Computational Statistics and Molecular Simulation: A practical cross-fertilization

Gabriel Stoltz (Ecole des Ponts and Inria Paris),
Luke Bornn (Simon Fraser University),
Christian Robert (Université Paris Dauphine and University of Warwick)

November 11-16, 2018

1 Overview of the Field

The aim of molecular simulation is to predict macroscopic properties of condensed matter from a microscopic description. This is done according to the laws of statistical mechanics which state that, in principle, any macroscopic property is obtained by averaging microscopic configurations according to some prescribed probability distribution, or *thermodynamic ensemble* [2]. In practice, these macroscopic properties are computed by time averages over sufficiently long trajectories generated by dynamical systems that are ergodic with respect to the probability measure at hand (see for instance [26] for a review). Quite often the rate of convergence is poor because the dynamics remain stuck in local minima of the energy landscape of the system—a phenomenon called *metastability*. Although rare, the passage from one local minimum to another is typically fast. This gives rise to a time scale separation in the system which can be exploited by suitable numerical methods to speed up the sampling of the distribution.

Similar issues arise in the field of computational statistics when the probability of rare events needs to be accurately estimated, e.g., in reliability analysis or in the performance analysis of large communication networks, and in the context of Bayesian statistics where high dimensional probability measures need to be sampled. If straightforward Monte Carlo techniques are implemented, estimating small probabilities requires one to simulate a possibly large system over very long computation times, a feature that is inefficient and maybe even unfeasible. The common problem in all cases is that the variance of the estimated quantities can be large. A standard statistical method for tackling these issues is importance sampling, a variance reduction technique which simulates the system under a different set of parameters or boundary conditions so that the rare events are no longer rare under the corresponding probability measures (see for instance [6] for an elementary introduction to variance reduction techniques).

As scientific issues are rather similar in both fields, it is not surprising that many numerical methods are used both for molecular simulation and statistical problems. The most prominent example is the Metropolis algorithm [29, 19] which was invented in the computational physics community before it was spotted and extended by statisticians to sample generally complicated, high-dimensional probability distributions. Yet other, more recent, examples are found in the study of Hybrid Monte Carlo methods for statistical inference [11, 28, 18] (nowadays known in fact as Hamiltonian Monte Carlo in the statistics literature), the use of free energy techniques from molecular simulation in importance sampling techniques in statistical applications [9], the mathematical and numerical analysis of the Wang-Landau technique from molecular simulations using techniques from statistics [16], the application of splitting techniques developed in statistics to sample rare events in molecular simulation [7, 8], etc.

It is fair to acknowledge that the communication between the communities of molecular simulation and computational statistics relies on a few individuals only, which is why it is not uncommon that similar (or even sometimes identical) numerical methods are (re-)discovered independently in both fields at a few years interval. For instance, the famous Metropolis adjusted Langevin algorithm (MALA) introduced by Roberts and Tweedie in the statistics literature in 1998 [34] had been known since 1978 in the molecular simulation community [35] (but not so much used), whereas advances on the estimation of normalizing constants in statistics, which considerably extended previous work from physics as developed by Bennett in 1976 [3], remained unnoticed by the molecular simulation community until the recent work of Chodera and Shirts in 2008 [36].

We felt, therefore, that it would be very beneficial for research and researchers in molecular dynamics and statistics to once again bring these communities in closer contact and to facilitate the exchange of ideas and expertise, at the uttermost practical level.

2 Recent Developments and Open Problems

Prior to the workshop, we asked all participants to contribute by one question/motivation which was driving their attendance to the workshop. These questions allow to give a perspective on the open problems of interest in the communities of computational statistics and molecular dynamics, and are in direct relationship with recent progresses in these fields.

2.1 Sampling in very high dimension

As mentioned in Section 1, probability measures to be sampled in computational statistics and molecular dynamics are in very high dimension. The design of efficient and reliable sampling methods, able to cope as much as possible with the multimodality of the target measure, is still an intense area of research. In the statistics community, a specific care is taken to consider methods which work well for situations when the target measure has heavy tails. Standard gradient-based samplers slow down in the tails and cannot efficiently sample configurations in these regions.

