configuration space where $\sum_i \mathbf{F}_i \cdot \mathbf{v}_i > 0$. If the latter sum becomes negative $\mathbf{v}_i$ is set equal to zero for all particles i. A simulation is run with

```
# timeStep specifies the time step of Newtonian dynamics
itg = rumd.IntegratorIHS(timeStep=0.0005)
```

The `timeStep` should be chosen smaller than usual MD simulations. After setting the integrator on the simulation object one can call `Run(...)` in the usual way to integrate for a fixed number of steps. A more typical mode of use is to run until some convergence criterion has been reached. Moreover, it is also typical to be running an ordinary MD simulation, and periodically make a copy on which energy minimization is performed. This can be done by using the `IHS_OutputManager` class, described in Section 6.5, which acts as a wrapper around `IntegratorIHS`.

4.9   IntegratorMMC

`IntegratorMMC` implements the Metropolis algorithm for canonical ensemble Monte Carlo simulations [N. Metropolis et al., J. Chem. Phys. 21, 1087 (1953)]. The integrator applies cubic-shaped all-particle trial moves and uses a random-number generator suitable for running efficient simulations on the GPU [GPU Gems 3, Chapter 37]. The integrator may be invoked with

```
# canonical ensemble Monte Carlo simulations
# dispLength = 0.01 is the sidelength of cubic trial displacements
# all particles are moved at once
itg = rumd.IntegratorMMC(dispLength=0.001, targetTemperature=1.0)
```

The acceptance rate of the MC algorithm can be extracted via

```
sim.Run(...)
```

```
# a magic number for the acceptance rate is 50%
ratio = itg.GetAcceptanceRate()
print("MC acceptance ratio: " + str(ratio))
```

5   Molecules

The setting up of simulations involving molecules is described by a basic example in the tutorial. Here we add some further details about the intra-molecular potentials and show how to set up position and topology files for a mixture of molecules.

5.1   Potentials for bonds, bond angles, dihedral angles

The potential describing intra-molecular interactions in RUMD contains three terms: one for bonds, one for bond angles and one for dihedral angles. The bonds can be described either by a harmonic potential or by the Finite Extensible Nonlinear Elastic (FENE) potential. The potential for harmonic bonds reads:

$$U_{\text{bond}}(\vec{r}) = \frac{1}{2} \sum_{\text{bonds}} k_s^{(i)} (r_{ij} - l_b^{(i)})^2, \tag{8}$$

where $k_s^{(i)}$ is the spring constant of bond type i and $l_b^{(i)}$ its zero force bond length. The potential for bond angles is given by:

$$U_{\text{angle}}(\vec{r}) = \frac{1}{2} \sum_{\text{angles}} k_\theta^{(i)} [\cos(\theta) - \cos(\theta^{(i)})]^2, \tag{9}$$

where $k_\theta^{(i)}$ is the angle force constant for angle force type i and $\theta_0^{(i)}$ the zero force angle. The parameters for the harmonic bond and bond angle potentials come from the Generalized Amber Force Field (GAFF) in Wang et al., "Development and testing of a general amber force Field", J. Comp. Chem., 25, 1157 (2004). The potential for the dihedral angles have the following form:

$$U_{\text{dihedral}}(\vec{r}) = \sum_{\text{dihed}} \sum_{n=0}^{5} c_n^{(i)} \cos^n(\phi), \tag{10}$$

where $c_n^{(i)}$ are the torsion coefficients for torsional force type i. The parameters for the dihedral angle potential come from Ryckaert et al., "Molecular dynamics of liquid alkanes", Faraday Disc. Chem. Soc., 66, 95(1978).

The FENE potential can also be used to describe bonds. It is given by:

$$U(\vec{r}) = -\frac{1}{2} k R_0^2 \sum_{\text{bonds}} \ln \left[1 - \left(\frac{r_{ij}}{R_0} \right)^2 \right], \tag{11}$$

where $k = 30$ and $R_0 = 1.5$ (in reduced units). At the moment the constraint method is applicable for molecules with few constraints.

Here is a typical python code for setting up the intra-molecular potential:

```
[set up simulation object and add non-bonding potentials as usual]

# read topology file
sim.ReadMoleculeData("start.top")

# create intramolecular potential objects
pot_harm = rumd.BondHarmonic() # or rumd.BondFENE()
pot_angle = rumd.AngleCosSq()
pot_dihedral = rumd.DihedralRyckaert()

# set their parameters
pot_harm.SetParams(bond_type=0, bond_length=0.39, stiffness=2700, exclude=True)

pot_angle.SetParams(angle_type=0, theta0=2.09, ktheta=537)
pot_angle.SetParams(angle_type=1, theta0=1.85, ktheta=413)

pot_dihedral.SetParams(dihedral_type=0, coeffs=[0, 133, 0, 0, 0, 0])
pot_dihedral.SetParams(dihedral_type=1, coeffs=[0,-133, 0, 0, 0, 0])
pot_dihedral.SetParams(dihedral_type=2, coeffs=[15.5000, 20.3050, -21.9170, \
                                        -5.1150, 43.8340, -52.6070])
# add the potentials
sim.AddPotential(pot_harm)
sim.AddPotential(pot_angle)
sim.AddPotential(pot_dihedral)
```

Note that there can be as many bond-, angle- and dihedral- types as wished. The function `DihedralRyckaert.SetParams()` has two arguments: the dihedral angle type i and a list or array of values of the torsion coefficients from $c_0^{(i)}$ to $c_5^{(i)}$. From version 3.1 it does not matter whether the non-bonding interactions are specified first or the intramolecular interactions. Prior to that it was necessary to specify first the non-bonding interactions, otherwise the exclusions would not be correctly applied. The way to obtain the topology file `start.top` used above will be described in section 5.4 below.

The argument `exclude` to `SetParams` on the harmonic bond indicates whether the existing interactions (likely some variant of ranged interaction like Lennard-Jones) should be excluded. This argument should generally be `True` for harmonic bonds, and `False` for FENE bonds.

Finally there are alternative angle and dihedral-angle potentials. The respective classes are named `AngleSq` and `PeriodicDihedral`. The form of the potential for `AngleSq` is a simple harmonic potential in the angle, $\frac{1}{2}k(\theta - \theta_0)^2$, so the parameters are analogous to those of the `AngleCosSq` case. For `PeriodicDihedral` the interaction energy is a function with four rather than six parameters. NEED REFERENCE TO GET FUNCTIONAL FORM.

5.2 Constraints on bond lengths

As an alternative to using intra-molecular potentials, molecules can be simulated using constraints, whereby a fixed distance between two particles is maintained via a constraint force. The actual implementation is described in [S. Toxvaerd *et al.*, J. Chem. Phys. 131, 064102 (2009)]. The procedure for setting up a rigid bond between two particles is analogous to adding a bond potential between the two particles, i.e.,

```
# the distance between all bonds of type = 0 is length = 0.584
cons_pot = rumd.ConstraintPotential()
cons_pot.SetParams(bond_type=0, bond_length=0.584)
```

```
sim.AddPotential(cons_pot)
```

The constraint forces are calculated by solving a set of non-linear equations. This is achieved by first making them linear and thus involves iteration. The number of times the linear set of equations is iterated can be controlled with

```
# iterate the linear set of equations 5 times
cons_pot.SetNumberOfConstraintIterations(5)
```

The more iterations the more costly the algorithm becomes. Since there is no convergence check, one should verify that the bond lengths are satisfied to within the desired tolerance via

```
sim.Run(...)

# prints the standard deviation of all constrained bond lengths
# with respect to the length values
pot_cons.WritePotential(sim.sample.GetSimulationBox())
```

A good value for the standard deviation would be less than $10^{-5} - 10^{-6}$. More detailed statistics on the bond length distribution may be obtained from the post-processing tool `rumd_bonds` (described below)

Due to an internal constraint on the linear equation solver, there may not be more than 512 constraints (fixed bonds) per molecule. This will be lifted in a future RUMD release. It should also be noted that for each bond constraint one degree of freedom is naively removed (for the purpose of evaluating the kinetic temperature). The implementation does not try estimate whether the system is, for instance, over-specified, but the user can manually specify the number of degrees of freedom via `sim.sample.SetNumberOfDOFs()`.

