HPC usage in the University of Luxembourg Soft Matter Theory Group

Joshua T. Berryman, Muhammad Anwar, Sven Dorosz, Marko Mravlak, Amirhossein Taghavi, Francesco Turci, Tanja Schilling

Overview

Computational Challenges in Soft Matter
- Free Energy Estimation
- Reaction Pathways

Methods In Use

Methods And Cluster Usage Patterns
- Codes Used
- Compilation
- Launch Scripts

Free Energy Estimation

$$Z \propto \int d\vec{x} d\vec{p}\, e^{-\mathcal{H}(\vec{x},\vec{p})}$$

$$A = -k_B T \ln(Z)$$

- PCA to get the normal modes of the dynamics: equivalent to fitting a multivariate Gaussian to Z.
- Many, many other methods . . .

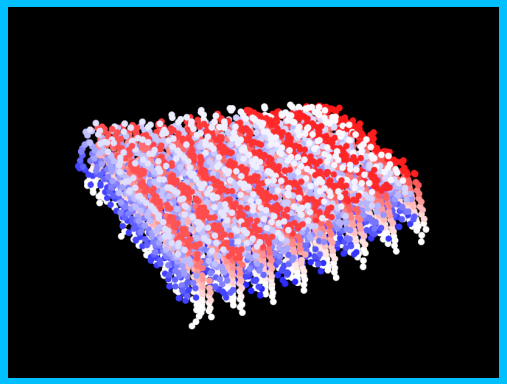

Lara, Reynolds, Berryman, Zhang, Xu, Mezzenga, "ILQINS Hexapeptide, Identified in Lysozyme Left-Handed Helical Ribbons and Nanotubes, Forms Right-Handed Helical Ribbons and Crystals." JACS 2014.