2.2 Non-reversible sampling

Many works recently considered non-reversible samplers, obtained by non-gradient perturbations of Langevin dynamics for instance, or relying on piecewise deterministic Markov processes (PDMPs) [10]. Examples include the Zig-Zag process [13, 4], the Bouncy Particle Sampler [32, 31, 5], Randomized HMC (originally known as Andersen dynamics [1]), etc. The latter topic is of clear interest to researchers both in statistics and molecular dynamics. There is some indication now that some PDMPs perform better than others, at least in certain regimes, but efforts are still needed in order to gain a complete picture of the field.

The variance reduction arising from sampling with non-reversible methods is now documented by several works [20, 21, 24, 33]. This corresponds to situations when the target measure has a known analytical expression. An interesting area in molecular dynamics concerns the computation of properties of nonequilibrium systems, for which the invariant measure is unknown (see [26, Chapter 5]). Variance reduction in this situation remains a challenge.

2.3 Rare event sampling

Sampling typical properties of systems in high dimension is already quite challenging. It is even more difficult to sample atypical events. These events can correspond either to specific molecular configurations or Bayesian parameters; or even to paths connecting two important modes of the target distribution of interest. Apart from the dimensionality of the problem, additional challenges arise from the fact that sample paths should link metastable states which may be unknown in advance in complex problems. Some importance sampling based on large deviation results can be performed, based on the works [14, 15, 17, 12], but ideally one would like to go beyond the regime of large deviations in order not to have to rely on a parameter being sent either to 0 or to infinity (time, temperature, etc).

2.4 Model misspecification and uncertainty quantification

An emerging topic in molecular dynamics is the quantification of the uncertainties arising from modeling issues. Modeling in computational statistical physics is done primarily through the design of a potential energy function which accounts for interactions between the elementary constituents of matter. Most researchers currently take the model as granted in the molecular dynamics community. Recent efforts are however invested in understanding how changes in the model affect the simulation results. Questioning the model is on the other hand a traditional task in computational statistics. Modeling is an important and natural task in computational statistics, which arises in the Bayesian setting through the choice of the likelihood function and prior distribution on the parameters.

2.5 Coarse-graining and constraints

A last topic of interest concerns reduction techniques. One situation is the sampling of measures concentrated near lower-dimensional manifold. Some constraints may even limit strictly the range of the simulation results, the target measure to be sampled being reduced in this case to a measure supported by a lower-dimensional space. This arises in molecular dynamics when molecular constraints are taken into account (e.g. fixed bond lengths or bond angles), and in statistics in Approximate Bayesian Computations (ABC) [27]. Recent works showed how to unbiasedly sample from probability measures on submanifolds when the log-likelihood of the unconstrained measure can be computed [38, 25]. This is however not the case for ABC, where one can only simulate from the model (i.e. draw samples according to the target measure), without accessing the expression of the target measure. The constraints arising in computational statistics can often be expressed through some summary statistics, which gathers some important information on the data. The equivalent of this function in the computational statistical physics world is known as a reaction coordinate. The name comes from the fact that this nonlinear function of the configuration of the system, which has values in a low dimensional space, indexes the transition of the system from one metastable state to another, as in chemical reactions.

Statistical coarse-graining techniques should also be of great interest to construct surrogate models in molecular dynamics. Such techniques include Generative Adversarial Networks, clustering, diffusion maps, auto-encoders, etc, with the aim of constructing better reaction coordinates for instance.

3 Presentation Highlights

There were two types of lectures: standard 30-minutes long talks, and longer "hands-on" sessions. We wanted to emphasize practical cross-fertilization that is not only based on attending lectures from researchers from other communities, but also operated through an active participation in these "live" hands-on sessions, where a researcher from one domain presented a numerical test case of interest to everyone, possibly ran simulations and commented on the results, and in any case underlined the actual challenges at hand in this problem. This presentation was followed, and in fact constantly interrupted, by all sorts of questions, comments, remarks from researchers of both statistics and molecular dynamics. We thus deem this innovation a success and plan to repeat it in future workshops involving researchers from several communities.

Let us now give a more chronological (and somewhat biased!) overview of the week, somewhat following the live posts Christian Robert wrote on his famous blog `https://xianblog.wordpress.com/` during the workshop.

