



HAL
open science

Pushing the limits of Multiple-Timestep Strategies for Polarizable Point Dipole Molecular Dynamics

Louis Lagardère, Félix Aviat, Jean-Philip Piquemal

► **To cite this version:**

Louis Lagardère, Félix Aviat, Jean-Philip Piquemal. Pushing the limits of Multiple-Timestep Strategies for Polarizable Point Dipole Molecular Dynamics. *Journal of Physical Chemistry Letters*, 2019, 10 (10), pp.2593–2599. 10.1021/acs.jpcllett.9b00901 . hal-02119888

HAL Id: hal-02119888

<https://hal.science/hal-02119888>

Submitted on 14 May 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Pushing the limits of Multiple-Timestep strategies for Polarizable Point Dipole Molecular Dynamics

Louis Lagardère,^{*,†,‡} Félix Aviat,^{¶,‡} and Jean-Philip Piquemal^{*,¶,§}

[†]*Sorbonne Université, Institut Parisien de Chimie Physique et Théorique, FR2622 CNRS, F-75005, Paris, France*

[‡]*Sorbonne Université, Institut des Sciences du Calcul et des Données, F-75005, Paris, France*

[¶]*Sorbonne Université, Laboratoire de Chimie Théorique, UMR 7616 CNRS, F-75005, Paris, France*

[§]*Institut Universitaire de France, F-75005, Paris, France*

E-mail: louis.lagardere@sorbonne-universite.fr; jean-philip.piquemal@sorbonne-universite.fr

Published in *J. Phys. Chem. Lett.*, 2019, 10, 2593–2599

DOI:10.1021/acs.jpcllett.9b00901

Abstract

We propose an incremental construction of multi-timestep integrators to accelerate polarizable point dipole molecular dynamics while preserving sampling efficiency. We start by building integrators using frequency-driven splittings of energy terms and a Velocity-Verlet evaluation of the most rapidly varying forces, and compare a standard bonded/non-bonded split to a three-groups split dividing non-bonded forces (including polarization) into short- and long-range contributions. We then introduce new

approaches by coupling these splittings to Langevin dynamics and to Leimkuhler’s BAOAB integrator in order to reach larger time-steps (6 fs) for long-range forces. We further increase sampling efficiency by: i) accelerating the polarization evaluation using a fast/non-iterative Truncated Conjugate Gradient (TCG-1) as short-range solver; ii) pushing the outer time-step to 10 fs using hydrogen mass repartitioning. The new BAOAB-RESPA1 integrators demonstrate up to a 7-fold acceleration over standard 1 fs (Tinker-HP) integration and reduce the performance gap between polarizable and classical force fields while preserving static and dynamical properties.

The most straightforward way to speedup molecular dynamics (MD)^{1,2} is to use larger time-steps. In this context, multi-timestep schemes emerged³, but the largest usable time-step is limited by resonance effects^{4,5}. As pointed out by various authors, it is possible to overcome these effects by using modified dynamics that still sample the correct measure, but these solutions alter the dynamical properties (Generalized Langevin Equation (GLE)⁶, stochastic isokinetic extended phase-space algorithm⁷⁻⁹). However, in practice, one would like to accelerate MD while also preserving the dynamic^{6,10}. This letter addresses this problem in the particular context of polarizable force fields (PFF)^{11,12}. This class of methods is more computationally expensive than classical force fields (FF) because of the need to evaluate a many-body polarizable energy^{13,14}. Multi-timestepping is therefore essential. The general consensus to ensure conserved properties is to limit the use of reversible Reference System Propagator Algorithm (**RESPA**) integrators³ to a bonded/non-bonded forces split and to use a 2 fs time-step for the non-bonded forces. Further splitting of the non-bonded forces is not straightforward, because of the many-body nature of polarization, but has been shown to be applicable^{8,15}. Indeed one can define a short-range polarization energy and evaluate, at an outer-timestep, the slowly varying difference between the actual polarization energy (and forces) and the short-range ones. More precisely, one has to evaluate both the short-range and total polarization terms at these outer time-steps. The reduced computational cost of the short-range polarization contribution and the less frequent evaluation of

the total one effectively reduces the computational effort. Since the upper limits of these strategies have not yet been evaluated by the community, we will, in this letter, assess this frontier to improve simulation performances while respecting two important constraints: i) the mandatory need to preserve static and dynamical properties; ii) the possibility of a black-box implementation allowing for strong computational speedups without dependence to the studied system. In everything that follows tests have been made using the AMOEBA polarizable force field¹⁶ and the Tinker-HP software¹⁷. Technical details as well as various algorithmic setups are provided in Supplementary Information (see section S1). A summary of our incremental strategy is depicted on Figure 1. Interested developers can also look at the code that will be available on the Tinker-HP website¹⁸ and later on Github¹⁹.

