Report on the Best Practices in Molecular Simulation Workshop

Hosted by at NIST, Gaithersburg, MD, August 24th-25th, 2017

Organizers: Michael Shirts, David Mobley, Daniel Siderius

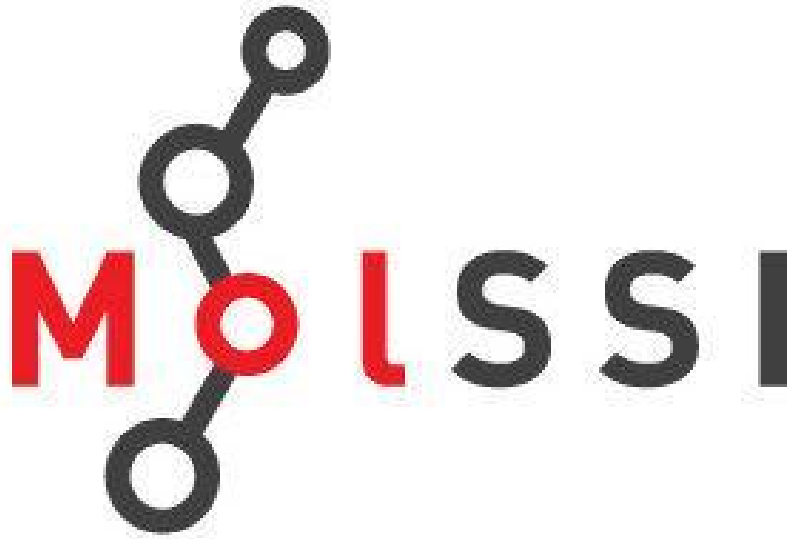

Purpose of Workshop

The motivating theme for the workshop is that improving the way that molecular software is used is as important as improving the software itself, and that improving the use of software will also drive further improvements in our science and software.

Molecular simulation has not achieved its full promise in applications to drug and other molecular design, physical property prediction, and materials design in academia or in the chemically-related industries for a number of reasons. Computational cost is historically the most cited reason for lack of wider adaption of molecular simulation usage in science and engineering. Consequently, a significant amount of effort has been expended making molecular simulation code very efficient both on serial and parallel computers. Significant effort has also been expended to develop methods that can accelerate the computation of thermodynamic quantities and rare molecular events, and in extending simulation packages to incorporate a wide variety of algorithms and features for diverse applications. These efforts to improve the power of molecular software have been extremely successful, making it possible to model far larger systems at far longer timescales than ever before.

However, this success means that other issues have become bottlenecks in performing useful molecular simulations. The wealth of features means that usability can become a problem, and naive users may select poor combinations of options which are incorrect or inconsistent with their goals, but still allowed by the software package. Thus, it takes significant expertise for researchers to generate usable molecular simulation results, and there is lack of reproducibility across the available simulation codes. It is therefore often difficult to validate new computational results against prior simulations or make realistic comparisons to experiment. This workshop focused on facilitating several possible solutions to these problems.

Summary of the Workshop Activities

Thursday Morning Session:

The plenary session began with a statement of the challenge the workshop tried to address: how could the participants of the workshop change the knowledgescape so that molecular simulation is done with maximum impact, according to available literature best practices? The quality of computational science can be improved dramatically if we, as a community, do a better job identifying best practices, documenting these, and educating about/disseminating these. This dissemination will allow the field to better and more effectively employ the immense computing power currently available. Doing this will involve not just making software better, but *using* software better. The organizers specifically intended the workshop to discuss and develop ways for *sustainable* mechanisms to produce high-quality materials that identify, document, and educate users about best practices in order to improve the use

of molecular simulation.

Ed Maginn (Notre Dame) started with an introduction which helped motivate this effort. He highlighted how current problems in the field include a lack of understanding of the algorithms and scientific foundations of the codes available, in addition to known problems in reliability and reproducibility. The discussion focused on how our current efforts can overcome these challenges, including a push for open source code and clear documentation, so any errors can be recognized and corrected.

The plenary session continued with a presentation by John Chodera (Memorial Sloan-Kettering Cancer Center) on molecular simulation usage in drug discovery. He highlighted the long list of user choices at every stage, many of which do not have obvious best choices; worse, we do not always know which parameters will affect results. For example, the choice of protein or ligand protonation states can be of critical importance, but users often simply choose default values with little awareness of the importance of these issues and no way to assess how their choices will affect their results. Dr. Chodera offered several suggestions to address these issues, including collecting “usual suspects” that can cause simulations to go astray, and a “simulation health report” that collects results of a standard battery of analysis tests in much the same way as a set of lab tests a doctor orders would report a variety of indicators of a patient’s health. He pointed out other efforts in this area, including HTMD.org and OpenBioSim, and we perhaps should work to coordinate such efforts. Discussion emphasized the desire to share information about the most common pitfalls, especially where there is and where there is not clear information indicating that one choice should be preferred. This information is useful to referees as well, and guidelines from journals such as Journal of Chemical & Engineering Data can inform checklists that we make as part of this effort.