For linear molecules, for instance chain molecules without side groups which are often used as course-grained models for polymers, an optimized linear solver has been implemented. This solver can be chosen by calling

```
# use optimized solver for linear molecules with constraints
cons_pot.SetLinearConstraintMolecules(True)
```

after creating the `ConstraintPotential` object. This optimization is possible because the system of linear equations for a linear molecule has a simple tridiagonal shape. For this optimized solver to work properly, the bonds in the topology file have to be ordered. Moreover, since the same solver is used for all molecules in the simulation, the solver only gives correct results if *all* molecules containing constraints are linear. This optimization has not been tested for a large range of molecule sizes, but for Lennard-Jones chains consisting of 9 constraint bonds, the simulation time is about 70% of the default linear solver time.

5.3 The topology file

An example of topology file can be found in the subdirectory `Conf/Mol`. It is called `mol.top`. It is a text file which specifies the bonds between the atoms. To define the bonds in a `.top` file we:

1. Specify the beginning of the bond section with the keyword

   ```
   [ bonds ]
   ```

   We will call this a section specifier.

2. Add a comment line starting with `;`, for example,

   ```
   ; My bonds
   ```

3. Specify the bonds. Each bond is described row-wise using four positive entries which are all integers: (i) The first entry is the the molecule index, (ii) the second and third integers are the indices of the particle forming the bond and (iii) the forth integer sets the bond type. This means that if atom 42 is bonded with 55 in molecule 2 with bond type 5 we write

   ```
   2 42 55 5
   ```

   and so forth.

Likewise for angles. Here the section specifier is `[ angles ]` and five entries are needed to specify the angle. The first again indicates what molecule this angle belongs to and the last specifies the angle type. For example, if atom with index 42 is the vertex of the angle and atoms 55 and 78 are the two end-points we can use

```
2 55 42 78 0
```

This angle is found in molecule with index 2 and is of type 0. As you might have guessed, to specify a dihedral angle you need 6 entries since a dihedral is defined via three bonds. The section specifier is `[ dihedrals ]`.

Here is a simple example on how to write a topology file of a single butane molecule based on a united-atomic-unit model, where the methyl (the two CH_3 groups) and methylene (the two CH_2 groups) are represented by a single particle:

```
[ bonds ]
; Bond section - ks = 311 kcal/mol lb = 3.4 ?~E
0 0 1 0
0 1 2 0
0 2 3 0


[ angles ]
; Angle section - ktheta = 63.7 kcal/(mol rad**2) theta0 = 109 degr.
0 0 1 2 0
0 1 2 3 0


[ dihedrals ]
; dihedral/torsion section - RB coefficients
0 0 1 2 3 0
```

5.4 Tools for setting up position and topology files

Writing position (`.xyz`) and topology (`.top`) files from scratch for a mixture of molecules can be a lot of work and it is very error prone to type in all the information. Fortunately, RUMD provides a useful tool to do that from the position and topology files of single molecules:

rumd_init_conf_mol This program allows you to set up the configuration for a molecular system (including molecular mixtures and mixtures of molecules and atoms) from `.xyz` and `.top` files of the individual molecules. The molecules are put by default on a cubic lattice with the appropriate spacing. The files `start.xyz.gz` and `start.top` are written and can be used as start files. Usage:

```
rumd_init_conf_mol [--seed N] [--temperature X] [--density X] [--norandomorientation] \\
[--spacing X] <m1.xyz> <m1.top> <num1> ...
```

The arguments in square brackets are optional.

`--seed N:` to give the seed (integer). The default value is 4622.

--temperature X: to give the temperature (float). The default value is 1.0.

--density X: to give the density (float). The default value corresponds to a spacing of twice the size of the largest molecule.

--norandomorientation: to ask for not randomizing the molecules orientation. By default, the orientations are random.

--spacing X: to give the spacing between two lattice sites (float). The default value is twice the size of the largest molecule. If both density and spacing are given by the user, that corresponding to the largest spacing is used.

m1.xyz: .xyz file for the atoms in molecule one (string).

m1.top: the topology file for molecule one (string). If this is an atom, this string should begin with the word "atom". The atom topology file, which would be empty, does not need to exist. Files containing both molecules and isolated atoms can be created, although RUMD does not yet support simulations of such systems.

num1: the number of molecules of this type (integer).

and so forth.

The atom velocities are chosen from a uniform distribution with zero mean and a variance corresponding to the desired temperature. The mass of the atom types involved in each molecule type can be given in the comment line of the .xyz file of each single molecule. For example, the .xyz file of a molecule containing atom type 0 and 2, with mass 1.0 and 2.4 respectively can have a comment line beginning by

```
numTypes=2 mass=1.0,2.4
```

If this line is not found in the .xyz file of the single molecule, the masses of the atom types declared in this molecule are set by default to 1.0, unless the masses of the same atom types are declared in the .xyz files of other molecules. The molecule centers-of-mass are positioned on a low density simple cubic lattice, and you will most likely need to compress your system to a more appropriate density before running proper simulations. This can be done with the python module RunCompress which gently compresses the system while letting the molecules interact with each other and move into more favourable positions.

Example

Assume you wish to simulate a system composed a mixture of 200 FENE molecules with 5 beads in each and of 100 atoms. You can write the single-molecule topology file of the FENE molecule as:

```
[ bonds ]
; Single FENE chain
0 0 1 0
0 1 2 0
0 2 3 0
0 3 4 0
```

We here let both the molecule index and the bond type be 0, but they need not to be. This file can be saved as, for example, single_fene.top. To specify the relative positions of the beads in the molecule write a single-molecule .xyz file

```
5
numTypes=1 mass=2.0 Single FENE
0 0 0 0
0 0 0 1
0 0 0 2
0 0 0 3
0 0 0 4
```

This corresponds to a simple linear chain of five identical particles of type 0 equally spaced along the z-axis. The mass of these particles is 2.0. This file can be saved as, for example, `single_fene.xyz`. There is no topology file for the single atom. The `.xyz` file for the single atom of type 1 and for which no mass is declared can look like:

```
1
single atom
1 0 0 0
```

It can be saved as, for example, `atom.xyz`. To create the mixture of 200 molecules and 100 atoms:

```
./rumd_init_conf_mol single_fene.xyz single_fene.top 200 atom.xyz atom.top 100
```

This command produces two files called `start.top` and `start.xyz.gz` with all the information needed to start a simulation. The mass of atom type 1 is set to 1.0. As mentioned above the density of the system is probably not correct. To compress and equilibrate the system, a loop consisting of a series of short simulations, increasing density by a small fraction each time, can be used. Such a loop is provided in the module `RunCompress`, which is used as follows:

```
from rumd.RunCompress import RunCompress

# create sim object, read configuration, set potential(s) etc
...

RunCompress(sim, final_density=0.9)
```

Here `final_density` is the desired density (in particles per units volume, not molecules per unit volume) that one wants to achieve before running a simulation proper. Other arguments to `RunCompress` are `num_scale_steps` (default 100) which is the number of times the system is uniformly scaled to get to the final density and `niter_per_scale` (default 1000) which is the number of time steps that are run for each scaling step.

