

UC Berkeley

UC Berkeley Previously Published Works

Title

M-Chem: a modular software package for molecular simulation that spans scientific domains

Permalink

<https://escholarship.org/uc/item/6j88b5ss>

Journal

Molecular Physics, 121(9-10)

ISSN

0026-8976

Authors

Witek, Jagna
Heindel, Joseph P
Guan, Xingyi
[et al.](#)

Publication Date

2023-05-19

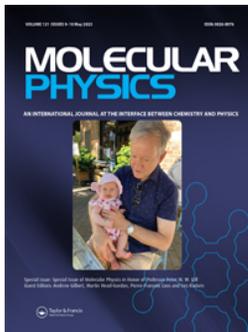
DOI

10.1080/00268976.2022.2129500

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed



Molecular Physics

An International Journal at the Interface Between Chemistry and Physics

ISSN: (Print) (Online) Journal homepage: <https://www.tandfonline.com/loi/tmph20>

M-Chem: a modular software package for molecular simulation that spans scientific domains

Jagna Witek, Joseph P. Heindel, Xingyi Guan, Itai Leven, Hongxia Hao, Pavithra Naullage, Allen LaCour, Selim Sami, M. F. S. J. Menger, D. Vale Cofer-Shabica, Eric Berquist, Shirin Faraji, Evgeny Epifanovsky & Teresa Head-Gordon

To cite this article: Jagna Witek, Joseph P. Heindel, Xingyi Guan, Itai Leven, Hongxia Hao, Pavithra Naullage, Allen LaCour, Selim Sami, M. F. S. J. Menger, D. Vale Cofer-Shabica, Eric Berquist, Shirin Faraji, Evgeny Epifanovsky & Teresa Head-Gordon (2023) M-Chem: a modular software package for molecular simulation that spans scientific domains, Molecular Physics, 121:9-10, e2129500, DOI: [10.1080/00268976.2022.2129500](https://doi.org/10.1080/00268976.2022.2129500)

To link to this article: <https://doi.org/10.1080/00268976.2022.2129500>



© 2022 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group



[View supplementary material](#)



Published online: 07 Oct 2022.



[Submit your article to this journal](#)



Article views: 1124



[View related articles](#)



[View Crossmark data](#)

M-Chem: a modular software package for molecular simulation that spans scientific domains

Jagna Witek ^a, Joseph P. Heindel ^{a,b}, Xingyi Guan ^{a,b}, Itai Leven^a, Hongxia Hao^a, Pavithra Naullage^a, Allen LaCour^{a,b}, Selim Sami ^a, M. F. S. J. Menger ^d, D. Vale Cofer-Shabica ^e, Eric Berquist ^f, Shirin Faraji ^d, Evgeny Epifanovsky^f and Teresa Head-Gordon ^{a,b,c}

^aDepartment of Chemistry, Kenneth S. Pitzer Theory Center, Berkeley, CA, USA; ^bChemical Sciences Division, Lawrence Berkeley National Laboratory, Berkeley, CA, USA; ^cDepartments of Bioengineering and Chemical and Biomolecular Engineering, University of California, Berkeley, CA, USA; ^dStratingh Institute for Chemistry, University of Groningen, Groningen, The Netherlands; ^eDepartment of Chemistry, University of Pennsylvania, Philadelphia, PA, USA; ^fQ-Chem, Inc., Pleasanton, CA, USA

ABSTRACT

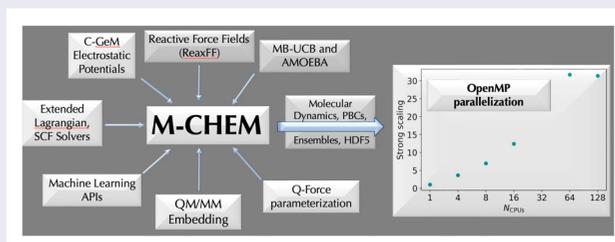
We present a new software package called M-Chem that is designed from scratch in C++ and parallelised on shared-memory multi-core architectures to facilitate efficient molecular simulations. Currently, M-Chem is a fast molecular dynamics (MD) engine that supports the evaluation of energies and forces from two-body to many-body all-atom potentials, reactive force fields, coarse-grained models, combined quantum mechanics molecular mechanics (QM/MM) models, and external force drivers from machine learning, augmented by algorithms that are focused on gains in computational simulation times. M-Chem also includes a range of standard simulation capabilities including thermostats, barostats, multi-timestepping, and periodic cells, as well as newer methods such as fast extended Lagrangians and high quality electrostatic potential generation. At present M-Chem is a developer friendly environment in which we encourage new software contributors from diverse fields to build their algorithms, models, and methods in our modular framework. The long-term objective of M-Chem is to create an interdisciplinary platform for computational methods with applications ranging from biomolecular simulations, reactive chemistry, to materials research.

ARTICLE HISTORY

Received 18 July 2022
Accepted 6 September 2022

KEYWORDS

Machine learning; force fields; molecular dynamics; QM/MM; simulation software



1. Introduction

Computer simulations are undoubtedly established as one of the main pillars of modern research, accelerating scientific progress in the molecular sciences including biomolecular modeling, computational quantum chemistry, and materials science research. Each of these scientific domains are supported by a variety of codes that incorporate a combination of software efficiencies and new algorithms that avoid bottlenecks and/or reduce unfavourable scaling, thereby allowing for exploration of ever increasing system sizes and longer timescales. This

opens up the question – why is there a need for another simulation software package?