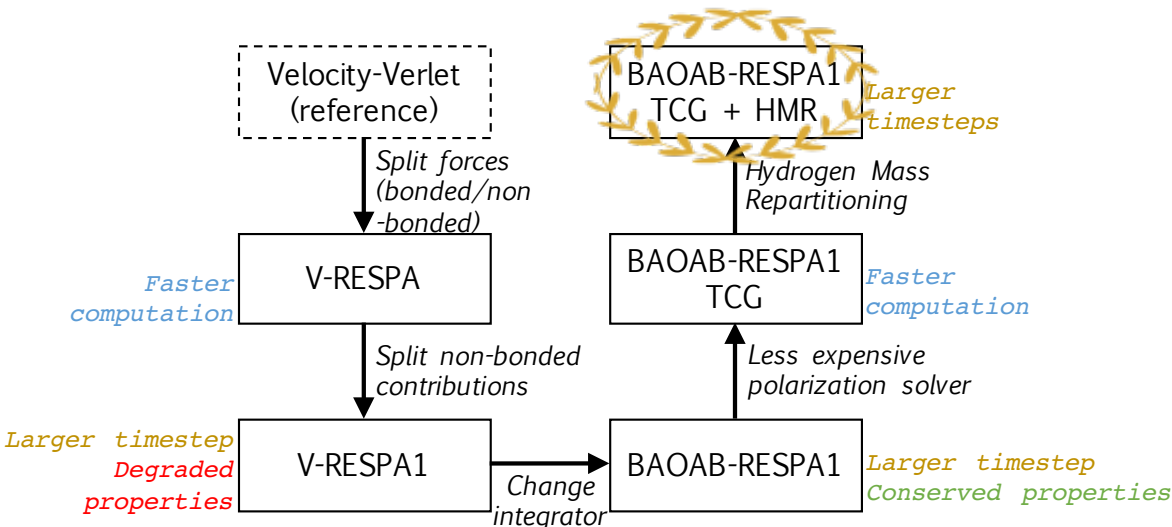


Figure 1: 'V-' (respectively 'BAOAB-') indicates that the numerical integration scheme is Velocity-Verlet²⁰ (respectively BAOAB^{21,22}). 'RESPA' and 'RESPA1' respectively mean the RESPA single-split (bonded vs. non-bonded) strategy³ and the RESPA1 double-split (bonded, short-range non-bonded, long-range non-bonded) one²³. 'TCG' is the acronym for Truncated Conjugate Gradient, a fixed-cost non-iterative polarization solver²⁴. 'HMR' stands for Hydrogen Mass Repartitioning¹⁰, implemented to avoid high-frequency motions.

A popular integrator: V-RESPA. Let us first evaluate the limits of the bonded/non-bonded RESPA integrators for which all the bonded terms are evaluated within a Velocity-

Verlet²⁰ loop (denoted as **V-RESPA** in the rest of the text) every 0.25 fs and for which the non-bonded terms (van der Waals, electrostatics and polarization) are evaluated first at 2 and then at 3 fs. To assess the accuracy of the associated integrators we ran simulations of 2 ns in the NVT ensemble at 300 K with two test systems: a cubic box of 500 water molecules, with a 24.66 Angstroms edge, and a 9737 atoms box with edges of 54.99×41.91×41.91 Angstroms containing a solvated protein (the ubiquitin). In both cases, periodic boundary conditions for electrostatics and polarization were evaluated with Smooth Particle Mesh Ewald (SPME)^{17,25,26} with standard parameters (see SI) as we chose a Pre-conditioned Conjugate Gradient (**PCG**) polarization solver using a diagonal preconditionner and a 10^{-5} convergence threshold^{13,14}. For each of these systems and for each integrator, we computed various static observables: average potential energy, average polarization energy and for the bulk water system, the radial distribution functions. In this last case, we also computed the self-diffusion coefficient, a dynamical property evaluated with the Einstein formula by averaging over a series of time origins²⁷. The self-diffusion coefficient is known to have a size dependency vanishing at the infinite size limit²⁷, but here these values are only used as means of comparison between integrators, hence these corrections were not applied. These tests are performed in the canonical ensemble for which the choice of the thermostat impacts the dynamics of the system. We ran these tests using the global velocity rescaling thermostat developed by Bussi *et al.* with a relaxation constant of 0.2 ps, for which the dynamical properties are close to the one obtained with pure hamiltonian dynamics²⁸. These values have been compared to the ones obtained with a Velocity-Verlet integrator used at a 0.5 fs time-step, which can be considered as a reference. In the rest of the text, we will be denoting the different time-step lengths as a/b , a being the bonded terms one and b the non-bonded terms one (both in fs).

For both systems, the V-RESPA integrator where non-bonded forces are evaluated at 2 fs gives similar results as the reference (within statistical uncertainty) with a difference of less than two percents for average energies (see Tables 1-3 of the SI). With an outer time-step

of 3 fs, the error on the total potential energy is still satisfactory (around one percent), but the error on the polarization energy grows significantly (more than 2.5%). This advocates for a careful use of this setup.

Concerning O-O radial distribution function for water, no significant differences with the reference Velocity-Verlet (0.5fs) ones can be observed among these different methods (see Figure 2).