Atomistic DNA in High Salt

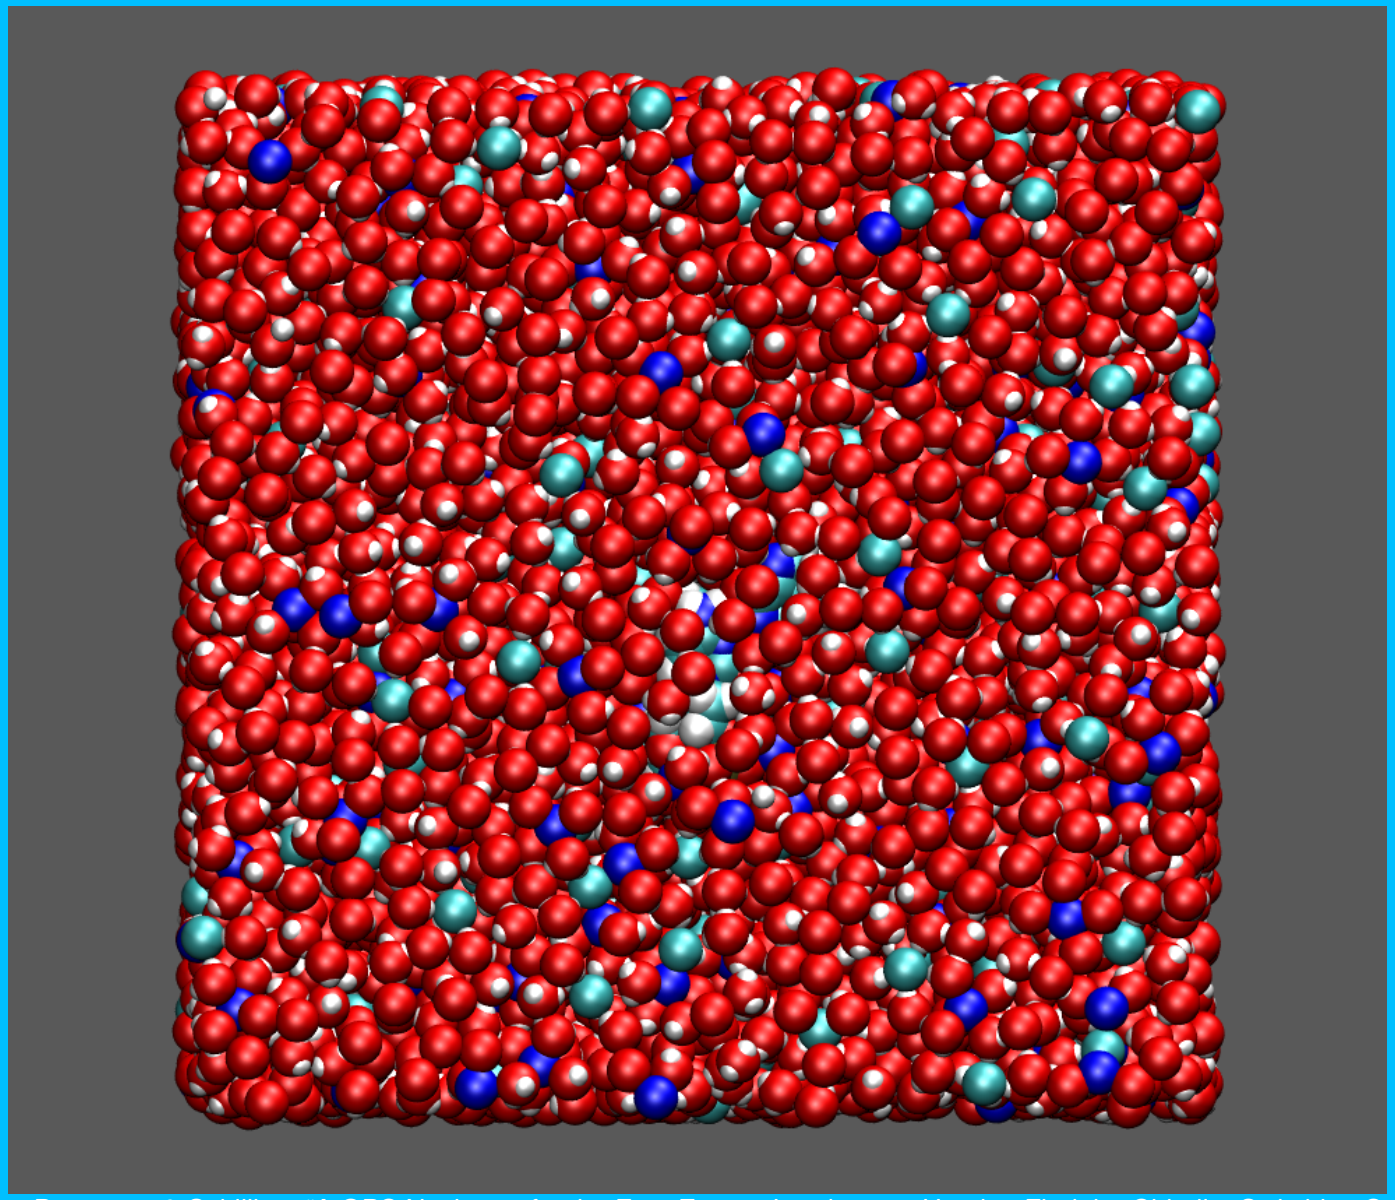

9441 waters

30113 atoms

10^{-19} mol...

10^{-8} sec...

Berryman & Schilling, “A GPS Navigator for the Free Energy Landscape, Used to Find the Chirality-Switching Salt Concentration of DNA” J. Chem. Theory Comput. 2013.