Monday. On the first day of the workshop, Arthur Voter (Los Alamos National Laboratory) gave a fantastic hands-on review of molecular dynamics for material sciences, which was primarily aimed at the statistician side of the audience and most helpful in their understanding of the concepts and techniques at the source of HMC and PDMP algorithms. This presentation included the BAOAB version of HMC, which sounded to neophytes like an improvement to investigate. A part on metastability was completed by a talk by Florian Maire, at a more advanced level.

The shorter and more traditional talks of Monday all brought new perspectives and information to the audience, although they were definitely more oriented towards their "own" side of the audience than the hands-on lecture. For instance, Jesús María Sanz-Serna (Universidad Carlos III de Madrid) gave a wide

ranging and accessible overview of numerical integrators in the context of HMC and Tony Lelièvre (École des Ponts) presented a recent work on simulating measures supported by manifolds via an HMC technique constantly projecting over the manifold, with proper validation. There was a talk by Josh Fass (Memorial Sloan Kettering Cancer Center) on simulating implicit solvent models that mixed high-level programming and reversible jump MCMC, with an earlier talk by Yong Chen (University of Michigan) on variable dimension hidden Markov models that could have also alluded to reversible jump. Angela Bitto (WU Wien) talked about using ASIS (Ancilarity-sufficiency interweaving strategy) for improving the dynamics of an MCMC sampler associated with a spike & slab prior, the recentering-decentering cycle being as always a magical step (as to why it works better despite introducing multimodality in this case), and Gael Martin (University of Monash) presented some new results on her on-going work with David Frazier about approximate Bayes computation (ABC) with misspecified models, the summary statistic being there a score function that relates the work to the likelihood free approach of Bissiri et al.

Tuesday. On the second day, Carsten Hartmann (BTU Cottbus-Senftenberg) exploited a representation of the log cumulant as solution to a minimisation problem over a collection of importance functions (by the Donsker–Varadhan principle), with links to cross-entropy and optimal control; a theme also considered by Alain Durmus (ENS Cachan) when considering the uncorrected discretised Langevin diffusion with a decreasing sequence of discretisation scale factors (Jordan, Kinderlehrer and Otto) in the spirit of convex regularisation à la Rockafellar. His presentation also included representing ULA as an inexact gradient descent algorithm. Murray Pollock (University of Warwick) presented a new and exciting technique called fusion that simulates from products of d densities, as in scalable MCMC (but not only so). The method stems from an (early) starting and startling remark that when simulating one realisation from each density in the product and waiting for all of them to be equal means simulating from the product, in a strong link to the (A)BC fundamentals. This is of course impractical and Murray proposes to follow d Brownian bridges all ending up in the average of these simulations, constructing an acceptance probability that is computable and validating the output.

The second "hands-on" lecture was delivered by Gareth Roberts (University of Warwick) on the many aspects of scaling MCMC algorithms, which started with the famous 0.234 acceptance rate paper of 1996 by Roberts, Gelman and Gilks. While some of us were aware of some of these results, the overall picture was impressive, including a notion of complexity some of us had not seen before. The lecture included a final section on PDMPs where Gareth presented very recent on the different scales of convergence of Zigzag and bouncy particle samplers, mostly to the advantage of Zigzag.

In the afternoon, Jeremy Heng (Harvard University) presented a continuous time version of simulated tempering by adding a drift to the Langevin diffusion with time-varying energy, which must be solution to the Liouville partial differential equation (PDE) $\mathrm{div}\pi_t f = \partial_t \pi_t$. This approach connects to a flow transport problem when solving the PDE under additional conditions. This talk was very much at the interface and hence completely in the spirit of the workshop. Jonathan Weare's (New York University) talk was about quantum chemistry which translated into finding eigenvalues of an operator using stochastic techniques, turning in to a change of basis in an inhumanly large space (10^{180} dimensions!). Matt Moore (University of Wollongong) presented the work on Raman spectroscopy he did while a postdoc at Warwick, with an SMC based classification of the peaks of a spectrum (to be used on Mars?) and Alessandra Iacobucci (Université Paris Dauphine) showed us the unexpected thermal features exhibited by simulations of chains of rotors subjected to both thermal and mechanical forcings.