We believe there are three primary reasons. The first is that simulation codes are siloed into scientific domains that insufficiently bridge across different simulation communities. Software packages such as Amber [1], Charmm [2], NAMD [3], and GROMACS [4] are popular in the biomolecular simulation community, whereas LAMMPS is viewed primarily as a materials research code [5], and many quantum chemistry codes [6] are not particularly compatible with classical simulations and/or require

CONTACT Teresa Head-Gordon  thg@berkeley.edu

 Supplemental data for this article can be accessed here. <https://doi.org/10.1080/00268976.2022.2129500>

plane wave implementations to describe electronic structure. A second reason is that software implementations for advanced force fields comprising many-body potentials or reactive force field models are insufficiently optimised, thereby limiting their wider adoption by consumers of molecular simulation software [7]. Finally it is well appreciated that large software packages benefit greatly from ground up rewrites, creating new opportunities for software efficiency, ease of development and use, exploiting new hardware paradigms, and seamless software integration and interoperability.

In that spirit we present a new software package called M-Chem that is designed to facilitate efficient molecular simulations on shared-memory multi-core architectures to address these issues. M-Chem, like all new software packages, is incomplete and yet is rapidly expanding its simulation capabilities. Hence some users may find that their research objectives may well be met with many of the features that we describe here. At present M-Chem enables users to run efficient MD simulations for molecular liquids using pairwise to many-body non-reactive force fields [8–10], ReaxFF simulations [11,12], coarse-grained electron (C-GeM) [13] and monoatomic-water (mW) simulations [14], and forces derived from machine learning models created with PyTorch by utilising the Torch C++ library. Standard, but critical, MD algorithms include periodic boundary conditions, particle mesh Ewald [15], targeted ensembles such as NVT and NPT using velocity scaling Berendsen thermostats, as well as extended system Nose-Hoover thermostats and barostats [16]. There is an emergent analysis toolkit of simulated properties, including the ability to generate electrostatic potentials at protein surfaces using C-GeM, whose accuracy is comparable to some of the best ab initio charge partitioning methods at orders of magnitude less cost [17]. Furthermore, we introduce Q-Force as an automated approach to small molecule force field development for M-Chem utilising ab initio quantum chemistry calculations [18]. Both standard conjugate gradient self-consistent field (CG-SCF) [19] and extended Lagrangian formulations [20] are available to solve the many-body energy and force solutions, including iEL/0-SCF [21] and SC-XLMD [22] that are iteration-free algorithms for classical polarisable and ReaxFF models, respectively. M-Chem also has multi-scale capabilities with mechanical and simple electrostatic (point charge) embedding, exploiting some of the common libraries between M-Chem and Q-Chem such that the interface is internal to M-Chem. All of these features are described in some detail below.

Finally, M-Chem is an attractive developer platform and environment. The M-Chem package is entirely written in C++, with an overall architecture that

facilitates excellent transferability and code readability. M-Chem is composed of independent modules (libraries) and a developer can decide which libraries they want to include in their project, which makes the code cleaner and reduces the compilation and linking time. From the computational software engineering perspective, the M-Chem project follows software best practices including extensive unit testing and full documentation. M-Chem has been optimised for excellent on-node parallel performance on modern hardware architectures using OpenMP, even on desktop machines, and thus provides a solid basis for extending the paradigm to efficient Graphical Processing Units (GPUs) and multi-node parallelisation with the Message Passing Interface (MPI) in the future. As we show, M-Chem strives for both accuracy and greater tractability, utilising novel methodology alongside code optimisation that provides faster turnaround and greater interoperability capabilities compared to other software packages. This manuscript serves the primary purpose as an open invitation to software developers who will receive M-Chem source code in exchange for software contributions that builds up capabilities that further benefit developers and users across multiple scientific domains.

2. Biomolecular, chemical, and materials force fields in M-Chem

The vast majority of all-atom MD simulation studies use classical pairwise fixed-charge electrostatics because the MD engine is so well-tuned and optimised for such force fields. M-Chem fully supports such partial charge models for liquids such as water or chain molecules such as proteins. Even so, the failures of pairwise additivity have been made clear in a number of simulation studies [23–26]. Therefore classical potential energy surfaces are evolving towards more intricate models based on the following energy (and force) contributions

$$E_{\text{int}} = E_{\text{valence}} + E_{\text{elec}} + E_{\text{Pauli}} + E_{\text{disp}} + E_{\text{pol}} \quad (1)$$

where E_{valence} corresponds to geometric terms that are designed to capture chemical connectivity but not reactivity, and the E_{elec} , E_{Pauli} , E_{disp} , and E_{pol} terms correspond to the non-bonded permanent electrostatics, Pauli repulsion, dispersion, and polarisation energies, respectively. These terms are the accepted areas for accuracy improvements of advanced force fields [9,26]. Furthermore, most many-body force fields have neglected certain interactions such as charge penetration and charge transfer, and thus they rely on, but do not always demonstrate, how cancellation of errors occurs among the remaining molecular interactions accounted for such as exchange repulsion, electrostatics, and polarisation. Our

recent (many-body) MB-UCB force field for water [10] and ions [27] explicitly accounts for the decomposed molecular interactions commensurate with the variational absolutely-localised molecular orbitals (ALMO) analysis [28,29], including charge penetration and transfer, and makes force field design choices that reduce the computational expense while remaining accurate. We have completed a bottom to top software creation of a MD engine in C++ in M-Chem that encompasses both simple partial charge and advanced non-reactive force fields with point multipole electrostatics and many-body polarisation as developed in force fields such as AMOEBA [9,30,31] and AMOEBA+ [32], EFP [33], GEM [34,35], HIPPO [36], SIBFA [37], MB-Pol [38,39] and the MB-UCB force-field [10,27,40].