The Living Journal of Computational Molecular Science (LiveCoMS) was proposed as an open access, online journal that provides a new take on our current publication model in order to improve the quality of modeling work. The journal aims to be distinctive in two key ways: (a) the “living” aspect - authors will be able to continuously update their publications, responding to comments, new information, and progress in the field; (b) the incentivization of article types that foster the highest standards in the community but otherwise might be difficult to publish for academic credit.

Significant planning has gone into the journal and mechanisms to keep articles “living.” GitHub provides a mechanism to do this. Copyright will remain with the authors. The journal itself will be hosted/published through Scholastica, a well-established online journal host. Types of articles will include: (1) Best practice guides, (2) “Updatable” reviews, (3) Tutorials, (4) Program comparisons and benchmarks, and (5) Lessons Learned, all of which benefit greatly from this “living” model of ongoing updates. For example, typical review articles are obsolete 3 months before they are published, in the current publication model; here, these can be updated

on an ongoing basis. The journal will begin with classical statistical mechanics and molecular simulations but is is expected to broaden to include quantum mechanics in the near future. Plans were developed for a Best Practices Tutorial and Guide on how to prepare Best Practices tutorials and guides. Copyright and legal considerations were briefly discussed.

Thursday Afternoon Session:

The first session of the afternoon was a discussion on how to improve the ability to perform comparisons between simulation programs, ensuring that simulation packages yield comparable results when they ought. That is, packages which can compute the same properties given the same data should yield consistent estimates of the computed properties. Comparisons to ensure this is the case should be strongly encouraged and should in principle be straightforward. However direct comparison of simulation toolkits and results is currently difficult to perform automatically, and the next session focused on a discussion of these problems. Current hindrances to comparison include technical (interoperability of file formats used by software, requirement of software-specific knowledge to be sure that the simulations were performed correctly) as well as social (coauthorship for grant reviews, "bad taste" of potentially unfavorable comparisons) aspects. We determined that an open access journal could provide relief in these problems by publishing studies that may not be easily published elsewhere (updates for benchmark-type studies, confirming consistency between new and old versions of the same software or different execution configurations).

One way MolSSI could potentially help this is effort is by gathering a consensus on and hosting a library of systems to use for comparison studies and as reference benchmarks. This reference library would likely need to span multiple application domains (materials science, chemical engineering, biophysics) to be useful for the entire simulation community (or communities). Another way MolSSI could make this possible is organizing occasional challenge-type "contests" for particular properties, systems, or computational methods, where MolSSI provides some standard systems, at some level of detail and asks people to run them.

The final portion of the afternoon was devoted to two one-hour sessions of working group meetings to plan the initial best practices documents and, time allowing, begin development of those documents. Five groups met in each hour block. Each work ing group identified a leader for discussion, and reviewed outlines of the planned best practices documents that had been developed collaboratively by the group through online meetings in the three months prior to this workshop. They then collaboratively expanded and developed those outlines into feasible starting points for writing the actual documents. Currently, the working group notes are contained as Google Docs, but will be transferred to Github repositories for further authoring using LaTeX.

The output of each working group is intended to be a full best practices guide that can stand independently, but which can be submitted to LiveCoMS when it begins accepting articles in the next few months. Groups agreed to set deadlines between September 30th and December 15th for the completion of the first version suitable for publication, with the expectation that most of the group would continue to revise with community feedback afterwards.

Friday Morning Sessions

The morning session began with presentations from one representative from each of working group. The topics of the group presentations (and group leaders who have agreed to manage the groups were) were.

1. Basic simulation training, background knowledge : Avisek Das
2. Potential of Mean Force/Free Energy Profiles: Heather Mayes
3. Polymer Simulation Setup: Michael Fortunato
4. Transport Properties: Richard Elliot
5. Software validation/Physical validation: Pascal Merz
6. Glass Transition Temperature: Paul Saxe
7. Sampling-uncertainty assessment: Dan Zuckerman
8. Simulation Setup - Interfacial: Jacob Monroe
9. Biomolecular Setup: Toni Mey
10. Alchemical Free Energy: John Chodera

Representatives of the group summarized their progress over the summer and the previous afternoon. A revised scope for each best practices document was laid out and reviewed collectively by all conference attendees, and comments and feedback on important points were provided.