Wednesday. The first talk of the Wednesday morning by Jianfeng Lu (Duke) was less accessible to the statisticians, partly due to less common notations, although the topic very much correlated to their interests including path sampling, with an augmented version of HMC using an auxiliary indicator. It also brought back BAOAB in the picture. Next, Marcello Pereyra (Heriot Watt University) spoke about Bayesian image analysis, with the difficulty of setting a prior on an image. In case of astronomical images there are motivations for an L^1 penalisation sparse prior. Since sampling is an issue, Moreau-Yoshida proximal optimisation is used instead, in connection with his MCMC survey published in Statistics & Computing two years ago. Transferability was a new concept for most of us, as introduced by Kerrie Mengersen (QUT Brisbane), to extrapolate an estimated model to another system without using the posterior as a prior. The talk included a great interlude about the crown of thorns starfish killer robot! The concept appears like a prior determination

based on historical data, in connection with recent (2018) Technometrics and Bayesian Analysis papers towards rejecting non-plausible priors. The use of summary statistics for prior calibration gave the approach an ABC flavour.

The hands-on session of the day was Jonathan Mattingly's discussion of gerrymandering reflecting on his experience at court on this very public topic. As such, it was hard to beat for an engaging talk reaching between communities and beyond. It was quite exciting to listen to Jonathan explaining his vision of the problem and on the resolution by the simulation of random electoral maps through MCMC, as well as debating the notion of fairness and of extreme electoral maps.

Thursday. The Thursday of the workshop saw a wee bit of reduction in the audience as some participants had already left Oaxaca. Meaning they missed the talk of Christophe Andrieu (Bristol) on hypocoercivity which could have been another hands-on lecture, given its highly pedagogical contents. The audience showed its appreciation of the talk by engaging into a lively discussion, as for instance on the exhibited connection between the guided random walks of Gustafson and Diaconis, and continuous time processes like PDMP. Then Samuel Power (University of Cambridge) gave another perspective on PDMPs, with another augmentation, connected with time, that he calls trajectorial reversibility. This has the impact of diminishing the event rate, but creates some kind of reversibility which seems to go against the motivation for PDMPs. A remark in the talk worth reiterating is the importance of figuring out which kinds of approximations are acceptable in these approximations. This somewhat connected with the next talk by Luc Rey-Bellet (University of Massachusetts) on a theory of robust approximations, in the sense of concentration inequalities (Poincaré, Gibbs, Bernstein, ...) and large deviations; with applications to rare events.

The fourth and final "hands-on" session was delivered by Miranda Holmes-Certon (New York University) on simulating under constraints, as motivated by on-going research on colloids for which the overdamped Langevin diffusion applies as an accurate model, maybe surprisingly. This perspective made for a major change from the other talks [i.e., most of the workshop!] relying on this diffusion as a mere tool. The lecture included an interesting intermede on molecular velcro made of DNA strands. Connected with this example, exotic energy landscapes are better described by hard constraints. The definition of the measure projected on the manifold defined by the constraints is obviously an important step in simulating the distribution, whose density is induced by the gradient of the constraints. The proposed algorithm is in the same spirit as the one presented by Tony Lelièvre (Ecole des Ponts) on Monday, namely moving along the tangent space then on the normal space to get back to the manifold. A solution that causes issues when the gradient is (near) zero. This great hands-on session induced massive feedback from the audience.

In the afternoon session, Gersende Fort (CNRS, Toulouse) gave a talk on a generalisation of the Wang-Landau algorithm, which modifies the true weights of the elements of a partition of the sampling space, to increase visits to low [probability] elements and jumps between modes. The idea is to rely on tempered versions of the original weights, learned by stochastic approximation; with an extra layer of adaptivity, and leading to an improvement with parameters depending on the phase of the stochastic approximation. The second talk was by David Sanders (Universidad Nacional Autónoma de México) on a recent paper published in Chaos about importance sampling for rare events of (deterministic) billiard dynamics. Diffusive limits have tails that are hard to evaluate, except by importance sampling. The last talk of the day was by Anton Martinsson (University of Edinburgh) on simulated tempering for a molecular alignment problem, with weights of different temperatures proportional to the inverse of the corresponding normalising constants, which themselves can be learned by a form of bridge sampling.