Reaction Pathways

Free energy is only properly defined at thermodynamic equilibrium: to study transitions in collective behaviour, need to take a view of ‘pathways’ instead of ‘landscapes’:

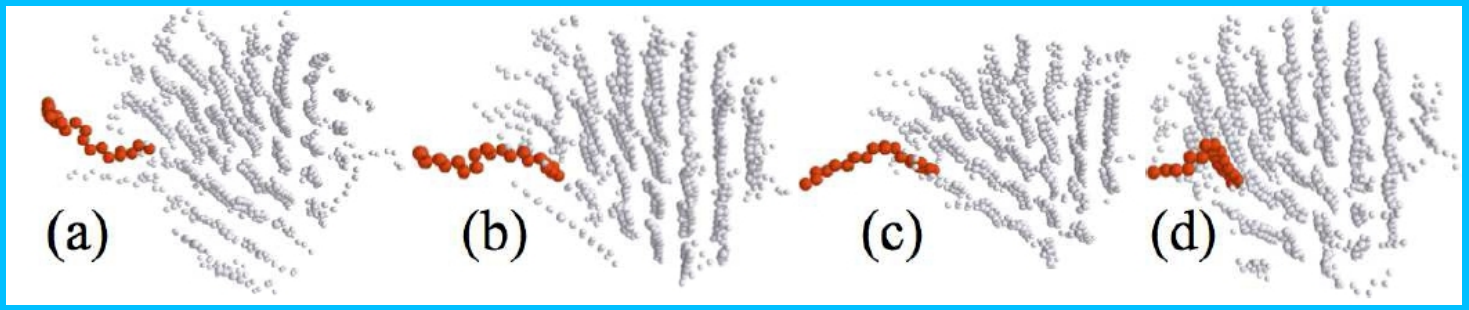

- Brute force MD (e.g. Alkane nucleation pathway above).
- Also rare event methods.

Muhammad Anwar, Francesco Turci and Tanja Schilling, “Crystallization mechanism in melts of short n-alkane chains”

J. Chem. Phys. 2013

Methods In Use

Topic	Method	Parallelism	Username	Papers 2013-2014	
Phase Diagrams	MC	Total	sdorosz	Dorosz et al.	Soft Matter 2014
				Schilling et al.	Eur. Phys. J. Special Topics 2013
	MD	4 cores/run	jberryman	Berryman & Schilling	J. Chem. Theory Comput. 2013
				Case et al.	AMBER 14 2014
				Berryman	Phys. Proc. 2014
		GPU		Lara et al.	JACS 2014
Reaction Paths	MD	12 cores/run	manwar	Anwar et al.	J. Chem. Phys. 2013
		Total	sdorosz	Dorosz & Schilling	J. Chem. Phys. 2013
				Dorosz & Schilling	J. Crystall. Proc. and Tech. 2014
			mradu	Radu et al.	Europhys. Lett. 2014
		Asynchronous	jberryman	Kratzer et al.	Comput. Phys. Commun. 2014

Username	CPU time 2013
sdorosz	195 years 307 days
manwar	128 years 105 days
jberryman	103 years 262 days

Codes Used

Codes:

- Anwar uses ESPResSoMD, own build. **icc+impi**. Standard 12-core one node job script.
- Sven uses his own codes. **icc**. Farms groups of serial jobs.
- I use AMBER, own build. **icc+impi+(CUDA sometimes)**. 4-36 cores. Job scripts to follow.
- . . . group is exploring asynchronous parallelism using FRESHS to act as a wrapper for all of the above.

Compilation (intel):

As we understand it, best practice *currently* for any code is to use intel compilers and MPI:

```
$ oarsub -I
$ module load imkl
$ # module load CUDA
$ export I_MPI_CXX=icc
$ export I_MPI_FC=ifort
$ export I_MPI_CC=icc
$ export I_MPI_F90=ifort
$ make
```

MPI Launch Scripts

Current best-practice MPI job script in our group isn't very pretty:

```
#bin/bash -l
oarsub -l "nodes=3/core=12,walltime=12" \
   ". /etc/profile.d/modules.sh; \
   export MODULEPATH=/opt/apps/HPCBIOS/modules/all; \
   module load imkl; \
   mpirun -hostfile $OAR_NODEFILE $my_exe_name"
```