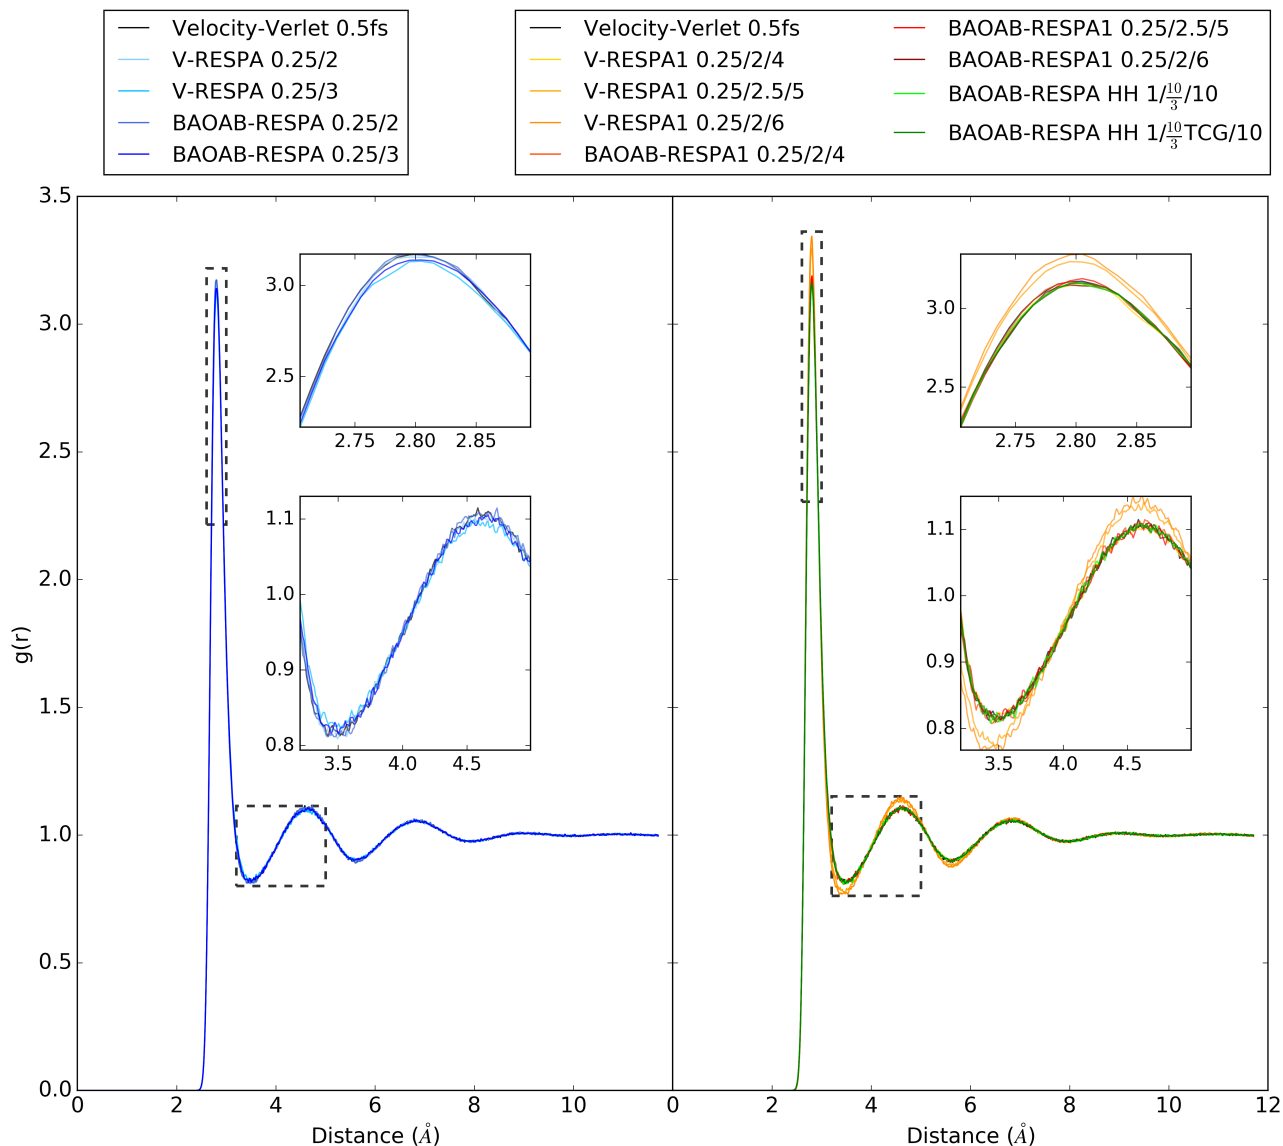


Figure 2: Oxygen-oxygen radial distribution function for various integrators (see Text for notations). Radial distributions appear correct with most of the setups. However, degraded results are obtained with the V-RESPA1 integrators using large outer-timesteps beyond 4 fs.

The self-diffusion coefficients of water is also nicely preserved for the V-RESPA integrator with a 2 fs outer time-step, though it is slightly off (around eight percent) with a 3 fs outer time-step (see Table 2, SI).

Further range-separation in polarizable force fields: V-RESPA1. As a second part, we will now evaluate the limits of other RESPAs integrators, for which the non-bonded terms are further split in two parts, the short- and long-range. We are now considering three terms: the bonded, the non-bonded short-range, and the non-bonded long-range terms. Regarding the split of the electrostatics and the polarization energies, we chose to use the **RESPA1** logic²³, where the short-range part of the electrostatic (and polarization) energy is defined as the short-range part of the real space SPME energy. In this case, it has been shown that the stability of the integrator is less dependant on the smoothing parameters used to switch between short- and long-range²⁹. Details on the definition of these short and long range forces as well as these smoothing parameters can be found in SI. We test various setups within this context: the bonded forces are always evaluated every 0.25 fs, but the short-range non-bonded ones are either evaluated every 2 or 2.5 fs, and the time-step of the long-range forces (that has to be a multiple of the previous one) is either 4, 5 or 6 fs. In the rest of the text, these integrators will be denoted as **V-RESPA1**. We will be denoting the different time-steps lengths of the integrators as $a/b/c$, a being the bonded terms time-step length, b the short-range non-bonded and c the long-range ones (all in fs). For the bulk water system (see Table 4 in SI), we observe that both the average potential and polarization energies are preserved within two percents of the reference value for the 0.25/2/4 and the 0.25/2.5/5 setups, but that the average polarization energy is more than two percents off for the 0.25/2/6 setup. Concerning the radial distribution functions of water, it is clear that only the 0.25/2/4 integrator gives satisfactory result as other choices diverge from the reference, as can be seen in Figure 2. Furthermore, if the self-diffusion coefficient is stable for the 0.25/2/4 integrator (see Table 5, SI), it exhibits a dramatic decrease for the other ones (falling at 1.34 instead

of 2.08 for the 0.25/2/6 setup). This shows not only that the dynamical properties are not well preserved with these setups, but also that the computational gains expected due to the use of a larger time-step are counterbalanced by a lower sampling rate^{6,10}.