Friday. On the last day of the workshop, there were only three talks and only half of the participants were there. The first talk by Andrea Agazzi (Duke University) was on large deviations for chemical reactions. The second talk by Igor Barahona (Universidad Nacional Autónoma de México) was somewhat disconnected from the rest of the conference, dealing with textual analysis by way of algebraic data analysis. In the last and final talk, Eric Vanden-Eijden (Courant Institute) made a link between importance sampling and PDMP, as an integral can be expressed via a trajectory of a path, which appears like a generalisation of path sampling, for almost any ODE. This is also a competitor to nested sampling, waiting for the path to reach an Hamiltonian level, without some of the difficulties plaguing nested sampling like resampling, and involving continuous time processes.

4 Scientific Progress Made

The numerous discussions we had during the workshops, at the end of the talks or during the hands-sessions, led to a lively final discussion among the participants, which allowed to highlight and make precise various points related to the open problems discussed in Section 2. In many situations, the outcome of the discussion suggests a transfer of good practices or techniques from one field to another; sometimes it leads to agreeing on a common challenge, to be tackled by techniques from possibly different toolboxes.

4.1 Setting benchmarks

A first point in order to test successful methods from one field in another is to decide on a few simple benchmarks, in which both the system to be considered and the aim of the simulation are made precise. In molecular simulation, systems of interest are small Lennard–Jones clusters, Ising-like spin systems, the migration of atoms on a metallic surface, and dialanine in a water solvent for biological applications. Almost all new numerical methods are tested on at least one of these systems, and their performance is compared to other reference methods. The situation is less clear in statistics since the applications cover a wide range of situations. One case which however remains challenging (in the opinion of the researchers present during the discussion) is the metastability which arises when fitting the data to a mixture model. The conclusion from the discussion we had was that it would make a lot of sense to agree on a few benchmark models in the computational statistics community and make them available to the molecular dynamics community.

4.2 Towards even more efficient sampling methods

Although the sampling objectives are quite similar in computational statistics and molecular dynamics, the habits for doing so differ to some extent. In particular, statisticians usually prefer to avoid any bias in the sampling, and therefore rely on Metropolis–Hastings algorithms or Gibbs samplers. On the other hand, researchers in molecular dynamics tend to avoid resorting to Metropolis-like algorithms since their rejection rates decreases exponentially with the system size. This could be taken care of by a Metropolis-within-Gibbs procedure (which amounts to evolving only a few degrees of freedom at the time), but is usually not performed since this would limit the parallel efficiency of large scale codes such as LAMMPS or NAMD. It is therefore customary in molecular dynamics to rely on discretizations of stochastic differential equations which are ergodic for the target measure. These discretizations lead to biases in the target measure (which can be quantified to leading order using Talay–Tubaro expansions [37]), but the corresponding bias, of order Δt^p (with p related to the weak order of the method), is usually much smaller than the statistical error, which is of order $\sigma/\sqrt{T}$ for a given physical simulation time $T = N_{\rm iter}\Delta t$. Metropolis schemes usually increase the variance σ due to the stagnation incurred by the rejections, and are therefore discarded in molecular dynamics. In order to estimate the bias and make sure that it is sufficiently small, one can rely on extrapolation methods where simulations are performed for various timesteps. The conclusion of the discussion was that this methodology should also be tested in a Bayesian statistics framework.

Another point which was raised during some talks was the need for good importance sampling functions to facilitate the exploration of modes. One natural approach, as in Wang–Landau algorithms, is to bias by the density of states which amounts to consider the energy or log-likelihood as a reaction coordinate/summary statistics. The problem with this approach is that it requires defining a priori bounds on the range of the log-likelihood – a difficult task in general, in particular for the upper bound. In molecular dynamics, mostly reaction coordinates based on chemical intuition are used. These reaction coordinates, although not perfect, often lead to reasonably efficient importance sampling functions, known as free energies. There is currently no such intuitive perception in Bayesian statistics. In the sampling of mixture models in Bayesian statistics, however, some studies suggested that hyperparameters controlling the variance of Gaussian modes may be a good variable to act upon [9] – the idea being that modes with large variances can easily move around and hence change the value of their centers. The conclusion on this topic is that further time should be invested in trying out various reaction coordinates and computing the associated free energies to bias the sampling of multimodal target measures in computational statistics.