Another addition to the model suite in M-Chem is the reactive force-field, ReaxFF [11,12]. We have implemented the ReaxFF model based on the 2008 description of the force field [41], and it is thoroughly validated against the implementation of the same force field in LAMMPS [42]. We also verified the integrity of the implementation through extensive testing of all gradient terms on a wide variety of test systems. Originally, ReaxFF utilised the electronegativity equalisation method (EEM) method [43–45] which manifests charge rearrangements by adopting atomic electronegativity and atomic hardness as fitting parameters to DFT derived Mulliken charges [44,45]. The 2008 ReaxFF model in M-Chem supports the EEM extension developed by Rappé and Goddard by replacing the standard Coulomb potential with a shielded electrostatic term, and using experimental atomic ionisation potentials, electron affinities, and atomic radii as the input data for optimising the charge rearrangements in response to nuclear displacements [46]. We have also developed a ReaxFF model that replaces EEM with the novel Coarse-Grained Electron Model (C-GeM) [13,47]. The C-GeM model represents atoms in terms of a core-shell model, with nuclei being positive cores surrounded by separable electronic shells, with both described by Gaussian charge distributions. This greatly improves the accuracy of the ReaxFF force field without any significant increase in computational time. At this point in time, the ReaxFF energy and gradient terms are all complete, but the code is being optimised, so we have not extensively timed it against other implementations. Additionally, we wish to incorporate recently introduced corrections [48] which smooth slight discontinuities in ReaxFF, thereby greatly improving energy conservation in the microcanonical ensemble.

Finally, M-Chem aims to be an open platform development for coarse-grained models. To illustrate, we have implemented the monoatomic-water (mW) model developed by Molinero and co-workers, which is a single

site water model that exhibits both cooperativity and directional hydrogen-bonding of water in various phases [14,49–51] based on the Stillinger Weber Potential for tetrahedral systems [52]. The mW model is found to be in good agreement with experimental data on water energetics, density, structure, and phase transitions [14], and in the limit of high cooperativity phase separation of liquid water itself was observed in the supercooled state. Recently the mW water model was combined with a CG protein model to study antifreeze proteins [53,54]. In this case the coarse-grained model is comprised of an attractive 2-body potential with a $(1/r^4)$ distance dependence screened with an exponential, and a repulsive 3-body potential with a cosine squared angular dependence that penalises non-tetrahedral configurations, and a product of two scaled exponentials with $\exp(1/r^2)$ distance dependence. This approach of coarse-graining to effective potentials has been shown to agree well with all-atom potentials for force matching [55–63].

3. Performance of non-reactive force fields on multicore nodes

Software implementations of models such as AMOEBA and MB-UCB, when insufficiently optimised, limit their greater adoption by consumers of molecular simulation software [7]. This is because the fastest MD codes for fixed charge models such as CHARMM [64], NAMD [3] and GROMACS [4] efficiently exploit heterogeneous computing, and hence define the aspirational goals for M-Chem to improve scaling performance for advanced many-body potentials and forces. Shared-memory CPUs are an obvious target for a new software platform given widely available computing resources ranging from mid-range compute clusters, cloud computing, and even supercomputer centers. This is recognised by the NAMD [3] and GROMACS [4] teams, whose software packages are leaders in performance in the 10–200-core regime as demonstrated in a recent comparison [64].

However, both packages lack the optimisation for advanced force fields that utilise multipolar electrostatics and point dipole polarisation, core-shell models such as C-GeM, and overall advanced force fields often scale poorly on multicore architectures. This limitation is significant since recent hardware offerings that arrived in late 2019, the 2nd generation AMD EPYC processors, can now support up to 128 CPU cores per node. With two-way simultaneous multithreading, with each core doubling the current peak rate to 33.6 GFLOPS, the tripling of the core count of previous generation CPU configurations is a great opportunity for modernised parallel software.

Parallel efficiency of codes is most commonly measured in terms of weak scaling. But as recently stated by Abraham and co-workers for CPU optimisation of European-based GROMACS: ‘When studying a protein system with 30,000 atoms, it is not relevant that a virus comprising 10 million atoms would scale better. Therefore, weak scaling performance is typically not of primary concern’ [4]. We therefore consider not just weak but strong scaling performance as well for the two most time-consuming parts of evaluation of an MD step in the many-body models: the fixed-radius nearest neighbour search and force field evaluation whose kernel is the non-bonded forces. In this section we compare parallel scaling benchmarks for the MD engine of M-Chem utilising OpenMP on a single node and compared to the high-performance TINKER-HP software package developed by the European group of Piquemal and co-workers [65], and written in Fortran with MPI parallelisation.

OpenMP parallelisation of fixed-radius nearest neighbour evaluations. Due to the computational complexity of the nonbonded terms in force fields, neighbour lists are essential to attain them with linear cost scaling by only evaluating a subset of atom pairs located within a defined cutoff radius. In particular, the real-space Ewald and van der Waals pair list need to be evaluated with a reasonable frequency to account for atoms’ movement during updates through integration of the equations of motion. When the pair list is evaluated for the whole system, the computational complexity becomes significant and therefore well scaling algorithms become critical. Additionally, due to the nature of the AMOEBA force field, positions of atomic centers for certain atoms evaluated with the van der Waals interaction are scaled with respect to the atomic centers, which effectively means that two pair lists need to be evaluated.