□

6   Output

The core of a molecular dynamics program is the algorithm, which needs to be as fast as possible. But the algorithm is useless without output, and the more control a user has over the output the more useful the software. An additional consideration is that with a very fast code one has to be careful not to drown in data, particularly where configurations are concerned. We use so-called *output managers* to control both how frequently a given kind of data is written (*scheduling*) and precisely what data is included in the files (called the *meta-data*, a specification of what data is in a given output file).

6.1 Output managers

Output is controlled by *output managers*. These are software components (classes) written in C++ or Python which manage the writing of various observables to files. A given output manager manages writing of data to one output file at a time, for example one containing energy-like quantities (energies file), or one containing configurations (trajectory file). The basic output managers are written in C++ but can of course be controlled using the Python interface. The mechanism for handling Python output managers exists in order for users to be able to write their own run-time analysis code in Python without having to recompile the main RUMD software (i.e. the C++ library).

We start by discussing the C++ output managers, in particular the so-called energies output manager and trajectory output manager. In order to avoid data corruption in the event of a long simulation being interrupted, for the purposes of output the simulation run is divided into output "blocks". The output manager starts a new file each time a new block is started, with an incremented block index, thus energies files have names like `energies0000.dat.gz`, `energies0001.dat.gz`, etc. and trajectory files have names `trajectory0000.xyz.gz` etc. (previously, up to version 2.0, `block0000.xyz.gz`: old data is still be readable). Note that the files are written in ASCII to be human readable but then compressed using gzip. These files reside by default in the directory `TrajectoryFiles`, but for standard analysis and post-processing it is not necessary to be aware of them (the analysis tools know where to find them). On the other hand it is straightforward to make a concatenated uncompressed energies file using `zcat TrajectoryFiles/energies* > energies.dat` if needed. The uncompressed file will take up 2-3 times as much disk space as the compressed ones did. As much as possible, analysis scripts should also read directly the compressed files (in Python this is easily done using the `gzip` module) rather than requiring the user to separately uncompress the files.

The output files appear in the directory `TrajectoryFiles` by default. If this directory exists it is renamed with a `.rumd_bak` extension (if the latter exists it is overwritten). From version 2.1.1 on it is possible to disable backup by calling `sim.sample.EnableBackup(False)`. Changing the output directory is done as follows:

```
sim.SetOutputDirectory("TrajFilesRun01")
```

Data files in general should be self-describing; in particular it should be possible to look at the file and know what each column represents. We implement this principle in both energies and trajectory files. Analysis programs should make use of the provided meta-data instead of expecting a certain order of columns.

If not explicitly set by the user via `sim.SetBlocksize()`, the block size will be chosen automatically when `sim.Run()` is called, such that the number of blocks is of order 100-1000 and the block size a power of 2.

6.2 Log-lin scheduling: Linear, logarithmic, limlin and mixed.

The output managers can save items (configurations to a trajectory file or lines to an energies file) at equally spaced intervals (linear saving), logarithmically spaced (in powers of two, plus the zeroth) or a "mixed" scheduling called log-lin. The scheduling for a given output manager is set via

```
SetOutputScheduling(manager, schedule, **kwargs)
```

which was described in the tutorial. The `manager` argument is the name of an output manager, for example "trajectory" or "energies". The `schedule` argument is one of "linear" (which requires a following keyword argument `interval`), "logarithmic" (which can take an optional following keyword argument `base`), "none", "loglin" (which requires keyword arguments `base` and `maxInterval`), or "limlin" (which requires keyword arguments `interval` and `numitems`). Since "linear" and "logarithmic" scheduling are special cases of "loglin", let us consider the meaning of the parameters `base` and `maxInterval`:

base This is the smallest interval of time steps considered. All intervals between writing items are multiples of base.

maxInterval The largest interval allowed between saves. Its purpose is to allow combined linear and logarithmic saving. It should be set to zero to do pure logarithmic saving. If non-zero, the interval between writes will, after increasing initially logarithmically, become fixed at maxInterval.

Some possible combinations are, assuming blockSize=128:

Pure logarithmic saving with base 1 Set base=1, maxInterval=0, and you get time steps 0, 1, 2, 4, 8, 16, 32, 64, 128 within each block. This is equivalent to "logarithmic" scheduling without specifying `base`.

Pure linear saving at interval 4 time steps Set base=4 and maxInterval=4 and you get time steps 0, 4, 8, 12, 16, . . . , 128. This is equivalent to "linear" scheduling with `interval`=4.

"log-lin" Set base=1 and maxInterval=8, and you get time steps 0, 1, 2, 4, 8, 16, 24, 32, 40, 48, 56, . . . , 128.

Note the following restrictions on the values of base and maxInterval: When maxInterval is non-zero it must be a multiple of base. Except for the case of linear saving (base=maxInterval), both base and maxInterval must divide evenly into blockSize.

The scheduling type "limlin" is not a special case og log-lin. It saves linearly spaced with specified interval, starting at the beginning of the block, but only until a total `numitems` saves have been performed. Nothing more is saved until the start of the next block. This can be used when short-time linearly spaced is necessary, as well as averaging over longer times, without having to save the entire trajectory at short intervals.

The meta-data in the comment line of a trajectory file contains an item `timeStepIndex` which gives the number of the time step *within the output block* and a string which looks like `logLin=0,10,0,4,10`. The five items in this list identify both the parameters of a particular log-lin saving scheme and the indexing of the current time step with that scheme. It can be useful occasionally to understand these parameters. A description of each one follows, in the order they appear in the string:

block Identifies which output block we are in. Will therefore be the same for all items written in a given block.

base Smallest interval of time steps considered. It does not change during a simulation run, but is a parameter controlled (perhaps indirectly) by the user.

index Labels items within a block, in that it starts at zero and increments by 1 for each item written. When maxInterval is set to zero (pure logarithmic saving), there is a simple relation between index and nextTimeStep: nextTimeStep = 0 if index = 0 and nextTimeStep = base $\times$ $2^{\text{index-1}}$ otherwise. Note that the block size is given (again for pure logarithmic saving, maxInterval=0) by base $\times 2^{\text{maxIndex-1}}$

maxIndex The value of "index" for the last item in the block. So if the saving scheme is such that there are 8 items per block, maxIndex will have the value 7 (since index starts at zero). Does not change during a simulation run[3]; is set based on parameters specified by the user.

maxInterval The largest interval allowed between saves. This is a user parameter which does not change. There is a more complicated relation between nextTimeStep and index in this case.

6.3 Trajectories and energies

Trajectory and energy files are the primary output of RUMD, so it is worth describing them in some detail. Trajectory files contain configurations, including the types and positions of all particles, and optionally their velocities, forces, and other quantities. Energy files consist of one line of output for a given time step, which lists global quantities, such as the total potential energy, or pressure, or a component of the stress tensor. To keep track of potential changes in the formats for these files there is a format tag `ioformat` which appears in

---

[3]The exception to this is linear saving where the interval does not evenly divide the blockSize.

the meta-data (first line of each file). The current value of `ioformat` is 2. The difference between `ioformat` 1 and 2 involves changes in the meta-data for trajectory (configuration) files. The software and analysis can handle either format transparently.