Indeed, as pointed out by Berendsen^{6,10}, such a decrease in the self-diffusion coefficient is expected to reduce the sampling efficiency by a similar amount because it is associated to an increase of water viscosity and thus to a slowing down of large scale motions. For the solvated ubiquitin, it is also clear that only the 0.25/2/4 setup corresponds to a satisfactory accuracy, as the other ones give average potential and polarization energies off by more than three percents (see Table 6, SI).

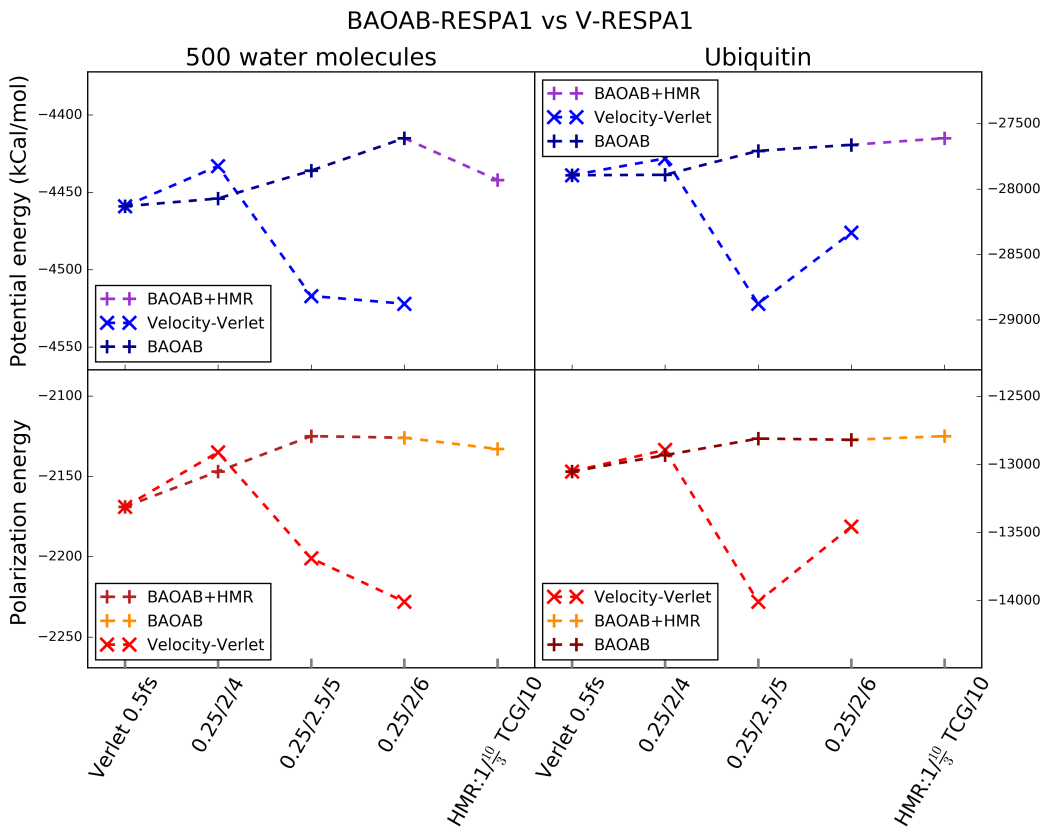


Figure 3: Average potential and polarization energies (in kcal/mol) for a 500-molecule water box and solvated ubiquitin computed using various integrators.

Recovering accuracy through Langevin dynamics: the new BAOAB-RESPA1 integrator for polarizable force fields. Thirdly, another way to sample the canonical

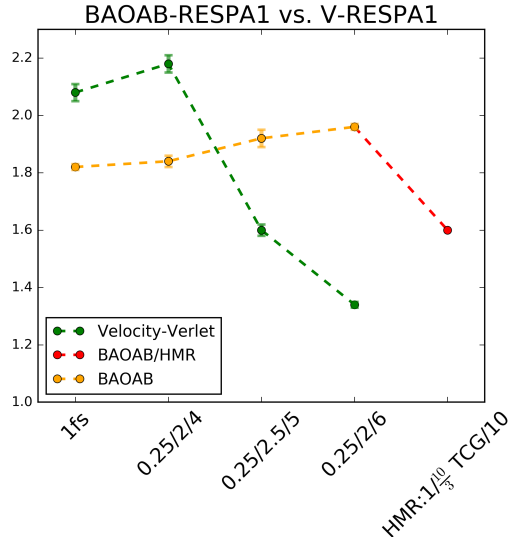


Figure 4: Self-diffusion (D) coefficient for various integrators (in $10^{-5} \text{ cm}^2\text{s}^{-1}$).