Designing an efficient parallel fixed-radius nearest neighbour (fRNN) algorithm with $O(n)$ memory cost scaling and $O(n \log n)$ computational scaling is difficult, however we are able to lean on recent developments for similar algorithms for image and volumetric data processing that have not been used in molecular simulation software previously [66]. The execution time of a single neighbour list evaluation for a system of over 2 million atoms is shown in Figure 1, comparing M-Chem’s octree fRNN algorithm with the linked cells fRNN method available in Tinker-HP. The execution time demonstrates that M-Chem’s implementation is approximately 4–5 times faster than Tinker-HP’s for serial execution on a single core, and outperforming TINKER-HP albeit with increasingly smaller gaps as core count increases. Even so, M-Chem yields a factor of 2 and 1.4 improvement over Tinker-HP for 64 and 128 cores, respectively (numerical values are provided in Table S1).

OpenMP parallelisation of force field evaluations per MD step. Next we consider the parallelisation of the MD step, equivalent to the force evaluations based on Equation (1). Provided that the topology of the molecules are known, parallel algorithms that run in linear time are straightforward to implement for the valence terms. Instead the rate-limiting step of Equation (1) are the non-bonded interactions, especially the long-range electrostatic terms that require computing multipole and induced dipole moments on a grid and performing fast Fourier transforms (FFTs) to complete the Ewald model [15]. While FFTs are also relatively straightforward with existing high-performance software libraries [67,68], the electrostatic multipole moments and inducible dipoles present several difficulties that requires a carefully designed algorithm. Our approach is similar to other codes by using spatial localisation to form batches of multipoles, and the grid assigned to processors followed by parallel accumulation on the final grid [4,69–71].

The M-Chem strong scaling parallel performance of force-field evaluation using Equation (1) is presented in Table 1 in which we compare the timings of 20 iterations of energy and gradient evaluations against Tinker-HP for methanol (978,000 atoms) and water (2,203,614 atoms). The liquid methanol system invokes two sets of induced dipoles (treated with p- and d-scaling to separately evaluate intra- from intermolecular polarisation [30]), and gives us a sense of what performance looks like for the proteinaceous component of a protein-solvent system. We have found it worthwhile to optimise the water force field separately, which only invokes p-scaling (i.e. no intramolecular polarisation is operative), as it is almost always the relevant solvent in simulation studies. M-Chem utilises the fact that systems where all molecules belong to a single polarisation group can be evaluated in a significantly simplified and more efficient manner.

It is important to note in this comparison that the SCF solution to polarisation is evaluated differently in the two codes. M-Chem is using conjugate gradients SCF (CG-SCF) [19] with the electric field from fixed multipoles as an initial guess, while Tinker-HP is evaluated with a divide-and-conquer Jacobi iterations with direct inversion in the iterative subspace extrapolation (DC/JI-DIIS) [72] which is more efficient than the CG-SCF method. Even so, and regardless of the number of cores used, M-Chem performs better. But to consider a more direct comparison between the codes, we consider the evaluation of the AMOEBA force field without mutual polarisation, but just direct polarisation as per the iAMOEBA model [73], such that the rate-limiting step is the 2-body multipolar permanent electrostatics. Table 1 shows that the multipolar electrostatics and

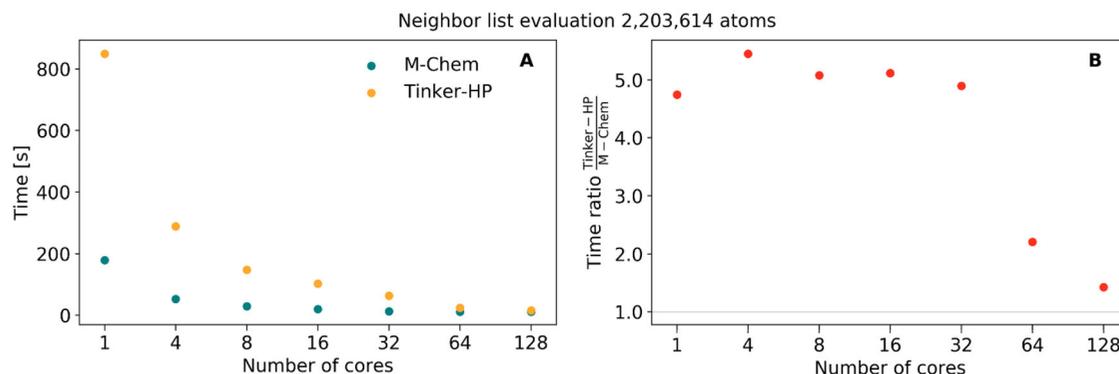


Figure 1. Strong scaling performance of neighbour list evaluation comparing M-Chem and Tinker-HP. A: Total execution time of a single fixed-radius neighbour list evaluation per number of cores on a large water system. B: Tinker-HP/M-Chem execution time ratio.

Table 1. Strong scaling execution time [s] of 20 iterations of the AMOEBA force field energy and gradient evaluation across different systems and algorithms.