Controlling what gets written, and the precision, in the energies and trajectory files is done via `SetOutputMetaData` as described in the tutorial. One can specify more than one variable in a given call, for example:

```
sim.SetOutputMetaData("energies", totalEnergy=False, stress_xy=True)
```

A look at the first few lines of an output file (using `zcat TrajectoryFiles/energies* | head`, for example) shows the meta-data. For the energies file and default settings, it might look like the following:

```
# ioformat=2 N=1000 Dt=0.005000 columns=ke,pe,p,T,Ps,Etot,W
0.78297 -6.92308 3.65423 0.522503 0 -6.14011 2.52269
0.779256 -6.9193 3.66705 0.520024 0 -6.14004 2.53585
0.77511 -6.91508 3.68272 0.517257 0 -6.13997 2.55168
...
```

The comment line contains the meta-data, including here the `ioformat`, time interval between lines (for energies files one typically uses linear saving), number of particles (useful for analysis) and the list of symbols identifying the columns. It is not a good idea to rely too much on a particular order of columns; when writing an analysis script in Python for example, it is recommended to read the meta-data and extract the column information so that the potential energy column, for example, is always correctly identified.

For trajectories, we use an extended xyz format, and a given file will contain all the trajectories from one output block. For logarithmic saving this will be a relatively small number of configurations, less than 20. Each configuration starts, as is usual with the xyz format, with the number of particles. The line following is considered a comment line in the xyz format. We use it to write the meta-data. This can include information about the simulation such as the simulation box and the integrator (including whatever parameters are needed to restart a simulation), the time step and log-lin parameters and a list of symbols identifying the columns. It can sometimes be necessary to look at the configuration file (for example in order to check the number of particles or the box size). Here is an example:

```
1000
ioformat=2 timeStepIndex=0 logLin=0,10,0,4,10 numTypes=2 integrator=IntegratorNVE,0.005 \
sim_box=RectangularSimulationBox,9.41,9.41,9.41 mass=1.0000,1.0000 \
columns=type,x,y,z,imx,imy,imz,pe,vir
0 4.6021 0.1670 3.7406 0 0 0 -7.4641 13.2556
0 -0.2562 -4.1929 4.5378 0 0 0 -6.9324 7.1333
0 -3.2822 2.1659 1.1286 0 0 0 -7.2647 15.1471
```

The columns labelled `imx` etc., give the image indices—what multiple of the box size must be added in a given direction to give the total displacement taking into account the number of times a periodic boundary has been crossed. It is essential for correct determination of mean squared displacement, for example. Table 4 shows the quantities that can be controlled using `SetOutputMetaData` on the `trajectory` output manager. So including `rm` for example in a call to `rm` causes the the individual particles' forces to be written to the trajectory configuration files.

6.4 External calculators: adding data to the energies file

In addition to including or excluding the standard energy-like quantities, it is possible to attach an "external calculator" which knows how to calculate some quantity, and have that quantity be included in the energies file. This can be more convenient than creating separate files for new quantities.

Table 3: Identifier and column-label for the main quantities that can be written to the energies file. For extensive quantities, such as potential energy or virial (though not volume), the "per-particle" value is written.

identifier	column label
potentialEnergy	pe
kineticEnergy	ke
virial	W
totalEnergy	Etot
temperature	T
pressure	p
volume	V
density [a]	rho
thermostat_Ps [b]	Ps
enthalpy	H
stress_xx[c]	sxx
stress_yy	syy
stress_zz	szz
stress_xy	sxy
stress_yz	syz
stress_xz	sxz
v_com_x	v_com_x
v_com_y	v_com_y
v_com_z	v_com_z
potentialVirial [d]	pot_W
constraintVirial [d]	con_W
simulationDisplacementLength [e]	dispLength
instantaneousTimeStepSq [e]	dt^2
euclideanLength [e]	eclLength

[a] From version 2.1 on.

[b] The extra degree of freedom used by the NVT integrator.

[c] Versions before 3.0 did not include the kinetic contribution to the atomic stress by default. Change using (Sample method) SetIncludeKineticStress(bool). Contributions to the atomic stress from bond-constraints, angle- and dihedral-forces are not included.

[d] potentialVirial : contribution to virial from the potential; constraintVirial : contribution to virial from constraints.

[e] Parameters associated with NVU dynamics (IntegratorNVU)

Table 4: Identifier and column-label for the quantities that can be optionally written or not to the trajectory file.

identifier	column label
images	imx, imy, imz
velocities	vx, vy, vz
forces	fx, fy, fx
pot_energies	pe
virials	vir

```
pot18 = Pot_IPL_n_18(...)
alt_pot_calc = rumd.AlternatePotentialCalculator(sample=sim.sample, alt_pot=pot18)
sim.AddExternalCalculator(alt_pot_calc)
```

The potential energy calculated by the IPL $n = 18$ potential will then appear in the energies files in the column designated by the "ID string" of the potential. The latter can be obtained via `pot18.GetID_String()` and changed via `pot18.SetID_String()`. Another external calculator calculates the hypervirial by numerical differentiation. Doing the following

```
hypervirCalc = rumd.HypervirialCalculator(sample=sim.sample, delta_ln_rho=0.0005)
sim.AddExternalCalculator(hypervirCalc)
```

will cause columns labelled `approx_vir` and `approx_h_vir` to appear in the energies file.

These represent the finite difference approximations to the virial and hypervirial, respectively. The virial is included in order to be able to compare it to the true virial as a check. The second argument to the constructor of `HypervirialCalculator` is the change in $\ln(\rho)$ (approximately the same as the fractional change in density) to be used for the finite difference calculation. Note: the above lines must appear after setting the potential!

6.5 Other output managers (C++)

An example of an additional output manager implement in C++ is one for computing inherent states (local minima of the potential energy function in 3N dimensions), using `IHS_OutputManager`. Here each time the `Update` function is called, the sample configuration is copied and an energy minimization calculation is run on the copy using the integrator `IntegratorIHS` (see Section 4.8). The energies are written to files `ihsEnergiesXXXX.dat.gz`, while the minimized configurations are written to files `ihsXXXX.xyz.gz`.

```
ihs_man = rumd.IHS_OutputManager(sample=sim.sample, timeStep=0.005)
ihs_man.SetMaximumNumberIterations(5000)
sim.AddOutputManager("ihs", ihs_man)
sim.SetOutputScheduling("ihs", "linear", interval=100)
sim.SetOutputMetaData("ihs", writeIHS_Config=False) # turn off writing of minimized configurations
sim.sample.AddOutputManager(ihs_man)
```

As with other output managers the `SetOutputScheduling` method of `Simulation` is used to control frequency of output. Similarly `SetOutputMetaData` sets whether the minimized configurations should be written or not. Note also that the energy minimization algorithm has not been used very much and probably could be improved, for example by using conjugate-gradient minimization to speed convergence near the minimum. The use of single precision in RUMD is likely to limit the convergence, however.

6.6 Python output managers

Sometimes it is desirable to do data analysis during run-time. This is for instance the case when doing post-production analysis would involve saving huge amounts of data. A python output manager is useful in this case, since it allows so-called call-back functions to be added to the output manager, converting simulation data to output.