FRESHS

- GPL python application for rare event sampling.
- Intended as a very open collaboration, currently Kratzer, Berryman, Taudt, Zeman & Arnold.
- http://www.freshs.org

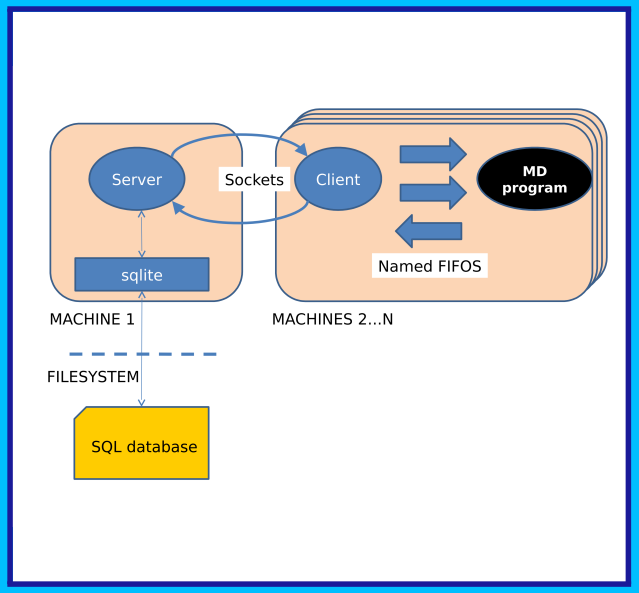

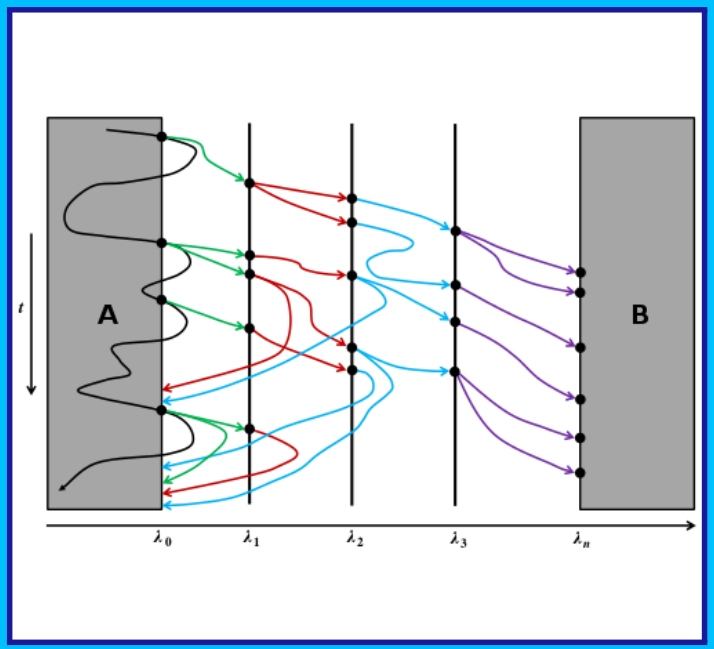

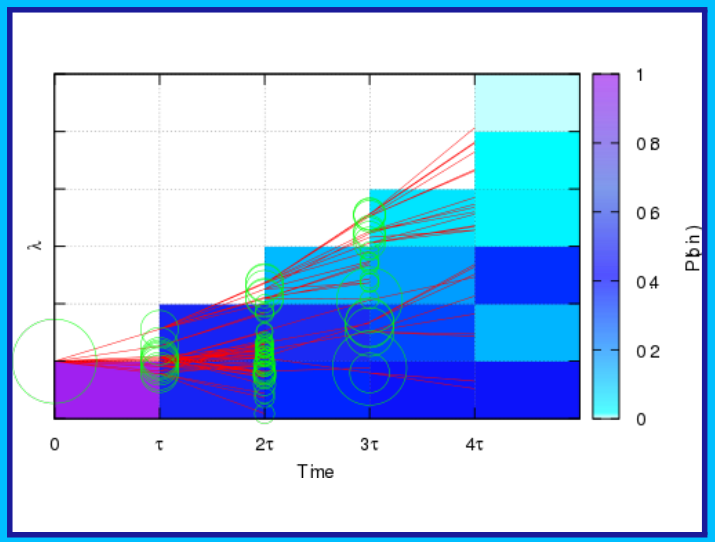

FRESHS Launch Script

Current best-practice FRESHS job script:

```
#bin/bash

##clip the first and last host ids from the list:
NODES=$(cat $OAR_NODEFILE)
SERVER_HOST=$(echo $NODES | awk 'print $1')
LAST_NODE=$(echo $NODES | awk 'print $NF')
NODES=$(echo $NODES | awk 'for(i=2;i<NF;i++)printf "%s ",$i')

##launch the server
oarsh $SERVER_HOST \
   "python $FRESHS/server/main_server.py \
   –db-folder $DB_STAGE –config-folder $CONF –config $inFile \
   >/dev/null 2>server.log"&
```

. . . continued

. . . continued from previous slide

```
##launch the clients

sleep 2
count=0
for node_host in $NODES
do
   oarsh “$node_host” \
      “nice python $FRESHS/client/main_client.py –server $SERVER_HOST \
         >client$count.log 2>&1” &
   count=$[count + 1]
done

oarsh “$LAST_NODE” \
   “nice python $FRESHS/client/main_client.py –server $SERVER_HOST \