4.3 Beyond sampling configurations and parameters

The primary concern of many researchers in molecular dynamics shifted towards sampling paths or realizations of a given dynamics – for instance trajectories linking configurations located in two different metastable modes of the target distribution, so-called reactive trajectories in analogy with paths in chemical reactions. There are by now various methods to this end, such as the accelerated molecular dynamics proposed by Arthur Voter at the end of the nineties (see the mathematical review [23]), or splitting techniques such as the Adaptive Multilevel Splitting method [8]. The techniques proposed by Arthur Voter do not require to know in advance the target mode when starting in a metastable state. Once a reactive trajectory is found, it is possible to examine the behavior of some function of interest along the path, for instance the log-likelihood or another summary statistics/reaction coordinate. This allows for instance to determine whether the barrier preventing the transition is of energetic type (i.e. the log-likelihood decreases a lot at the transition point separating the two metastable modes), or of entropic type (i.e. the log-likelihood does not change a lot, the metastability arising from a constriction of the available configuration space). Entropic barriers are often encountered in complex systems with many degrees of freedom. In this case, the system typically has enough energy to overcome the energetic barriers it can encounter, but has not, somehow, got its energy concentrated in the right modes or directions. It is expected that entropic barriers increase with the dimensionality of the system (think of a random walk in a high-dimensional space).

Of course, in real systems of interest, no barrier is either purely entropic or energetic. Having some information on the nature of the barrier is however of prime interest in order to select the best numerical method to use, in particular when considering tempering/annealing methods. Indeed, when the barrier is energetic, the exit time out of metastable states grows exponentially with the inverse temperature, and some methods specifically rely on this behavior and/or on the typical shape of the paths out of the well (as given by the Freidlin–Wentzell theory of large deviations). On the other hand, exit times scale linearly with the inverse temperature for purely entropic barriers, and the typical exit paths have the same structure for all temperatures. Tempering is not really an option in this situation – think of a Brownian motion in high-dimension, where increasing tempering simply amounts to rescaling time. It should be more efficient in this case to add some drift to the dynamics in order to bias it towards exit points.

In any case, it seems that an analysis of the nature of the obstruction to efficient sampling in Bayesian statistics (as provided by the nature of the barrier) would be an interesting task in order to have a better vision on the type of enhanced sampling techniques to use.

4.4 Theoretical convergence results

On the theoretical side, several talks reported results on the convergence of PDMPs. It is interesting to note that approaches based on the functional analytic study of generators, standard in molecular dynamics, have percolated in computational statistics. Some work however remains to be done, in particular to better understand the role/necessity of velocity refreshments. In addition, general frameworks have been developed to write more general PDMP models (beyond zigzag, bouncy particle sampler and randomized HMC), such as event chain Monte Carlo [30]. These research efforts can be performed using the complementary tools and viewpoints from both computational statistics and computational statistical physics. As a side note, many efforts are also invested in using a set of assumptions closer to actual models (e.g. singular potentials in molecular dynamics, heavy tails in statistics, etc).

4.5 Miscellaneous topics

Is (weak) identifiability an issue? Researchers in computational statistics emphasize the identifiability of the models, i.e. the fact that the mapping from parameters to the results is as non-degenerate as possible, ideally one-to-one. This issue is on the other hand not a concern in molecular dynamics (and other fields such as machine learning, where many sets of weights in neural networks for instance can give the same results). There is no clear-cut research at the moment on the necessity/importance/relevance of identifiability.

Uncertainty quantification. Model misspecification is an emerging topic in molecular dynamics. This was an explicit topic in at least two presentations during this week, but was also implicit in others. The

discussions indicated that efforts should clearly be invested in this direction, possibly building upon results from computational statistics like those obtained by Kennedy and O'Hagan [22].

5 Outcome of the Meeting

The aim of the workshop was to gather a mixed audience made both of statisticians who are using and developing computational methods, researchers involved in computational statistical mechanics and its applications (e.g., materials science, biophysics), and of applied mathematicians studying numerical methods used in the field of application from a mathematical viewpoint. Gathering assorted practitioners (physicists, statisticians, and mathematicians) alongside theoreticians clearly fostered new insights and cross-fertilization.