978,000 atoms methanol box	N _{CPUS}							
	1	4	8	16	32	64	128	
M-Chem iEL/O-SCF	2330.0	643.0	346.0	200.0	124.0	89.0	94.0	
M-Chem CG-SCF	3054.0	834.0	447.0	259.0	150.0	110.0	108.0	
Tinker-HP DC/JI DIIS	3478.0	1010.0	545.0	288.0	180.0	123.0	120.0	
M-Chem Direct	2268.0	620.0	324.0	186.0	105.0	74.0	66.0	
Tinker-HP Direct	2713.0	794.0	422.0	221.0	136.0	92.0	90.0	
61,000 atoms methanol box	N _{CPUS}							
	1	4	8	16	32	64	128	
	M-Chem iEL/O-SCF	141.0	41.0	22.0	13.0	8.0	7.0	8.0
	M-Chem CG-SCF	181.0	52.0	28.0	17.0	10.0	8.0	11.0
	Tinker-HP DC/JI DIIS	183.0	49.0	30.0	15.0	10.0	7.0	–
	M-Chem Direct	141.0	38.0	20.0	11.0	7.0	5.0	6.0
	Tinker-HP Direct	144.0	40.0	22.0	12.0	7.0	7.0	–
2,203,614 atoms water box	N _{CPUS}							
	1	4	8	16	32	64	128	
	M-Chem iEL/O-SCF	5605.0	1529.0	791.0	438.0	250.0	167.0	162.0
	M-Chem CG-SCF	6753.0	1855.0	972.0	545.0	301.0	213.0	215.0
	Tinker-HP DC/JI DIIS	11008.0	3610.0	1692.0	713.0	522.0	369.0	388.0
	M-Chem Direct	5095.0	1401.0	723.0	398.0	234.0	152.0	153.0
	Tinker-HP Direct	8760.0	2899.0	1488.0	698.0	462.0	289.0	256.0
96,000 atoms water box	N _{CPUS}							
	1	4	8	16	32	64	128	
	M-Chem iEL/O-SCF	219.0	61.0	32.0	17.0	10.0	7.0	7.0
	M-Chem CG-SCF	254.0	70.0	37.0	21.0	13.0	10.0	11.0
	Tinker-HP DC/JI DIIS	282.0	77.0	43.0	22.5	14.0	11.0	–
	M-Chem Direct	206.0	57.0	29.0	16.0	9.0	7.0	7.0
	Tinker-HP Direct	225.0	61.0	34.0	18.0	11.0	7.0	–

Notes: The cutoffs for the van der Waals and real-space Ewald terms were set to 12.0 and 7.0 Å respectively. The induced dipole convergence was set to 10^{-5} D.

corresponding evaluation of electric fields is $\sim 30\%$ faster for the methanol system and $\sim 90\text{-}100\%$ faster for the large water box.

Weak scaling performance for force field evaluations. The comparison of weak scaling between M-Chem and Tinker-HP is presented in Figure 2 and Table S2, where the size of the system per CPU is approximately 15,000 atoms for methanol and 17,000 atoms for water. In the top panel of Figure 2 we first consider the evaluation of AMOEBA with direct polarisation (direct SCF in Table S2), such that the rate-limiting step is again the 2-body multipolar permanent electrostatics. It is evident that the parallel efficiency for M-Chem is superior to

TINKER-HP with performance gains increasing with increasing core count, with the best performance gains of 60% (methanol) and 90% (water) using 128 cores. For the timings with polarisation in the bottom panel of Figure 2, M-Chem outperforms Tinker-HP for all core counts, with the biggest gap of 33 seconds on 128 cores for the $\sim 1\text{M}$ atom methanol box, and saving 193 seconds and approximately 90% of the time used by TINKER-HP for the 2 million atom water system (numerical values are given in Table S2).

Often the overall performance of an MD code is measured in terms of number of ns of simulation executed per day. The standard system used for such benchmarking