      >client$count.log 2>&1”
```

FRESHS Load Types

FRESHS Hard-Sphere Nucleation calculation by Sven:

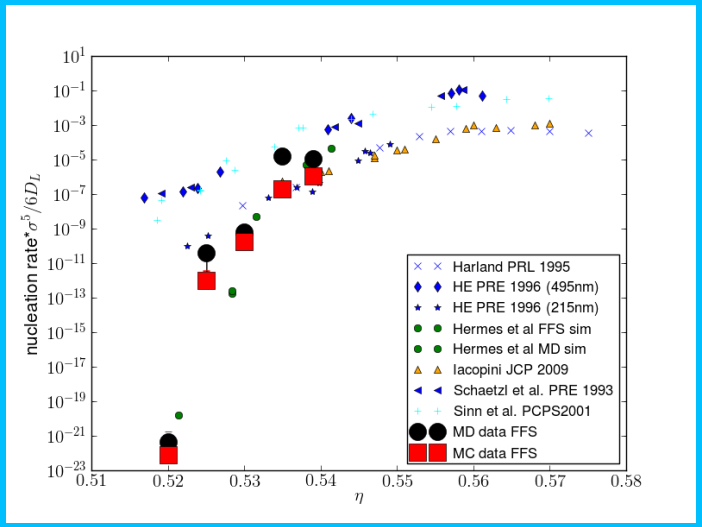

- Usage on ganglia: ≈10% on 1 node. Code spends most of its time in comms or blocking waits.
- Time to run: ≈1 day.
- Time to run by brute force: never.

FRESHS Load Types

FRESHS calculations can be compute-bound, comms-bound or i/o-bound (SPRES).

- compute-bound: haven't yet observed this.
- comms-bound: FFS, typically.
- i/o bound: SPRES, typically.

The best strategy for i/o bound calculations so far has been to save to node-local SSD drives, then compress-and-move to project directories as a background process.

The whole thing has been complicated by NFS and Lustre bugs, as we were apparently the first people to stress the filesystems.

CUDA Performance and Constraints

- AMBER
  - Basic features only are available so far.
  - $\approx$10$\times$ speedup for (1 core + 1 GPU) vs. (12 cores).
  - Memory limitations: $\approx$30k atoms. Cards have approx 6GB (vs. 24GB for nodes) so this is odd.
- ESPResSoMD:
  - advanced features only are available so far. . .