This example is a large scale, production quality loop search with LMOD on beta secretase (1W50). Two long loops (21 and 12 residues, respectively) known to flex significantly in the protein, are being sampled with the rest of the protein structure frozen. The results are displayed in two jpeg images, one showing the distribution of low-energy loops and another, closeup view showing the distribution with reference to the X-ray structure colored purple. The starting structure was the X-ray structure, the goal was simply to generate a number of low-energy alternative loop conformations. Nonetheless, the first three lowest energy conformations are within 1.14, 0.94, and 1.00 Å C-alpha RMSD distance from the X-ray structure. The job results can be found in 1W50_run_lmod.log and the LMOD generated low-energy conformations are saved in individual PDB files as well as in a multi-PDB file.

To run this job, one simply needs to compile 1W50_lmod.nab (which will read 1W50_4nab.pdb and 1W50_4nab.prmtop). It is highly recommended to compile it in parallel using mpinab, and run it it on 4-8 cores. 20 LMOD iterations with 200 total LMOD search steps took a little over 8 hours running on 8 cores of a 3.6 GHZ i7 processor. In general, 100 LMOD iterations are recommended for seriously exploring a protein system, which definitely requires parallel execution with as many cores as there are available. LMOD scales quite well thanks to the efficient force parallelization engine in AmberTools.

The structure of a typical LMOD nab script and associated job logfile is fully explained in the main AmberTools documentation under NAB: Molecular Dynamics and Mechanics/Low-MODe (LMOD) optimization methods, here some specifics are pointed out.

1. The general LMOD workflow for proteins includes i) fully flexible minimization of the entire system, ii) applying LMOD search on the partially frozen system meaning that the low frequency modes are those associated with the flexible (specified as “moving”) loops in the presence of the external field (electrostatic, van der Waals, and solvation terms) exerted by the frozen atoms, and iii) re-minimizing the low-energy loop conformations after unfreezing the whole system. (In this example iii) was not actually performed, but it is straightforward to edit the nab script to include it in the calculation.)
2. It is good practice to use a hardware generated seed for the random number generator and save this seed value in the file names associated with an LMOD job. In the example script the number of seconds passed since zero hour, January first, 1970 is used for this purpose.
3. Noteworthy LMOD parameters: `lo.nmod=15` is the number of lowest frequency modes used, `lo.kmod=5` means that out of the 15, every new LMOD iteration explores 5 randomly selected modes (and follows them in both directions), and `lo.nrotran_dof=0` refers to the presence of frozen (or tethered) atoms. `lo.energy_window=50` is set intentionally very high to generate a variety of different loop conformations and not only focus on the lowest energy ones. `lo.conf_separation_rms=1.0` means that in the final set of LMOD generated loops every single pair of them will be at least 1.0 Å superposition RMSD apart. This RMSD calculation includes all moving atoms (LMOD has no knowledge of atom types, etc.). This parameter is quite useful in controlling the diversity of low-energy loops, or conformations in general. `lo.nof_lmod_steps=0` instructs LMOD to try determining a barrier passing event automatically (see documentation). For example, something like this in the logfile
```
9 / 3 E = -1159.966 ( 0.079)  Rg = 4.250  rmsd= 0.083  p= 0.8737
9 /10 E = -1154.146 ( 0.077)  Rg = 4.249  rmsd= 0.575  p= 0.1255
```
means that starting from a particular conformation, low-frequency mode #9 is explored in both directions. In one direction LMOD determined that after only 3 LMOD steps (so-called zig-zag curvilinear perturbation, see documentation) a barrier was passed whereas the other direction required 10 steps. After minimization, the resulting new loop conformations are, respectively,

displaced by 0.083 and 0.575 Å RMSD from the same starting conformation. Rg is the radius of gyration (of the moving atoms) and p is the Boltzmann probability of a conformation with respect to the starting conformation in that particular LMOD iteration. `lo.mc_option=2` sets the Monte Carlo search to "total quenching" which means that every new LMOD iteration continues the search with the lowest-energy conformation found in the previous iteration, irrespective of probability. This option is usually the best for finding the lowest energy structures. For a broad search, however, `lo.mc_option=1` is recommended, which is the traditional Metropolis criterion applied to the minimized energies.

With any questions or comments please contact me at istvan@kolossvary.hu or ikolossv@bu.edu.