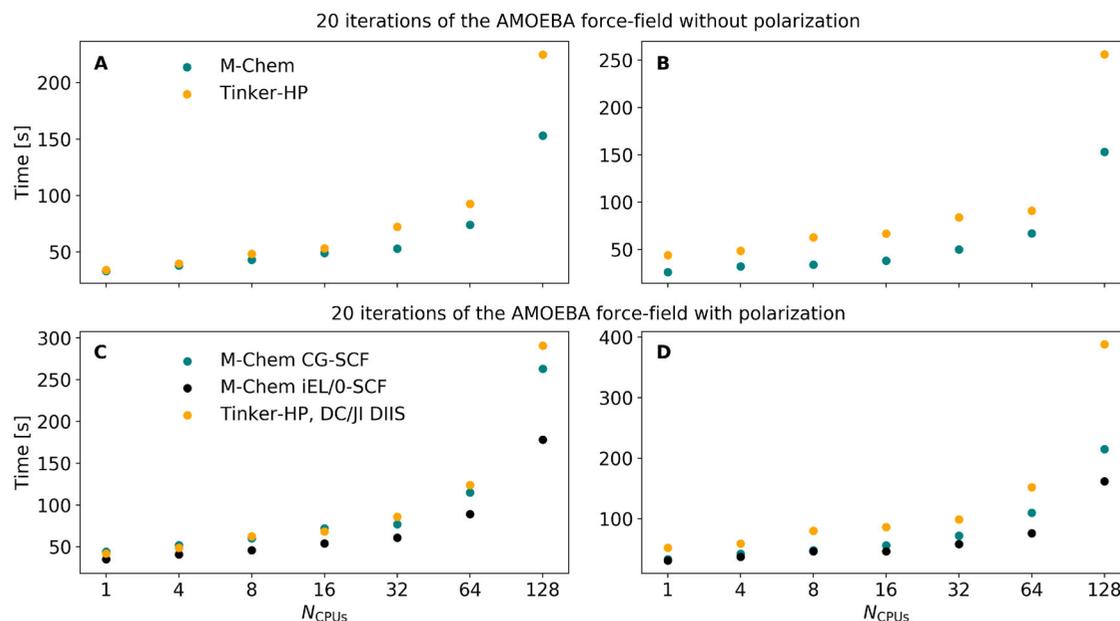


Figure 2. Weak scaling performance of AMOEBA force field components comparing M-Chem and Tinker-HP. Total execution time per number of cores for 20 iterations of force-field evaluations for a (A) methanol box (approx. 15,000 atoms per CPU) and (B) water box (approx. 17,000 atoms per CPU), both with only direct polarisation [73], and largely probing timings for multipolar electrostatics. Total execution time per number of cores for 20 iterations of the complete AMOEBA force-field (including polarisation) for a (C) methanol box (approx. 15,000 atoms per CPU) and (D) water box (approx. 17,000 atoms per CPU). The cutoffs for the van der Waals and real-space Ewald terms were set to 12.0 and 7.0 Å respectively. The induced dipole convergence was set to 10^{-5} D. This figure also compares different SCF algorithms: CG-SCF and iEL/0-SCF for M-Chem and DC/JI DIIS for Tinker-HP.

is the prototypic dihydrofolate reductase (DHFR) solvated in explicit water from the joint AMBER/CHARMM benchmark. Since a full protein front-end features have not yet been fully implemented in M-Chem, we performed the benchmark on a methanol box (23,556 atoms), and the best performance results using OpenMP were obtained on 64 cores, yielding the same performance of the MPI version of Tinker-HP using 64 cores of ~ 1.2 ns/day, and a factor of 6X compared to the TINKER reference code. Even so, the next step for M-Chem is to pursue hybrid OpenMP/MPI to improve on node performance and to advance to cross-node MPI parallelisation that will benefit large systems.

One of the goals of M-Chem is to also support a diverse set of models and algorithms to expand the timescales and lengthscales of MD simulation and to provide necessary statistical mechanics capabilities. The M-Chem MD engine is thus already imbued with standard velocity scaling including Nose Hoover thermostats and barostats [16], periodic boundary conditions using cubic cells, and conjugate gradients for minimisation and appropriate pre-conditioners for conjugate gradients used in the self-consistent field step for solving sets of linear equations (such as for polarisation).

The choice of the algorithm to treat polarisation or charge equilibration is of paramount importance to overall efficiency of force field evaluation for non-reactive

and reactive force fields. Albaugh and co-workers have introduced a number of extended Lagrangian approaches that reduces SCF cycles by half (iEL/SCF) [20] or eliminates the self-consistent field step altogether (iEL/0-SCF) [21] for non-reactive force fields such as AMOEBA or MB-UCB. Most recently we combined iEL/0-SCF with a stochastic integration scheme that allows for a longer time step using a multi-time stepping algorithm, SIN(R), developed by Tuckerman and co-workers [74,75]. Depending on system and desired accuracy, the iEL/0-SCF and SIN(R) combination yields lower bound computational speed-ups of 6–8 relative to a standard Verlet integration step using a standard SCF solver [75]. This benefits not only polarisation, but the simple induction-like model for charge transfer of the MB-UCB model. The results in Figure 2 show that, for the cases of both methanol and water, M-Chem’s default extended Lagrangian integrator decreases execution time further (black dots) than the SCF solver. Furthermore, the EEM or charge equilibration method (CEM or qEq) electrostatics of the ReaxFF model in M-Chem can be solved using either a preconditioned conjugate gradient method, and we are in the process of implementing the iEL/SCF [76] and SC-XLMD [22] extended Lagrangian approaches that should increase simulation timescales dramatically for ReaxFF simulations.

4. Interoperability in M-Chem: PyTorch and Q-Chem

Machine learning interfaces. The rise of machine-learned models in biology, chemistry, and materials science [77,78] is tackling everything from materials property prediction to learning the energies and forces of complex system with good to excellent ab initio quality. This exciting and burgeoning area has resulted in many different neural network-based force field models [79–82], and hence we have defined a generic interface for ML FF models built with the PyTorch framework [83]. This is made possible by the ability of PyTorch to compile its own models into a byte-code representation which can then be loaded and evaluated by libtorch, which is a full-featured front-end API for PyTorch written in C++. This means that all we need to do is link against the libtorch library and we can evaluate ML models on CPUs, GPUs, or even multiple GPUs with minimal programming effort. Our aim is to allow users to write code within the PyTorch ecosystem and use their models in M-Chem without having to write any C++ code.

We have done this successfully with NewtonNet [82], a message passing equivariant NN model, demonstrating the viability of this approach. In order to ensure the interface has been implemented properly, we compared timings of evaluating the force field 2000 times, including calculation of both energies and gradients using the PyTorch version and the same model exported to C++ and run in MChem. To do this, we used a version of NewtonNet trained on a hydrogen-combustion dataset [84] and evaluated the model on many configurations of $H_2O_2(OH\ddot{O}H)$. We find that evaluating energies and gradients with the model on the CPU takes an average of 6.4 ms using libtorch as integrated into M-Chem as opposed to the 7.7 ms using PyTorch. Similarly, on the GeForce GTX 1650 mobile GPU, which is not intended for fast double precision arithmetic as required here, evaluation with libtorch/M-Chem takes 8.7 ms vs 10.2 ms using PyTorch. Overall, by simply evaluating a neural network through the libtorch C++ API, we can evaluate the model $\sim 20\%$ faster using the PyTorch Python API. While these relative timings will depend on the structure of the model being used and the system being evaluated, we emphasise that it is reasonable to expect moving your model to MChem is likely to slightly speed up a model rather than slow it down.