Friday Afternoon Session:

The last session consisted in a group discussion of how to make the projects discussed in the workshop possible. A number of ideas for editorial policies and mechanisms for LiveCoMS were discussed, including adding a “Lessons Learned” category of articles, identifying ways to make sure molecular simulation comparisons were fair, how to promote civility in discussions about document revisions, how to keep articles from becoming “dead” instead of being updated. We also discussed what would be required or recommended for each of the articles. These comments were transferred into the authors instructions section of the LiveCoMS GitHub repository, where the are recorded as issues that will be incorporated into the LiveCoMS official policies.

We also discussed other ways that the ideas brought up during the workshop could be carried forward. We discussed how MolSSI could serve as a neutral arbiter of simulation comparisons of the sort published in the journal, as well as maintain a database of validation test input files.

We also discussed how these best practices that we discussed. It was suggested there should be a summer school on best practices use of software, like the

summer school on simulation development currently hosted by MolSSI. We talked about sponsoring best practices sessions at ACS (there is already one at AIChE that occurs from time to time). We discussed how there should be collaboration among the many people across the country (and world) offering simulation courses to collaborate to develop an online course or set of courses. David Kofke, David Mobley, and others are beginning to pull together and adapt existing materials along these lines, and also working to make them easier to update via migration to GitHub.

Summary of Outcomes

A first outcome of this workshop is a plan for the the creation of LiveCoMS, the *Living Journal of Computational Molecular Science*, (http://www.livecomsjournal.org) a peer-reviewed overlay journal allowing versioned updates. Authors post documents to their preprint server of choice, gaining the benefits of traditional publishing. By doing this, LiveCoMS provides (1) timely and high quality documents to advance our science which are (2) updatable, with (3) incentives provided so authors receive appropriate credit for their contributions. The editorial board is partially constituted, and will continue to be constituted, especially with more international authors. The journal will publish a range of documents that can and should be updated, including 1) best practice guides, 2) perpetually updated reviews, 3) comparisons between molecular simulation codes, 4) lessons learned, and 5) tutorials. The approach will use a "paper writing as code development" model, with authors maintaining a GitHub repository for their paper so that community members can provide feedback by filing GitHub issues with suggestions, comments, and concerns.

Another important outcome of the workshop is a set of working drafts of best practice guides that are being developed as a resource for the community. The topics are listed in the section above. In some cases, the topics will be restricted in scope to provide sufficiently deep coverage of the topics, so they can truly be best practices. The groups are committed to providing the documents by December 15th, 2017, and to submit them to LiveCoMS assuming the journal is running at that point.

Another conclusion of the workshop is that there are number of activities for which LiveCoMS is not useful mechanism to improve best practices, but that MolSSI is well-positioned to help lead and facilitate the collaborations of a number of researchers (both participants at the workshop and beyond) to developing best practices for using molecular simulation.

One idea discussed was that MolSSI could host a meeting for people who have developed molecular simulation courses across the country (or in other countries) and help them get together and together develop a consistent online course or set of courses that would be useful worldwide. Another activity that MolSSI could head would be to provide resources (coding infrastructure, computing resources) to run validation tests between codes, serving as a neutral arbiter between codes. Finally MolSSI could lead the development of a repository of standard benchmark systems for the community.

Use cases developed

Because the focus of the workshop is the *use* of software rather than the writing of software, the main products are instructions for how to perform certain molecular simulation tasks. Specifically, the participants will produce the documents, suitable for submission to LiveCoMS, which were started in the working groups listed above. These will be posted on a preprint server at submission time, so will be available to start being edited at that point. LiveCoMS is the planned mechanism by which these documents will be publicized and improved iteratively and with community input, though they will stand on their own. The journal model introduces a mechanism for an ongoing process that produces many more high-quality documents.

Schedule

Thursday, August 24th:

9-10:30 The challenge:

- 9:00-9:15 Introduction (Michael Shirts)
- 9:15-9:40 Keynote: “Making molecular simulation a reliable and reproducible tool” (Ed Maginn)
- 9:40-9:50 Discussion
- 9:50-10:15 Challenges in reproducibility and reliability of simulation setup (John Chodera)
- 10:15-10:30 Discussion

10:30-11:00 Break

11-12:30 The means to solve the problem:

- 11:00-11:20: LiveCoMS: A best practices overlay journal, i.e. How we are planning to get people credit for doing best practices documents (Michael Shirts, David Mobley, and Dan Zuckerman)
- 11:20-11:50: Discussion and brainstorming for the overlay journal
- 11:50-12:30: Discussion: How to best educate people beyond a journal: Tutorials, wikis, workshops, simulation requirements in other journals

12:30-1:45 Lunch

1:45-2:30 Discussion: How do we run/incentivize comparisons between codes?