We also made an effort to promote the presence of younger researchers at this workshop. It is indeed paramount that young researchers in our research areas are exposed as early as possible to researchers from other research areas in order to start looking for possible numerical methods and analysis outside their immediate scientific surrounding. Changing the habits of the younger researchers is one of the most performant ways of cross-fertilizing fields. The younger participants were particularly active and interacted a lot with more senior researchers.

All participants very much enjoyed the workshop, and many of them expressed an enthusiastic interest for a follow-up event, which would again include "hands-on sessions" or "tutorials". An idea is to have such sessions in parallel, e.g. tutorial from a researcher in statistics aimed at an audience of researchers in molecular dynamics, and conversely. The topics which emerged from the discussion held during the week include, for researchers in computational statistical physics,

- an introduction to statistical modeling (types of prior, likelihood, inclusion of latent variables of interest in the sampling, etc);
- a tutorial on model misspecification and the associated uncertainty quantification;
- a "hands-on" on ABC and the use of surrogate/coarse-grained models in general (as arising in Sequential Monte Carlo);

as well as, for researchers in statistics,

- a "hands-on" on the computation of free energy based importance sampling;
- a "hands-on" on the computation of reactive paths.

Following up on a series of conferences at the interface of molecular dynamics and computational statistics ("Computational methods for statistical mechanics — at the interface between mathematical statistics and molecular simulation" organized in 2014 by Carsten Hartmann, Greg Pavliotis, and Gabriel Stoltz at ICMS in Edinburgh, UK; and the 2016 workshop "COmputational Statistics and MOlecular Simulation" organized by Arnaud Guyader, Tony Lelièvre and Gabriel Stoltz in Paris, France), it seems relevant to plan an event in the near future where these topics would be adressed.

References

[1] H.C. Andersen, Molecular dynamics simulations at constant pressure and/or temperature, *J. Chem. Phys.* **72**(4) (1980), 2384–2393.

[2] R. Balian, *From Microphysics to Macrophysics. Methods and Applications of Statistical Physics*, volume I - II, Springer, 2007.

[3] C. H. Bennett, Efficient estimation of free energy differences from Monte-Carlo data, *J. Comput. Phys.* **22** (1976), 245–268.

[4] J. Bierkens, P. Fearnhead and G. Roberts, The Zig-Zag process and super-efficient sampling for Bayesian analysis of big data, to appear in *Annals of Statistics* (2018).

[5] A. Bouchard-Côté, S. J. Vollmer and A. Doucet, The bouncy particle sampler: a nonreversible rejection-free Markov chain Monte Carlo kethod, *Journal of the American Statistical Association* **113**(522) (2018), 855–867.

[6] R. E. Caflisch, Monte Carlo and quasi-Monte Carlo methods, *Acta Numerica* **7** (1998), 1–49.

[7] F. Cérou and A. Guyader, Adaptive multilevel splitting for rare event analysis, *Stoch. Anal. Appl.* **25** (2007), 417–443.

[8] F. Cérou, A. Guyader, T. Lelièvre and D. Pommier, A multiple replica approach to simulate reactive trajectories, *J. Chem. Phys.* **134** (2011), 054108.

[9] N. Chopin, T. Lelièvre and G. Stoltz, Free energy methods for Bayesian inference: Efficient exploration of univariate Gaussian mixture posteriors, *Statist. Comput.* (2012) **22**, 897–916.

[10] M. H. A. Davis, Piecewise-deterministic Markov processes: a general class of nondiffusion stochastic models, *J. Roy. Statist. Soc. Ser. B*, 46(3) (1994), 353-–388.

[11] S. Duane, A. Kennedy, B. J. Pendleton, and D. Roweth, Hybrid Monte Carlo, *Physics Letters B*, 195(2) (1987), 216–222.

[12] P. Dupuis and H. Wang, Importance sampling, large deviations, and differential games, *Stoch. and Stoch. Reports* **76** (2004), 481–508.

[13] A. Faggionato, D. Gabrielli and M. Ribezzi Crivellari, Non-equilibrium Thermodynamics of Piecewise Deterministic Markov Processes, *J. Stat. Phys.* **137**(2) (2009), 259–304.

[14] W. Fleming, Exit probabilities and optimal stochastic control, *Appl. Math. Optim.* **4** (1977), 329–346.

[15] W. Fleming and H. Soner, *Controlled Markov processes and viscosity solutions*, Springer, 2006.