To use a python output manager, a call-back function has to be defined. The call-back function should work on the sample object, and return a string with the data, for instance:

```
def momentum(sample):
    vel = sample.GetVelocities()
    P   = sum(vel)
    return "%4g %4g %4g " % (P[0], P[1], P[2])
```

Then a new python output manager can be made, and the call-back function should be added to it:

```
sim.NewOutputManager("analysis")
sim.SetOutputScheduling("analysis", "linear", interval = 100, blockSize = 10000)
sim.RegisterCallback("analysis", momentum, header = "Px, Py, Pz, ")
```

This writes data files named `analysis0000.dat.gz`, `analysis0001.dat.gz` etc. to the same directory as the other output managers.

Note that `RegisterCallback()` can be called multiple times on the same python output manager but with different call-back functions. The output and header strings of the different call-back functions are concatenated to a single line (do not use any end of line characters in these strings).

The scheduling of the python output manager is set in the same way as for the C++ output managers. They do not enforce the block size to be the same, however (related to this, is that the restart mechanism does not take account of Python output managers). To use a different block size for a python output manager, the optional keyword argument `blockSize` can be used.

There is another (useful) difference between the python and C++ output managers: when the schedule of a python output manager is set to "none", it is necessary to specify an interval. The schedule "none" is for for a python output manager the same as the schedule "linear", with the difference that no output is written to files. So the call-back functions that have been added to the python output manager are still called on the specified time steps. This is for instance useful for scaling the box or changing potential parameters step by step, or for doing time averaging of simulation data.

The call-back function can get simulation data from the sample object. Useful functions on the sample object include:

`rm` Returns the atom positions as a numpy array.

`rm` Returns the atom images as a numpy array.

`rm` Returns the atom velocities as a numpy array.

`rm` Returns the atom forces forces as a numpy array.

`rm` Returns the total potential energy.

`rm` Returns the total virial.

A more extensive list of `Sample` methods is given in section 12.

7 Optimizing performance with the autotuner

The basic usage of the autotuner is described in the tutorial. It is convenient to distinguish between "user parameters" and "internal parameters" (also called technical parameters). User parameters specify the kind of simulation the user wants to run: the potential and its parameters, the integrator and its parameters (time step and temperature for example), the number of particles (also of each type) and the density. The type of graphics card is also considered a user parameter, though it is not necessarily something the user can choose explicitly. The autotuner's task is to find the optimal technical parameters for a given set of user parameters. After the call to `Tune` the simulation object is left in the same state as before, but the internal parameters have been set to the optimal values, which are also written to the file `autotune.dat`. If the file exists already when the autotuner is run, it is checked to see if the user parameters match the current ones. If so, the optimal technical parameters are read from the file instead of going through the autotune-process again. (Note that a different GPU type is sufficient to cause the autotuner to run again.)

7.1 Technical parameters determining performance

Here is a description of the technical parameters optimized by the autotuner, mentioning also their default values and how they can be set "manually" by the user.

pb The number of particles per thread-block, for neighbor-list generation, force calculation and integration step, default value 32. A value for `pb` can be specified by the user as a constructor argument to `Simulation`: `sim=Simulation("start.xyz.gz",pb=32)`, or set later (together with `tp`) via `sim.sample.SetPB_TP(pb, tp)`.

tp The number of threads per particle, default value 4, for neighbor-list generation and force calculation (not integration). The total number of threads in a thread block is therefore `pb` times `tp`. It may also be specified as a constructor argument to `Simulation` or via `sim.sample.SetPB_TP(pb, tp)`

NB skin The thickness of the extra shell (beyond the maximum cut-off from the pair potential) searched when constructing the neighbor list. Default value 0.3. If there is more than one pair potential then in principle they could have different skin values, but it is not clear that this would be advantageous; the autotuner enforces a common value. To manually change the skin use `pot.SetNbListSkin(skin)`

Neighbor-list method Which method is used to construct the neighbor-list. The main two possibilities are "n2" meaning an order-(N^2) algorithm is used, which is generally optimal for small system sizes, and "sort" which involves sorting the particles in memory according to their spatial coordinates, allowing for an order-(N) algorithm. Another possibility is "none" which means no neighbor-list is used. This is relevant only for very small systems, where the cutoff is almost as large as it can be without exceeding half the box size. To manually change the NB method use `pot.SetNB_Method("n2")` or `pot.SetNB_Method("sort")`.

sorting scheme Versions of RUMD prior to 3.0 used sorting differently, uncoupled to the building of the neighbor-list. This old-style sorting is still available, because it can help to spatially sort the particles at regular intervals when using the "n2" NB-method (typically for sizes just below when the "sort" NB-method becomes optimal). Different sorting schemes are available: whether to sort in 1 (X), 2 (XY) or 3 (XYZ) dimensions. To set the sorting scheme manually do `sim.sample.SetSortingScheme(rumd.SORT_X)` for X sorting, etc (the options are `rumd.SORT_X`, `rumd.SORT_XY` and `rumd.SORT_XYZ`).

sort-interval The number of time steps between sorting when using the old sorting procedure, default 200. Use `sim.SetSortInterval(...)` to change. Setting the value to be 0 causes old-style sorting to be disabled (this is default since version 3.0).

7.2 Fine control of the autotuner

It can be useful to keep one or more parameters fixed while tuning with respect to all the others. This is done by passing keyword arguments to the constructor of the autotuner, such as

```
at = Autotune(pb=32, NB_method="sort", ...)
```

For this purpose the names of the technical parameters need to match the variable-names used in the Python script itself. They are: `pb`, `tp`, `skin`, `NB_method`, `sorting_scheme`, `sort_interval`.

Another optional keyword argument is `ignore_file=True`, which will cause the autotuner to ignore an existing file `autotune.dat`. In this case it will still (over-)write that file after the tuning process is complete.

8 Post-production analysis tools

The supplied tools are both stand-alone C++ programs that can be called from the command line, and (in many cases) Python modules that can be imported into your script in order to do analysis on the fly (for example if a script involves many simulations at different state points and the data is not otherwise needed afterwards one can save disk space by doing the analysis after each state point.) Use of the most important tools was described in the tutorial, although not all details were given. Here we give a more complete description of their options and output files. There are is not complete harmony of different tools; the "common" options listed below are not available on all tools, for example. There should be some kind of convergence in later versions of rumd. The output files generally start with a comment line which identifies the columns, unless otherwise specified.

8.1 Common command-line options

Help -h gives a list of options for the tool. For example

```
[bead30] ~%rumd_rdf -h
Usage: rumd_rdf [-n <number of bins>] [-m <minimum time between configurations>]
[-f <first_block>] [-l <last_block>] [-v <verbose>] [-d directory]
[-e <writeEachConfig>] [-p <particles per molecule>]
```

Restricting to part of the simulation run Including `-f 10` means data is read from block 10 onwards. Including `-l 20` means data from blocks after block 20 will not be read.

Non-standard output directory If data is not in the default directory `TrajectoryFiles` (either because it's been renamed or because SetOutputDirectory was used to change it from the start), use the `-d <dir>` command-line option to specify where the analysis program should look for the data

Verbosity Including `-v 1` allows various messages to be written to standard output. `-v 0` (the default) disallows (non-essential) output.