2:30-3:30 Small group breakout session 1: critique half of the working documents in parallel

3:30-4:00 Break

4:00-5:00 Small group breakout session 2: critique other half of the working documents in parallel

5:00-5:30 Daily summary and wrap up

6:30-8:00 Dinner (Rook's Corner Restaurant, at the Hilton)

Friday, August 25th:

9-10:30 Discussion of best practices: Presentation of work so far:
each group picks their speaker or speakers to present their work.

- (first half of topics) 10 min per presentation, 5 min discussion.

10:30-11:00 Break
11-12:30 Discussion of best practices: Presentation of work so far
Each group picks their speaker or speakers to present the groups' work

- (second half of topics) 10 min per presentation, 5 min discussion.

12:30-1:30 Lunch
1:30-2:30 Discussion of plans for moving forward. Topics include:

- Publicity for a overlay journal
- Recruiting more people for further documents
- Organization for a code comparison effort
- Other topics

2:30-3:00 Collaboratively writing outcomes document
3:00 Workshop end

Budget

A total of $15,000 was budgeted for the conference:

- $4100 was spent on airfare and other transportation costs for applicants (approximately 70% of nonlocal and nonfederal attendees were reimbursed for airfare).
- Approximately $2300 was spent on dinner, lunch, and break snacks and coffee for attendees
- NIST provided the facility free of charge
- Approximately $6200 was spent on lodging (2 or 3 days) for attendees. All non-federal attendees had their hotel paid for.

The conference thus ended well within the budget.

Participants

Participants came from a range of career points and subject areas including chemical engineering, chemistry, material science biophysics, and applied math. As the workshop was at NIST, a number of other NIST employees participated in parts of the workshop. In two cases, participants had to cancel the last week

Attendee	Institution	Department
Rommie Amaro	UC San Diego	Chemistry and Biochemistry
Alan Grossfield	U. of Rochester Med. Center	Chemistry and Biophysics
John Chodera	Sloan Kettering Institute	Computational Biology
Coray Colina	University of Florida	Chemistry & Materials Science
Avisek Das	University of Michigan	Chemical Engineering
Andrew Dienstfrey	NIST	Mathematical Analysis & Modeling

Eliseo Marin-Rimoldi	MolSSI	
Richard Elliott	University of Akron	Chemical and Biomolecular Engr.
Kristen Fichthorn	Penn State University	Chemical Engineering
Michael Fortunato	University of Florida	Chemistry
Justin Gilmer	Vanderbilt University	Materials Science and Engr.
Anthony Hazel	Georgia Tech	Physics
Sunny Hwang	Georgia Tech	Materials Science and Engineering
David Kofke	SUNY University at Buffalo	Chemical & Biological Engineering
Ed Maginn	University of Notre Dame	Chemical & Biomolecular Engr.
Heather Mayes	University of Michigan	Chemical Engineering
Pascal Merz	University of Colorado Boulder	Chemical and Biological Engineering
Richard Messerly	NIST	Material Measurement Laboratory
Antonia Mey	University of Edinburgh	Chemistry
David Mobley	UC Irvine	Pharmaceutical Sci. and Chemistry
Jacob Monroe	U. of California, Santa Barbara	Chemical Engineering
Jessica Nash	MolSSI	
Conor Parks	U. of California, San Diego	Chemistry and Biochemistry
Paul Patrone	NIST	Mathematical Analysis & Modeling
Baron Peters	U. of California, Santa Barbara	Chemical Engineering
Julia Rice	IBM Research	
Daniel Roe	NIH	NHLBI
Samarjeet	Johns Hopkins / NIH	NHLBI
Paul Saxe	MolSSI	
Andrew Schultz	SUNY at Buffalo	Chemical & Biological Engineering
Michael Shirts	University of Colorado Boulder	Chemical and Biological Engineering
Daniel Siderius	NIST	Chemical Sciences Division
Matthew Spellings	University of Michigan	Department of Chemical Enginering
William Swope	IBM Research	
Xiongwu Wu	NIH	NHLBI
Daniel Zuckerman	Oregon Health & Science U.	Biomedical Engineering