Furthermore, this interface opens many interesting possibilities. Indeed, one could implement an entire force field of any kind using PyTorch (in python), serialise the model, load it into M-Chem, and use this force field in MD simulations without having to write any C++. The

refinement of a neural network force field could also be accomplished on the fly, by backpropagating a NN model and updating the weights on the C++ side. One can hence envision advanced QM/ML simulations in which the QM method can be dynamically swapped in or out when a NN has low confidence in its prediction, allowing the ML model to be updated based on the results of the QM calculation without stopping the simulation. This complex type of simulation is still a work in progress but is the type of ML interface capabilities that MChem aims to make routine, thereby creating an easy interface to create reference or prototype code, and would be an excellent opportunity for a first contribution. In the future, we may explore similar interfaces with TensorFlow, another popular ML library.

QM/MM mechanical and electrostatic embedding. Biologically or technologically relevant systems are typically very large and thus computationally too demanding to be treated at any level of quantum mechanical method that accounts for the explicit role of the electrons and have high predictive power. At the same time, commonly used molecular mechanical force fields are not sufficiently flexible to model processes that involve electronic rearrangement, e.g. bond formation/breaking and electronic excitations. Fortunately, many chemical processes of interest occur in spatially localised regions and only a small region must be treated at the quantum mechanical level. Therefore, one can overcome the limitations of both methods by constructing an approximate reduced-dimensionality Hamiltonian. In such a hybrid scheme, the system is divided into a small region of interest, which undergoes electronic rearrangement and is treated quantum mechanically, and the remaining part, the environment, which is treated classically; the so-called quantum mechanics/molecular mechanics framework. Utilizing such model Hamiltonians not only makes the calculations computationally feasible, but also leads to a clear distinction between essential and non-essential parts of the system. As a practical matter, for a successful QM/MM simulation, a molecular dynamics software (here M-Chem) must be combined with an electronic structure software (here Q-Chem) that requires coordination between many different pieces of code to enable various flavour of QM/MM schemes. There are three QM/MM schemes, which describe the coupling between the QM and MM regions to different extents; i) mechanical embedding, ii) electrostatic embedding, and iii) polarisable embedding. Currently our QM/MM interface supports mechanical embedding. Work in the direction to more elaborate embedding such as multi-polar electrostatic embedding and polarisable embedding are in progress in our team.

5. Simulation utilities in M-Chem: ESPs, FF parameterization, and file formats

It is desirable for a single biomolecular simulation software package to support diverse models, methods, and analysis tools beyond the MD engine itself. An important utility relevant for all kinds of chemical applications is the ability to accurately create an electrostatic potential (ESP) on a molecular surface. ESPs are used to help understand sites of biomolecular association and docking, drug binding, interfacial surface potentials, and visualisation of noncovalent interactions such as hydrogen bonding, halogen bonding, and pi-pi interactions. Another example is that the point of transferable force field models is that protocols must be developed that allow a systematic search for a consistent set of force field parameters for new molecule chemistries. These software libraries associated with the M-Chem code are described here.

Electrostatic potentials using C-GeM. Although C-GeM is already part of the ReaxFF force field implementation in M-Chem, it can also be run standalone for the fast generation of electrostatic potentials. We have tested C-GeM against 600 molecules each with 10 separate configurations that are being used for benchmarking empirical approaches for ESP rendering of drug molecules containing 30–65 atoms evaluated at the B3LYP/6-311G** level of theory. C-GeM performed quite accurately across the data set with mean absolute error (MAE) of 2.8 kcal/mol, compared to 7.2 kcal/mol MAE for the default EEM method used in the docking software AutoDock, and slightly better than new machine learning approaches generated by the Forli group. Compared to the original C-GeM model [17] which works in gas phase, in the M-Chem implementation we added dielectric settings to allow for electrostatics predictions in solvent. Figure 3 shows the C-GeM electrostatic potential surface for crambin in solvent(a) and in gas phase(b).

It is our hope that, given the interface with Q-Chem, we can support alternative approaches such as the rendering of fitted charges (ESPC) in which the partial charges on atoms are derived from a partitioning scheme of an expensive electronic wave functions such as Mulliken [85], Hirshfeld-I [86], DDEC/c3 [87], or natural population analysis [88]. ESPC models are the main methods used in all the major biomolecular force fields including AMBER [1] and CHARMM [89], and thus should also be supported here in M-Chem. Electrostatic potential surfaces rendered from Poisson Boltzmann models are also possible, and the Head-Gordon lab has developed PB-AM and PB-SAM in APBS [90], but will be brought in-house in future versions of M-Chem.

QM-based force fields with Q-Force. Fully QM-based force fields provide an alternative to atom-types-based

force fields with molecule-specific parameters that result in a good match between QM and MM potential energy surfaces. The Q-Force toolkit [18], which derives QM-based force fields through an automated pipeline, is an open-source software that has been interfaced with M-Chem and Q-Chem. Q-Force automatically derives standard MD models for any (novel) molecule based on QM Hessians, the electrostatic potential, torsional scans and atomic bond orders used to parametrise these potentials. It has been shown that such an approach can drastically improve the match between QM and MD potential energy surfaces and consequently increase the predictive power of standard MD simulations without additional computational cost after the initial parameterisation. Additionally, the implemented molecular fragmentation procedure allows parameterising large molecules with significantly reduced computational cost, which allowed Q-Force to be easily applied to a variety of technologically relevant applications. Extending Q-Force to advanced force fields is currently an active research line in the Head-Gordon lab and will be developed in M-Chem.