Molecular quantities In some tools calculating structural and dynamical quantities, the number of particles per molecule may (or must) be specified (by `-p`[4]). If this is greater than 1 (the default) then corresponding structural and dynamical quantities for the molecular center of mass are also computed.

---

[4]For some tools `rm` was previously used for this.

9 List of analysis tools

9.1 Analysis of the trajectories

rm Compute time averaged radial distribution function. Usage:

```
rumd_rdf [-n <number of bins>] [-m <minimum time between configurations>]
 [-f <first_block>] [-l <last_block>] [-v <verbose>] [-d directory]
[-e <writeEachConfig>] [-p <particles per molecule>]
```

Output file: `rdf.dat`, whose first column is the pair-distance r, with subsequent columns giving $g(r)$ for all pair combinations of atom types; for example if there are two types, then the columns are r, g_{00}, g_{01}, g_{10}, g_{11} (note $g_{01} = g_{10}$). If the option `-e 1` is specified, then a separate rdf file for each configuration is written. The comment line in `rdf.dat` gives the concentrations of the different atom types. The range of r values is the box size in the x-direction (but note that for a cubic periodic box of side L non-zero values do not appear for r greater than $L\sqrt{3}/2$). A minimum (simulation) time (default 0) may be specified using `-m` when configurations have been saved logarithmically in order to reduce processing time and avoid too closely spaced configurations receiving too much weight in the average.

rm Basic dynamical correlation functions: intermediate scattering function, mean-squared displacement, non-Gaussian parameter. Usage:

```
rumd_msd [-h] [-p <particles per molecule>] [-f <first_block>]
[-l <last_block>] [-d directory] [-e <extra_times>] [-s <subtract_cm>] [-v <verbose>] [-t <allow_type
```

Note that a file `qvalues.dat` must be present, containing the q-value at which the self-intermediate scattering function is to be calculated (one value for each atom type). New in version 2.1 is that the restriction to configurations saved logarithmically has been relaxed. When the argument `-e 1` is present, the functions are evaluated more relative times (within a block), giving greater time resolution (but possibly noisier curves). When the argument `-s 1` is present, the center of mass motion is subtracted off, which is necessary when using Langevin dynamics. Output files: `msd.dat` (mean-squared displacement versus time for different atom types), `Fs.dat` (self-intermediate scattering function versus time for different types), `FsSq.dat`, `alpha2` (non-Gaussian parameter versus time for different types). Note that when Lees-Edwards boundary conditions are present, only transverse displacements are included in the calculation.

rm Calculate the static structure factor by Fourier-transforming the radial distribution function in `rdf.dat`. No attempt to handle the finite box-size, and hence integration range, intelligently is made; artifacts of this (extra oscillations) are present in the output. Usage: `rumd_sq <q_start> <q_final> <density>`. 500 values of q ranging between the two user-supplied end-points are calculated. New in version 2.1 is handling an arbitrary number of types. Output files: `Sq.dat` (first column: q, subsequent columns: $S_{AB}(q)$ for different pairs AB of types), `qmax.dat` (location of maximum of $S_{AA}(q)$ for each type A, can be used as input to `rumd_msd` by renaming to `qvalues.dat`).

rm (New in version 2.1.1) Compute bond length probability densities. Usage:

```
rumd_bonds [-h] [-n <number of bins>] [-m <minimum time between configurations>]
          [-f <first_block>] [-l <last_block>] [-d directory] [-t topologyfile]
          [-e <writeEachConfig>] [-v <verbose>]
```

Output file: `bonds.dat`, where the first column gives the bond lengths (the bins) and the rest of the columns contain the probabilities of the bond lengths for each type of bond specified in the topology file. The distributions have been normalized so that the surface under the distribution is 1. See `rumd_rdf` for explanations of the other command line options.

rm Calculate Rouse modes autocorrelation functions and the orientational autocorrelation of the end-to-end vector of linear molecules. (see M. Doi and S. F. Edwards, The Theory of Polymer Dynamics, Oxford Science Publications, 1986) Usage:

```
rumd_rouse [-h] [-p <particles per molecule>] [-f <first_block>]
[-l <last_block>] [-d directory] [-v <verbose>]
```

Output files: `rouse_X0Xt.dat` (autocorrelation functions of Rouse modes) `rouse_X0X0.dat` (variances of Rouse modes) `rouse_R0Rt.dat` (orientational autocorrelation of end-to-end vector) `rouse_R0R0.dat` (mean squared end-to-end vector and squared radius of gyration). Note that this analysis tool only gives correct/meaningful output for systems with one kind of (linear) molecule. The data in the trajectory files should be ordered such that the particle data are in the same order as the particles in the chain. (The program does not read the topology file.)

rm Load one configuration from each block into the molecular visualization program `vmd`.

rm View a configuration using the ray-tracing software POV-ray.

```
rumd_image [-h] [-f <block index>] [-i <input configuration>]
    [-p <povray header>] [-o <output base filename>]  [-W <width of output image>]
    [-a <angle>] [-y <height>] [-r <radial distance>]
    [-L <position that camera should look at>]
```

Use `rumd_image -h` for more usage information.

9.2   Analysis of the energies

rm Read data into `xmgrace` from the energies files as block data, and display one or more columns. The purpose of the tool is to uncompress the energies files into a temporary ASCII file for `xmgrace` to read. Command-line arguments are passed to `xmgrace`. At least one argument of the form `-bxy 0:2` is required to choose columns to be plotted (against each other, or as here, against line number).

rm Basic statistics (means, variances, covariances, drift) of the data in the energies files. Usage:

```
rumd_stats [-h] [-f<first_block>] [-l<last_block>] [-v <verbose> ]
[-d<directory>] [-b base_filename]
```

Written to standard output: total line count, then mean, variance and standard deviation for each column of data in the energies files. Any columns in the energies files which are not among the "standard" quantities, for example data computed by an external calculator, are flagged with a line "Non-standard column identifier: ...". Output files: `energies_mean.dat` (mean values), `energies_var.dat` (variances), `energies_mean_sq.dat` (mean squared values, useful for combining data from independent runs to get overall mean and variance), `energies_covar.dat` (covariances between all pairs of columns in the energies files), `energies_drift.dat` (linear regression estimate of the linear drift of each quantity with respect to line number; divide by the time interval between lines to get a proper rate). The option `-b base_filename`, available from version 2.1.1, allows you to do basic statistics on a different set of files in the output directory. They should be numbered like energies files, just with energies replaced by a different basename, and each one should start with a comment line containing "columns=a,b,c" to identify the columns.

rm Autocorrelation functions of the data in the energies files. Usage:

```
rumd_autocorrelations [-h] [-n] [-l <log_binning>] [-w <write_fraction>]
                      [-c <column_name> [<column_name> ...]]
```

The autocorrelation functions of all columns are calculated and saved to `autocorrelations.dat.gz` by default. The first column of the output file contains the time. With the option `-c column_name` it is possible to choose for which columns in the energies files to calculate the correlation function. The name of the column as in the energies files should then be given. Only the first 0.01 fraction of the autocorrelation is written to the output file by default to save space. It is possible to change this with the `-w write_fraction` option. Another way to reduce the output file size is by performing logarithmic data binning. This is done with the option `-l log_binning`, where the number of bins per decade should be given. The option `-n` is used to normalize the autocorrelation with the variance, so it is unity at time zero.

rm A crosscorrelation function of two columns in the energies files. Usage:

```
rumd_crosscorrelation [-h] [-n] [-l <log_binning>] [-w <write_fraction>]
                      <column_name> <column_name>
```

The two names of the columns for which the crosscorrelation is calculated should be given. The names of the columns will be used for the output file. The command `rumd_crosscorrelation pe W` for instance will result in an output file with the name `crosscorrelation-pe-W.dat.gz`. The other command-line options are the same as for the autocorrelations.

rm Calculate the response function of one or two columns in the energies files. Usage:

```
rumd_response [-h] [-l <log_binning>] <column_name> [<column_name> ...]
```

One or two names of the columns should be specified. The response function is then calculated as the Fourier transform of the derivative of the cross- or autocorrelation function. Again, the name(s) of the column(s) will be used in the output file, so if one column name is given, for instance `rumd_response pe` than the output file is `response-pe.dat`. By default, the data are smoothed with logarithmic data binning with 20 bins per decade. To turn off data binning, write `-l 0`. The output file contains three columns with the angular frequencies, the real parts, and the imaginary parts of the response function.