ensemble is Langevin dynamics where the coupling to a heat bath is made through additional local frictions and dissipative stochastic terms. The dynamics is then known to be altered compared to pure Hamiltonian one, but this impact is expected to be small for relatively small friction constant (less than 1 ps^{-1}). In this context, Leimkuhler and collaborators proposed an integrator for Langevin dynamics based on operator splitting and a particular ordering of the terms of the equations of motion named **BAOAB**^{21,22}. They showed in various contexts²¹ that this integrator has an improved accuracy for configurational sampling compared to other ones. We thus also tested the previously presented splittings using this new integrator, with a 1 ps^{-1} friction constant (they will be denoted **BAOAB-RESPA** and **BAOAB-RESPA1** in the rest of the text ; see SI for an additional description of these integrators) and noted a significant improvement in terms of accuracy (reported on Tables 7-9 of SI). Indeed, for the bulk water system (Table 7 of SI), the errors for the BAOAB-RESPA integrator with a bonded/non-bonded split and a 3 fs outer time-step is limited to less than one percent for the average total potential energy and two percents for the average polarization energy, staying in both case within statistical error, compared to 1.1 and 2.8% for the Velocity-Verlet based integrator. Figure 2 also shows an improved agreement with

the reference for the water radial distribution compared to RESPA. The same behaviour can be observed on the solvated ubiquitin for which both these values stay respectively below one and two percents – within statistical error (see Table 9, SI). Because the dynamics is modified when running NVT trajectories with Langevin, even if the differences are expected to be small for a small friction, comparing values of the self-diffusion coefficients obtained with these integrators only makes sense by taking as a reference a numerical scheme integrating Langevin dynamics with conservative parameters. This is why we chose, as a reference for these values, the ones obtained with a 1 fs BAOAB integrator and a 1 ps^{-1} friction, which as expected gives a self-diffusion close to the reference Velocity-Verlet one (1.82 versus 2.08). For the BAOAB-RESPA integrators, we see that errors on the self-diffusion coefficients (see Table 8, SI) are limited to six percents with a 3 fs outer time-step compared to eight percents with a similar time-step and a Velocity-Verlet inner loop. The better performances of BAOAB-based integrators with respect to the Velocity-Verlet ones becomes obvious within a RESPA1 split. Indeed, we computed the same observables for the BAOAB-RESPA1 integrators (see Tables 10-12, SI and Figures 3-4), equivalent to the V-RESPA1 integrators, and we see that the average potential and polarization energies are strikingly more stable and always within two percent error with respect to their reference value, that is to say within statistical uncertainty. Similar comments can be made on the radial distribution functions: unlike the Velocity-Verlet integrators, they almost perfectly overlap with their references even with a 5 or 6 fs outer time-step. When using the Velocity-Verlet based RESPA1 integrators with an outer-time-step larger than the most conservative (4 fs) one, the diffusion coefficient showed a dramatic decrease (see Figure 3, and Table 5 in SI). Yet in the BAOAB-RESPA1 case, this dynamical observable remains far more stable for all integrator setups: even for the choice of evaluating long-range non-bonded forces every 6 fs, the error is less than eight percents, whereas it exceeds thirty-five percents in the equivalent V-RESPA1 setup. This highlights again that BAOAB-RESPA1 integrators are not only more accurate, but also ensure a conserved sampling rate, which is not the case for the V-RESPA1 ones.

Concerning the effective speedups in our implementation, Table 1 displays the gains obtained for the BAOAB-RESPA and the BAOAB-RESPA1 integrators (compared to a regular 1 fs Velocity-Verlet). They are the same as the one obtained for the V-RESPA and the V-RESPA1 integrators. We show two entries in table 1: one where a guess based on Kolafa’s Always Stable Predictor Corrector³⁰ (ASPC) is used for the induced dipoles (standard Tinker-HP setting), but only at short-range for the RESPA1 integrators, and one where the ”direct field” guess is used^{13,14}, showing that up to a 2.53 speedup (3.7 without ASPC) is achieved. Note that for the RESPA1 schemes, additional gains are made in the long-range polarization solvers by using, at the same time-step, the short-range dipoles obtained as guess for the long range ones, effectively reducing the number of iterations required to converge. For the BAOAB based integrators, the benefits of using the RESPA1 splitting are clearly demonstrated as the 0.25/2.5/5 and 0.25/2/6 frameworks are both faster than the 0.25/3 RESPA integrator for a similar accuracy.

Speeding up BAOAB-RESPA1: TCG-1 solver for the short-range polarization and Hydrogen Mass Repartitioning.

When using a RESPA1 multi-timestep integrator and a polarizable force field, the sole purpose of the short-range polarization energy is to eliminate the high-frequency part of the total polarization energy^{8,31}. This is why an approximate but less computationally expensive and non-parametric expression of the polarization energy can be used to fill this role and provide an additional speedup. In that context we decided to use the recently introduced Truncated Conjugate Gradient (**TCG**)^{24,32} as short-range solver. TCG can be chosen to be minimal in cost (TCG-1 with a diagonal preconditioner, but without any guess and without peek step) to be the fastest possible. This coupling provides an additional computational gain at a conserved accuracy (see Tables 13-15, SI), corresponding to a final speedup of more than four times compared to a regular molecular dynamics of 1 fs with a Velocity-Verlet integrator (see Table 1). Such polarization setup offers full energy conservation and the static and dynamical properties are marginally

affected by this choice^{13,14,24}.