[16] G. Fort, B. Jourdain, T. Lelièvre and G. Stoltz, Convergence and efficiency of adaptive importance sampling techniques with partial biasing, *J. Stat. Phys.* 171(2) (2018), 220–268.

[17] M. Freidlin and A. Wentzell, *Random Perturbations of Dynamical Systems*, Springer, 1984.

[18] M. Girolami and B. Calderhead, Riemann manifold Langevin and Hamiltonian Monte Carlo methods, *Journal of the Royal Statistical Society, Series B (Methodological)*, 73(2) (2011), 1–37.

[19] W. K. Hastings, Monte Carlo sampling methods using Markov chains and their applications, *Biometrika* **57** (1970), 97–109.

[20] C.-R. Hwang, S.-Y. Hwang-Ma and S. J. Sheu, Accelerating Gaussian diffusions, *Ann. Appl. Probab.* **3** (1997), 897–913.

[21] C.-R. Hwang, S.-Y. Hwang-Ma and S. J. Sheu, Accelerating diffusions, *Ann. Appl. Probab.* **15** (2005), 1433–1444.

[22] M.C. Kennedy and A. O'Hagan, A., Bayesian calibration of computer models, *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, **63** (2001), 425–464.

[23] T. Lelièvre, Accelerated dynamics: Mathematical foundations and algorithmic improvements, *Eur. Phys. J. Special Topics* **224** (2015), 2429-2444.

[24] T. Lelièvre, F. Nier and G. A. Pavliotis, Optimal non-reversible linear drift for the convergence to equilibrium of a diffusion, *J. Statist. Phys.* **152** (2013), 237–274.

[25] T. Lelièvre, M. Rousset and G. Stoltz, Hybrid Monte Carlo methods for sampling probability measures on submanifolds, *arXiv preprint* **1807.02356** (2018).

[26] T. Lelièvre and G. Stoltz, Partial differential equations and stochastic methods in molecular dynamics, *Acta Numerica* **25** (2016), 681–880.

[27] J.-M. Marin, P. Pudlo, C. Robert, and R. Ryder, Approximate Bayesian computational methods, *Stat. Comput.* **22** (2012), 1167–1180.

[28] B. Mehlig, D.W. Heermann, and B.M. Forrest, Hybrid Monte Carlo method for condensed-matter systems, *Phys. Rev. B*, **45**(2) (1992), 679.

[29] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller, Equations of state calculations by fast computing machines, *J. Chem. Phys.* **21**(6) (1953), 1087–1091.

[30] M. Michel, S. C. Kapfer and W. Krauth, Generalized event-chain Monte Carlo: Constructing rejection-free global-balance algorithms from infinitesimal steps, *J. Chem. Phys.* **140**(5) (2014), 054116.

[31] P. Monmarché, Piecewise deterministic simulated annealing, *ALEA Lat. Am. J. Probab. Math. Stat.* **13**(1) (2016) 357–398.

[32] E. A. J. F. Peters and G. de With, Rejection-free Monte-carlo sampling for general potentials, *Phys. Rev. E* **85** (2012), 026703.

[33] L. Rey-Bellet and K. Spiliopoulos, Irreversible Langevin samplers and variance reduction: A large deviations approach, *Nonlinearity* **28** (2015), 2081–2103.

[34] G. O. Roberts and J. S. Rosenthal, Optimal scaling of discrete approximations to Langevin diffusions, *J. Roy. Stat. Soc. B* **60** (1998), 255–268.

[35] P. J. Rossky, J. D. Doll, and H. L. Friedman, Brownian dynamics as smart Monte Carlo simulation, *J. Chem. Phys.* **69**(10) (1978), 4628–4633.

[36] M. R. Shirts and J. D. Chodera, Statistically optimal analysis of samples from multiple equilibrium states, *J. Chem. Phys.* **124**(12) (2008), 124105.

[37] D. Talay and L. Tubaro, Expansion of the global error for numerical schemes solving stochastic differential equations, *Stoch. Proc. Appl.* **8** (1990), 94–120.

[38] E. Zappa, M. Holmes-Cerfon and J. Goodman, Monte Carlo on manifolds: sampling densities and integrating functions, *Communications in Pure and Applied Mathematics* (2018), in press.