9.3 Calling analysis tools via the Python interface

Some of the analysis tools can be accessed within Python through the `rumd.Tools` module. These are `rumd_stats`, `rumd_rdf` and `rumd_msd`. Before the analysis tools can be called, it is necessary to make sure the relevant output files are closed. This is done via `sim.sample.TerminateOutputManagers()`. Here is some example code for `rumd_stats`

```
import rumd.Tools
...
sim.Run(n_run_steps)
sim.sample.TerminateOutputManagers()
# create a rumd_stats object
rs = rumd.Tools.rumd_stats()
rs.ComputeStats()
meanVals = rs.GetMeanVals()
meanSqVals = rs.GetMeanSqVals()
pressure = meanVals["p"] # the keys are the symbols in the file meta-data
```

```
    pe = meanVals["pe"]
    pe_var = meanSqVals["pe"] - pe**2
    # to write the file that the command-line rumd_stats program writes:
    rs.WriteStats()
    # to print the stats summary to standard output:
    rs.PrintStats()
```

and here is some code for `rumd_rdf`:

```
    sim.sample.TerminateOutputManagers()

    rdf_obj = rumd.Tools.rumd_rdf()
    # constructor arguments: number of bins and minimum time
    rdf_obj.ComputeAll(1000, 100.0)
    # include the state point information in the rdf file name
    rdf_obj.WriteRDF("rdf%5.3f_%5.3f.dat" % (rho, T))
    # to get as numpy array, first the r-values
    rVals = rdf_obj.GetRadiusValues()
    # and then the actual g(r) values for type 0-type 1 pairs
    gVals = rdf_obj.GetRDFArray(type1=0,type2=1)
```

and for `rumd_msd`:

```
    sim.sample.TerminateOutputManagers()
    msd_obj = rumd.Tools.rumd_msd()


    msd_obj.SetQValues([7.25, 5.75]) # set qvalues (new in V2.1.1, otherwise
    # reads from the file qvalues.dat)
    msd_obj.SetExtraTimesWithinBlock # include extra time differences within
    # an output-block (new in V 2.1.1; corresponds to -e argument)

    msd_obj.ComputeAll()

    # get a nDataPoints x 2 numpy array containing times and mean squared
    # displacement values for  particles of type 0
    msd0 = msd_obj.GetMSD(0)
    # get the intermediate scattering function for particles of type 0
    isf0 = msd_obj.GetISF(0)
    # get the non-Gaussian parameter alpha2 for particles of type 0
    alpha2_0 = msd_obj.GetAlpha2(0)
    msd_obj.GetChi4(0) # get the variance of the intermediate scattering
    # for particle type 0 (new in version 2.1)

    # write data to a file (all types)
    msd_obj.WriteMSD("msd.dat")
    msd_obj.WriteISF("Fs.dat")
    msd_obj.WriteAlpha2("alpha.dat")
```

The module `Python/analyze_energies.py` (new in version 2.1.2) defines a class `AnalyzeEnergies` for the analysis of the energies. Methods for the calculation of correlation functions and response functions are

implemented. The module also defines some generic functions for calculating correlation and response functions of any 1D array. Some example code:

```
import rumd.analyze_energies as analyze
# Create analyze_energies object, read relevant columns from energy files
nrgs = analyze.AnalyzeEnergies()
nrgs.read_energies(['pe', 'W'])

# calculate the normalized autocorrelation
time, auto = nrgs.correlation_function('pe', normalize=True)
# calculate the first 10% of the crosscorrelation
length = int(1.1 * len(nrgs.energies['pe']))
time, cross = nrgs.correlation_function('pe', 'W', length=length)
# calculate response function
omega, response = nrgs.response_function('W', 'W')

# Save correlation functions to file
import collections
output = collections.OrderedDict()
output['time'] = time
output['pe'] = auto
output['pe-W'] = cross
analyze.write_columns_to_file("mycorrelations.dat.gz", output)
# Do logarithmic data binning on the response function and save to file
output = collections.OrderedDict()
bins = analyze.create_logarithmic_bins(omega[1], omega[-1], bins_per_decade=12)
binned_omega, binned_response = analyze.data_binning(omega, response, bins)
output['omega'] = binned_omega
output['real'] = binned_response.real
output['imag'] = binned_response.imag
analyze.write_columns_to_file("myresponse.dat", output)
```

10 Additional features/tricks

10.1 Changing particle masses

As a new feature in version 3.1 it is possible to change the mass of a particle type directly from Python (rather than hacking the configuration file). This is achieved via the `Sample` methods`SetMass` and `SetMasses`:

```
sim.sample.SetMass(0, 4.5) # set mass for particle type 0 to be 4.5
sim.sample.SetMasses(numpy.array([3.6,1.6,0.8])) # set masses for all types using a numpy array (size must
```

10.2 Momentum resetting

Due to the inevitable numerical round-off error due to finite precision arithmetic (single precision in the case of rumd), it is necessary to periodically reset the center of mass velocity to zero. By default this is done every 100 time steps. This interval can be changed via

```
sim.SetMomentumResetInterval(1000)
```

and turn off entirely by passing zero to the above method.

10.3   User defined run-time actions

New in version 2.1 is a facility for allowing users to include extra steps in the main integration loop. It operates similar to python output managers in that the user supplies a function or bound method and specifies how frequently it should be called. The method should not take any arguments so typically it will be a method bound to an object of a Python-class defined in the user's script, which contains references to the sample or simulation or some other relevant object (such as a potential). For example if the sample should be gradually compressed from a low density (this might be done to prepare a molecular configuration), then we could define a class which has a reference to the sample object, and a method `Compress`, as follows

```
class Compressor:
    def __init__(self, sample, scaleFactor):
        self.sample = sample
        self.scaleFactor = scaleFactor
    def Compress(self):
        self.sample.IsotropicScaleSystem(self.scaleFactor)


num_scale_intervals = 100
density_factor = rho_desired/rho_start
scale_factor = math.exp(-math.log(density_factor)/3/num_scale_intervals)
comp = Compressor(sim.sample, scale_factor)

scale_interval=1000
sim.SetRuntimeAction("compress", comp.Compress, scale_interval)
sim.Run(scale_interval * num_scale_intervals)
```

The run-time actions are called in order they were added, before the force calculation in the main loop. The list always includes sorting the particles in memory and setting the total momentum to zero. Weak references are used to store the list of run-time actions in order to allow them to refer to the sample or simulation object without causing problems with circular references.