At this point, we reached the performance limits if one wants to preserve a tight accuracy on the dynamics. One of the most natural ways to further increase the size of the usable time-step when simulating a large biological system is to redistribute the mass of the heavy atoms on the hydrogens they are carrying (a method named Hydrogen Mass Repartitioning, or **HMR**¹⁰), thus limiting the high frequency stretching motions of these atoms while keeping the same configurational potential energy surface. In the following, we show that this redistribution allows to use even larger time-steps while maintaining satisfactory accuracy with a BAOAB-RESPA1 integrator and the same TCG-1 short-range polarization solver as before (Table 16-18, SI). As can be seen in Table 1, the approach appears to be a very good compromise: large speedups can be obtained by pushing the bonded force time-step to 1 fs, the short-range non-bonded forces time-step up to $\frac{10}{3}$ fs and the outer one up to 10 fs. Details on how the mass repartitioning is done can be found in SI (section S1). A resulting acceleration of 4.72 (6.8 without ASPC) is obtained, keeping the errors on the average energies below two percents, maintaining an accurate evaluation of radial distribution functions and a good enough evaluation of the self-diffusion coefficient so that sampling efficiency is preserved. Since PCG, as a Krylov method, is systematically improvable¹³, additional small speedups can be obtained by focusing on the long range PCG solver performances. For example, besides using a diagonal preconditioner, one could use more advanced techniques such as those proposed by Skeel³³ or by Beran³⁴. Improved performances of three to four percents are observed, reaching a global acceleration of more than seven, with the same accuracy as a 1 fs Velocity Verlet scheme without ASPC. Finally, beside the net acceleration, another advantage of the TCG use lies in the absence of use of a dipole history (as in predictor-correctors such ASPC), leading to a method free of time-reversibility and volume preservation issues²⁴. Finally, the small decrease (8 %) of the self-diffusion constant (see Figure 4) observed in the most aggressive setup has to be compared with actual available large step methods⁹ that, despite their qualities, are not able to maintain accuracy on dynamical properties provid-

ing diffusion constants reduced by a factor of 5⁹. Our approach does not suffer from these problems: it remains operational, maintaining sampling efficiency.

To illustrate the robustness of these approaches, we performed several tests taking advantage of our massively parallel AMOEBA production implementation in Tinker-HP¹⁷. First, we checked the stability of the dynamics using the fastest available setup. We provide a 15 ns simulation of ubiquitin (see SI, section S3): the potential and polarization energies normally fluctuate around their mean values, demonstrating the stability of the approach. Furthermore, we computed the average molecular dipole moments for the bulk water systems and confirmed their full stability (see SI, section S4). Second, we ran simulations on large systems of biological interest, namely the solvated dihydrofolate reductase protein (DHFR, 23358 atoms) and the solvated Satellite Tobacco Virus (STMV, 1066628 atoms). The discussed speedup of 7 (vs. a 1 fs/Velocity Verlet/PCG-10⁻⁵) is conserved as we obtained a production of 22.2 ns/day on 680 cores for DHFR and 1.2 ns/day on 10800 cores for STMV. Such results are of major interest, as a 7-fold acceleration will enable to save millions of hours of computing time while enabling long and accurate polarizable molecular dynamics studies on large systems. Finally, we computed a more involved property which is of key importance in biological simulations: hydration free energies. We applied the Bennett Acceptance Ratio method³⁵, a commonly used approach to compute free energies differences³⁶ to evaluate the solvation free energy of a sodium cation in water. Results are shown in Table 2), and practical details on the choice of alchemical free energy difference windows can be found in SI (section S1). Even for the fastest setup, the values obtained are within 0.1 kcal/mol over 89.7 kcal/mol for the reference³⁷, demonstrating the validity of these acceleration schemes and their capability to preserve accuracy.

To conclude, after examining the limits of a standard Velocity-Verlet integrator for Polarizable Force Fields used in combination with a RESPA1 split, we introduced new BAOAB-RESPA1 Langevin integrators coupled to fast short-range non-iterative TCG-1 polarization solver and Hydrogen Mass Repartitioning, achieving all together large computational

speedups. Two optimal BAOAB-RESPA1 setups were presented and compared to a 1 fs Velocity-Verlet reference: i) one (namely 0.25/2(TCG)/6) for which all properties are preserved while providing a global speedup of more than four; ii) a second ($1/\frac{10}{3}$ (TCG)/10+HMR) for which dynamical properties are slightly affected but where sampling remains efficient, offering a strong acceleration up to a seven-fold. As accuracy is maintained and sampling efficiency is preserved while being system independent, the proposed methodology can be used as a black-box in our Tinker-HP framework, benefiting from its massive parallelism implementation and offering therefore further computational gains¹⁷. Such findings are game-changing, as they extend the applicability of polarizable Molecular Dynamics to longer timescale simulations and larger systems. In practice, the resulting performance gain helps reducing the computational gap between point dipole polarizable force fields such as AMOEBA and more tractable models such as Drude³⁸ or even non-polarizable force fields such as CHARMM³⁹ or AMBER⁴⁰.

Table 1: Speedup of BAOAB-RESPA and BAOAB-RESPA1 integrators calculated with respect to the Velocity-Verlet integrator at 1fs. The types of RESPA integrators are defined by: R1=RESPA1 and R=RESPA. Speedups obtained with V-RESPA and V-RESPA1 integrators are identical.*=replacement of the PCG diagonal preconditioner by an improved technique^{33,34}, see text.

Splits	0.25/2	0.25/3	0.25/2/4	0.25/2.5/5	0.25/2/6	0.25/2(TCG)/4	0.25/2.5(TCG)/5	0.25/2(TCG)/6	$1/\frac{10}{3}/10$	$1/\frac{10}{3}$ (TCG)/10
ASPC	1.75	2.37	1.72	2.43	2.53	2.32	2.7	2.9	4	4.72 (4.91*)
No ASPC	2.53	3.42	2.5	3.5	3.7	3.4	3.9	4.2	5.8	6.8 (7.0*)
RESPA-type	R	R	R1	R1	R1	R1	R1	R1	R1(HMR)	R1(HMR)

Table 2: Hydration free energies for the Na⁺ cation

	Velocity-Verlet 0.5fs	$1/\frac{10}{3}/10$ -HMR	$1/\frac{10}{3}$ (TCG1)/10-HMR
$\Delta G_{\text{hydrat}} \text{ Na}^+$ (kCal/mol)	89.7 (+/-0.13)	89.7 (+/-0.13)	89.6 (+/-0.13)

Acknowledgement

This work was supported in part by French state funds managed by CalSimLab and the ANR within the Investissements d' Avenir program under reference ANR-11-IDEX-0004-02. We thank GENCI for computer time within the Grand Challenge on the Irène Joliot-Curie (TGCC, Bruyères le Châtel, France).