10.4   Changing the density

Here is a function that will change the simulation box in order to give a specified density.

```
def SetDensity(sim, new_rho):

    # scale to get right density
    nParticles = sim.GetNumberOfParticles()
    vol = sim.GetVolume()
    currentDensity = nParticles/vol
    scaleFactor = pow(new_rho/currentDensity, -1./3)
    sim.ScaleSystem(scaleFactor)
```

When simulating molecules, it might be more reasonable to keep the distances within the molecules fixed, for instance when simulating rigid bonds. In this case, use the command `sim.ScaleSystem(scaleFactor, CM=True)`.

10.5   Instantaneously scaling the velocities

When changing temperature abruptly in the simulation script, the thermostat will take some time to get the average temperature to the right value. This can be circumvented by scaling the velocities instantaneously:

```
# Scale velocities from temperature Told to T
velocity_factor = ( T/Told )**0.5
sim.sample.ScaleVelocities( velocity_factor )
sim.itg.SetTargetTemperature( T )
```

10.6   Starting a simulation from where another left off

It can be as simple as taking the file, say "end.xyz.gz", that you wrote at the end of the last simulation and using this as the start file for the next. If you want perfect continuity, you might want the integrator meta-data to be used. This will only happen if the file is read after the integrator has been set, so you need to re-read the start file:

```
sim = Simulation("start.xyz.gz")
...
itg = ...
...
sim.SetIntegrator(itg)
sim.sample.ReadConf("start.xyz.gz")
```

11   Simulation methods

To get a presentation of the internal documentation use `help(Simulation)` within Python or `pydoc Simulation` on the command line. The docstrings for `Simulation` methods intended to be called by the user are reproduced here for convenience.

rm Add data computed by an external calculator class to the energies files.

   Example

```
alt_pot_calc = AlternatePotentialCalculator(...)
sim.AddExternalCalculator(alt_pot_calc)
```

rm Add an existing output manager object (typically a C++ output manager), specifying a name which will be used to refer to the object when calling SetOutputScheduling and SetOutputMetaData. The name can be anything, but the convention is that it matches the names of the output files associated with this manager.

rm Add a potential for this system (the total force vector and potential will be the sum over all added potentials). Set `clear=True` to remove previously added potentials.

rm Return the number of particles in the sample object.

rm Return the volume defined by the simulation box.

rm Create a new (Python) output manager, which will write files in the output directory starting with name.

rm Read data describing how particles are connected as molecules and create the MoleculeData object.

rm Register a data analysis function (which takes a Sample as argument) to be called with a specified time-step interval. Callback functions that are used to generate output should return a string without newline characters.

rmRemoveExternalCalculator(self, calc) Remove/disassociate the calculator object `calc` from the energies output manager

rm Remove an item from the list of runtime-actions

rm Sets the total momentum to zero via a Galilean velocity transformation

rm Run `num_iter` time steps, by default initializing/resetting output. Specify `restartBlock` to restart a simulation which has been interrupted. Set `suppressAllOutput` to `True` to avoid any output (for example for equilibration). Set `initializeOutput` to `False` to continue the last run without re-initializing all of the output managers.

rm Scale the simulation box and particle positions by scaleFactor in all directions, if direction is omitted; otherwise only in the direction direction. For molecular systems, the centers of mass can be scaled while keeping intra molecular distances constant.

rm Specify the block-size for output (None means automatic). Should be a power of 2 for logarithmic scheduling.

rm Set the integrator for this system, and call its momentum-resetting function.

rm Set how many time steps should go between resetting of center of mass momentum to zero.

rm Change the directory for output files, default TrajectoryFiles.

rm Access to the output manager to control precision and what gets written.

Examples

```
sim.SetOutputMetaData("trajectory", precision=6, virials=True)
sim.SetOutputMetaData("energies", potentialEnergy=False)
```

rm Set scheduling information for an output manager.

manager must be one of the current managers, which include "energies" and "trajectory" and whatever other managers have been added.

schedule must be one of "none", "linear","logarithmic","loglin"

extra keyword arguments may/must be supplied where relevant, e.g. interval=100 for linear scheduling (required). base=10 for logarithmic scheduling (optional, default base is 1) base=1, maxInterval=16 for loglin (required)

rm Specify that a user-supplied bound method, taking no arguments, should be called during the main loop (before the force calculation) every `interval` time steps. `name` is a string used to refer to this action subsequently.

rm Set an alternative simulationBox.

rm Turn on/off most messages in `Simulation` and sample objects. sim.SetVerbose automatically calls sim.sample.SetVerbose; call to the latter explicitly to only affect sample messages.

rm Write the current configuration to a file.

rm Write a file containing the pair potential and pair force for each associated pair potential.

rm Create the simulation object, and the underlying sample object, after reading in a configuration from the named file. Assign default values to various simulation- and output-related parameters. Parameters pb and tp set the number of particles per block and threads per particle respectively. verbose=False suppresses messages to standard output.

12 Sample methods

Here is a list of sample methods which may occasionally be useful in a user-script. Methods which correspond to ones in `Simulation` (which simply call the `Sample` ones) are not listed.

rm Return the number of particle types present.

rm Return the mass associated with a given type.

rm Return the concentration-weight average mass.

rm Return the number of degrees of freedom.

rm Set the number of degrees of freedom manually. This is normally done automatically, but manual setting may be necessary in systems with a lot of overlapping constraints.

rm Return the positions as a numpy array.

rm Return the velocities as a numpy array.

rm Return the forces as a numpy array.

rm Return the images as a numpy array.

rm Carry out force calculation to make sure the forces and energies are up to date. Include `calc_stresses=True` if stress(es) are required.

rm Return the total potential energy. Assumes `CalcF` has been called since positions were last changed.

rm Return the total virial. Assumes `CalcF` has been called since positions were last changed.

rm Return the system's (atomic) stress as a numpy array of length 6. Assumes `CalcF` has been called since positions were last changed.

rm Determine whether the (atomic) stress includes the kinetic contribution.

rm Scale isotropically by linear factor `Rscal`

rm Scale along the coordinate direction given by `dir` by the amount `Rscal`

rm Scale a molecular system isotropically by linear factor `Rscal`. The molecules' centers of mass will be rescaled, but the size of molecules is not changed.

rm Make a homogeneous simple shear by the given amount of strain.

rm Scale all velocities by a given factor.

rm Return the `SimulationBox` object. For example if you want to replace it with a `LeesEdwardsSimulationBox`, you pass it as a constructor argument to the latter.

rm Get name of the output directory (TrajectoryFiles by default)

rm Enable/disable making of back-up output directories.

rm Specify log-lin parameters for an output manager beyond the standard options of linear (corresponding to base = maxInterval = interval) and logarithmic (corresponding to maxInterval=0, with base=1 by default, although other values can be passed).

rm Change verbosity for sample, independently of the Simulation object.

rm Make sure all output files are closed. Necessary if doing analysis is to be done in the same python script (or session) as the simulation.

rm Read a new configuration into the sample from the given file.

rm Check for CUDA errors once every time step for debugging purposes.

rm Return p_b

rm Return t_p

13   Some internal technical details of possible relevance to users

13.1   Copying

Copying of simulation objects is important for the autotuner, which needs to re-run the same simulation many times and leave the simulation object in its original state.

When copying a sample object, it is assumed the data on the device is the most up to date. Therefore a call to CopyDataFromDevice is made on the object being copied; then the host data is copied over to the new object.