Supporting Information Available

References

- (1) Allen, M. P.; Tildesley, D. J. *Computer Simulation of Liquids*; Clarendon Press: New York, NY, USA, 1989.
- (2) McCammon, J. A.; Harvey, S. C. *Dynamics of Proteins and Nucleic Acids*; Cambridge University Press, 1987.
- (3) Tuckerman, M.; Berne, B. J.; Martyna, G. J. Reversible multiple time scale molecular dynamics. *J. Chem. Phys.* **1992**, *97*, 1990–2001.
- (4) Biesiadecki, J. J.; Skeel, R. D. Dangers of Multiple Time Step Methods. *Journal of Computational Physics* **1993**, *109*, 318 – 328.
- (5) Ma, Q.; Izaguirre, J.; Skeel, R. Verlet-I/R-RESPA/Impulse is Limited by Nonlinear Instabilities. *SIAM Journal on Scientific Computing* **2003**, *24*, 1951–1973.
- (6) Morrone, J. A.; Markland, T. E.; Ceriotti, M.; Berne, B. J. Efficient multiple time scale molecular dynamics: Using colored noise thermostats to stabilize resonances. *The Journal of Chemical Physics* **2011**, *134*, 014103.
- (7) Leimkuhler, B.; Margul, D. T.; Tuckerman, M. E. Stochastic, resonance-free multiple time-step algorithm for molecular dynamics with very large time steps. *Molecular Physics* **2013**, *111*, 3579–3594.
- (8) Margul, D. T.; Tuckerman, M. E. A Stochastic, Resonance-Free Multiple Time-Step Algorithm for Polarizable Models That Permits Very Large Time Steps. *Journal of Chemical Theory and Computation* **2016**, *12*, 2170–2180, PMID: 27054809.
- (9) Albaugh, A.; Tuckerman, M. E.; Head-Gordon, T. Combining Iteration-Free Polarization with Large Time Step Stochastic-Isokinetic Integration. *Journal of Chemical Theory and Computation* **2019**, *15*, 2195–2205, PMID: 30830768.

- (10) Feenstra, K. A.; Hess, B.; Berendsen, H. J. C. Improving efficiency of large time-scale molecular dynamics simulations of hydrogen-rich systems. *Journal of Computational Chemistry* **1999**, *20*, 786–798.
- (11) Shi, Y.; Ren, P.; Schnieders, M.; Piquemal, J.-P. *Reviews in Computational Chemistry Volume 28*; John Wiley and Sons, Inc, 2015; pp 51–86.
- (12) Gresh, N.; Cisneros, G. A.; Darden, T. A.; Piquemal, J.-P. Anisotropic, Polarizable Molecular Mechanics Studies of Inter- and Intramolecular Interactions and Ligand-Macromolecule Complexes. A Bottom-Up Strategy. *Journal of Chemical Theory and Computation* **2007**, *3*, 1960–1986, PMID: 18978934.
- (13) Lipparini, F.; Lagardère, L.; Stamm, B.; Cancès, E.; Schnieders, M.; Ren, P.; Maday, Y.; Piquemal, J.-P. Scalable evaluation of polarization energy and associated forces in polarizable molecular dynamics: I. Toward massively parallel direct space computations. *Journal of Chemical Theory and Computation* **2014**, *10*, 1638–1651.
- (14) Lagardère, L.; Lipparini, F.; Polack, E.; Stamm, B.; Cancès, E.; Schnieders, M.; Ren, P.; Maday, Y.; Piquemal, J. P. Scalable Evaluation of Polarization Energy and Associated Forces in Polarizable Molecular Dynamics: II. Toward Massively Parallel Computations Using Smooth Particle Mesh Ewald. *Journal of Chemical Theory and Computation* **2015**, *11*, 2589–2599.
- (15) Masella, M. The multiple time step r-RESPA procedure and polarizable potentials based on induced dipole moments. *Molecular Physics* **2006**, *104*, 415–428.
- (16) Ren, P.; Ponder, J. W. Polarizable Atomic Multipole Water Model for Molecular Mechanics Simulation. *The Journal of Physical Chemistry B* **2003**, *107*, 5933–5947.
- (17) Lagardère, L.; Jolly, L.-H.; Lipparini, F.; Aviat, F.; Stamm, B.; Jing, Z. F.; Harger, M.; Torabifard, H.; Cisneros, G. A.; Schnieders, M. J.; Gresh, N.; Maday, Y.; Ren, P. Y.;

- Ponder, J. W.; Piquemal, J.-P. Tinker-HP: a massively parallel molecular dynamics package for multiscale simulations of large complex systems with advanced point dipole polarizable force fields. *Chem. Sci.* **2018**, *9*, 956–972.
- (18) Tinker-HP’s Website. <http://tinker-hp.ip2ct.upmc.fr/>, Accessed:2019-04-24.
- (19) Tinker-HP’s Github. <https://github.com/TinkerTools/Tinker-HP>, Accessed:2019-04-24.
- (20) Swope, W. C.; Andersen, H. C.; Berens, P. H.; Wilson, K. R. A computer simulation method for the calculation of equilibrium constants for the formation of physical clusters of molecules: Application to small water clusters. *The Journal of Chemical Physics* **1982**, *76*, 637–649.
- (21) Matthews, C.; Leimkuhler, B. Rational Construction of Stochastic Numerical Methods for Molecular Sampling. *Applied Mathematics Research eXpress* **2012**, *2013*, 34–56.
- (22) Leimkuhler, B.; Matthews, C. Robust and efficient configurational molecular sampling via Langevin dynamics. *The Journal of Chemical Physics* **2013**, *138*, 174102.
- (23) Zhou, R.; Harder, E.; Xu, H.; Berne, B. J. Efficient multiple time step method for use with Ewald and particle mesh Ewald for large biomolecular systems. *The Journal of Chemical Physics* **2001**, *115*, 2348–2358.
- (24) Aviat, F.; Levitt, A.; Stamm, B.; Maday, Y.; Ren, P.; Ponder, J. W.; Lagardère, L.; Piquemal, J.-P. Truncated Conjugate Gradient: An Optimal Strategy for the Analytical Evaluation of the Many-Body Polarization Energy and Forces in Molecular Simulations. *Journal of Chemical Theory and Computation* **2017**, 180–190.
- (25) Essmann, U.; Perera, L.; Berkowitz, M. L.; Darden, T.; Lee, H.; Pedersen, L. G. A smooth particle mesh Ewald method. *The Journal of Chemical Physics* **1995**, *103*, 8577–8593.

- (26) Stamm, B.; Lagardère, L.; Polack, E.; Maday, Y.; Piquemal, J.-P. A coherent derivation of the Ewald summation for arbitrary orders of multipoles: The self-terms. *The Journal of Chemical Physics* **2018**, *149*, 124103.
- (27) Maginn, E. J.; Messerly, R. A.; Carlson, D. J.; Roe, D. R.; Elliott, J. R. Best Practices for Computing Transport Properties 1. Self-Diffusivity and Viscosity from Equilibrium Molecular Dynamics [Article v1.0]. *Living Journal of Computational Molecular Science* **2018**, *1*, 6324–.
- (28) Bussi, G.; Donadio, D.; Parrinello, M. Canonical sampling through velocity rescaling. *The Journal of Chemical Physics* **2007**, *126*, 014101.
- (29) Morrone, J. A.; Zhou, R.; Berne, B. J. Molecular Dynamics with Multiple Time Scales: How to Avoid Pitfalls. *Journal of Chemical Theory and Computation* **2010**, *6*, 1798–1804, PMID: 26615840.
- (30) Kolafa, J. Time-reversible always stable predictor–corrector method for molecular dynamics of polarizable molecules. *J. Comput. Chem.* **2004**, *25*, 335–342.
- (31) Liberatore, E.; Meli, R.; Rothlisberger, U. A Versatile Multiple Time Step Scheme for Efficient ab Initio Molecular Dynamics Simulations. *Journal of Chemical Theory and Computation* **2018**, *14*, 2834–2842, PMID: 29624388.
- (32) Aviat, F.; Lagardère, L.; Piquemal, J.-P. The truncated conjugate gradient (TCG), a non-iterative/fixed-cost strategy for computing polarization in molecular dynamics: Fast evaluation of analytical forces. *The Journal of Chemical Physics* **2017**, *147*, 161724.
- (33) Wang, W.; Skeel, R. D. Fast evaluation of polarizable forces. *J. Chem. Phys.* **2005**, *123*, 164107.

- (34) Nocito, D.; Beran, G. J. O. Fast divide-and-conquer algorithm for evaluating polarization in classical force fields. *J. Chem. Phys.* **2017**, *146*, 114103.
- (35) Bennett, C. H. Efficient estimation of free energy differences from Monte Carlo data. *Journal of Computational Physics* **1976**, *22*, 245 – 268.
- (36) Rackers, J. A.; Laury, M. L.; Lu, C.; Wang, Z.; Lagardère, L.; Piquemal, J.-P.; Ren, P.; Ponder, J. W. TINKER 8: A Modular Software Package for Molecular Design and Simulation. *Journal of Computational Chemistry* **2018**,
- (37) Schnieders, M. J.; Baltrusaitis, J.; Shi, Y.; Chattree, G.; Zheng, L.; Yang, W.; Ren, P. The Structure, Thermodynamics, and Solubility of Organic Crystals from Simulation with a Polarizable Force Field. *Journal of Chemical Theory and Computation* **2012**, *8*, 1721–1736, PMID: 22582032.
- (38) Jiang, W.; Hardy, D. J.; Phillips, J. C.; MacKerell, A. D.; Schulten, K.; Roux, B. High-Performance Scalable Molecular Dynamics Simulations of a Polarizable Force Field Based on Classical Drude Oscillators in NAMD. *The Journal of Physical Chemistry Letters* **2011**, *2*, 87–92, PMID: 21572567.
- (39) Huang, J.; Rauscher, S.; Nawrocki, G.; Ran, T.; Feig, M.; de Groot, B. L.; Grubmueller, H.; MacKerell, A. D., Jr. CHARMM36m: an improved force field for folded and intrinsically disordered proteins. *NATURE METHODS* **2017**, *14*, 71–73.
- (40) Maier, J. A.; Martinez, C.; Kasavajhala, K.; Wickstrom, L.; Hauser, K. E.; Simmerling, C. ff14SB: Improving the Accuracy of Protein Side Chain and Backbone Parameters from ff99SB. *Journal of Chemical Theory and Computation* **2015**, *11*, 3696–3713, PMID: 26574453.




Figure 5: TOC