File Formats. For saving trajectories, we use a storage mechanism based on the Hierarchical Data Format version 5 (HDF5) standard. It is similar in spirit to the MDTraj HDF5 layout, in particular with storing the topology alongside time-ordered information, such as coordinates and optional velocities. However, due to the data requirements of non-traditional force fields such as ReaxFF and machine-learned potentials, it is not a direct reimplementation of the MDTraj format. We instead define the data layout in an extensible JSON Schema-like format called qarchive for which compile-time invariants on the layout are enforced, with bindings to other languages forthcoming.

6. Discussion and conclusions

Typically, and over time, software components and algorithms become tightly coupled into a monolithic code that is difficult to port to a new platform or do not offer ease of entry for novel development work. It is now understood that much better portability can be attained at a very modest performance trade-off, with attention to modularity. Furthermore, present-day computers require careful algorithm design and performance tuning to harness the power of multiple levels of parallelism: data-parallel instructions at the lowest level, multitasking at the shared-memory multi-core level, and across-node distributed-memory parallel code execution.

In summary, we have presented a current snapshot of the capabilities of the M-Chem software package, a modular and feature diverse molecular simulation code

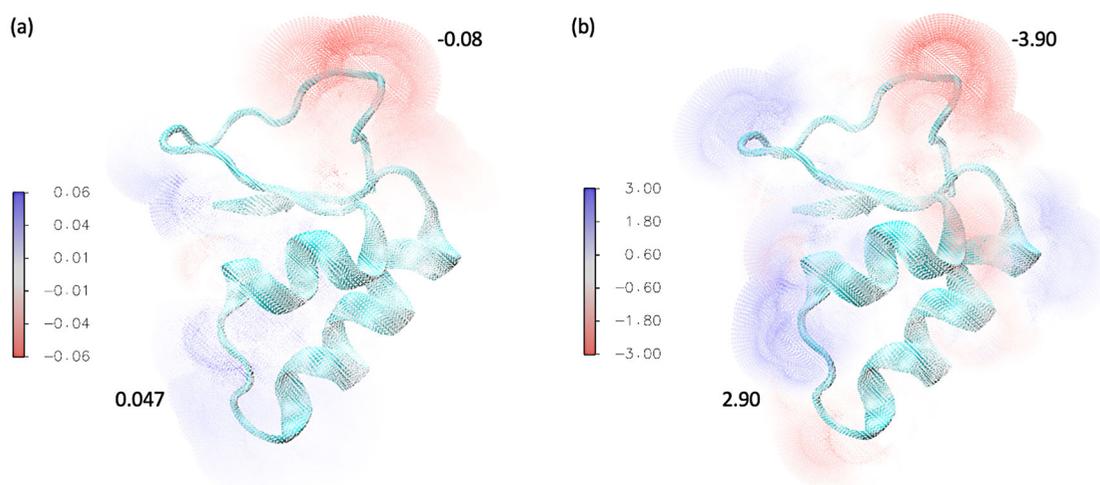


Figure 3. Electrostatic potential surface predicted with C-GeM on crambin(1CRN) (a) with dielectric constant for solution $\epsilon_s = 78.54$ and (b) in gas phase $\epsilon_s = 1$.

written from the ground up in C++. M-Chem has already optimised serial and on-node performance on multicore CPUs, but the resulting software design is extensible to cross-node and GPU parallelisation which defines the next phase of code optimisation. The M-Chem infrastructure has also led to software hooks and APIs that permit integration of diverse models and methods such as a range of force fields, ML, and QM/MM capabilities. Our hope is that the ease with which we can change the many-body electrostatic model associated with any of the above force fields, or to interface or embed alternative ML and physics-based models, demonstrates our goal of making the physics modular where possible without sacrificing accuracy or speed.

Of course M-Chem as a user-friendly code is far from complete in many ways including front-end user interfaces to set up interfacial systems, large macromolecules, and extremely large all-atom simulations (~ 500 million atoms) or back-end analysis tools that address the needs of diverse communities. M-Chem is missing critical algorithms such as free energy calculations, Poisson Boltzmann solvers, Monte-Carlo methods, other coarse-grained models, or advanced sampling methods. Even so, enough code capability is present and useable such that M-Chem is a software platform ready for bringing in a completely new set of developers of academic researchers and setting the stage for better informed partnerships with academic users and industrial consumers of M-Chem.

Acknowledgments

As Prof. Peter Gill is one of the founders of Q-Chem, which was also built from scratch over 25 years ago, we are happy to

honour him with this debut of the M-Chem software package (of which we hope Peter becomes a developer!). We thank Farnaz Heider-Zadeh for the dihedral angle and C-GeM reference codes

Data availability

The datasets generated during and/or analyzed during the current study are available from the corresponding author on reasonable request.

Disclosure statement

E. E. is a part-owner of Q-Chem, Inc.

Funding

We also thank the National Science Foundation under grant CHE-1955643. M-Chem software development was also supported under a SBIR grant from the National Institutes of Health [2R44GM128480-02].

Code availability

The source code is available to all developers.

ORCID

Jagna Witek <http://orcid.org/0000-0003-1476-6130>

Joseph P. Heindel <http://orcid.org/0000-0002-9748-1730>

Xingyi Guan <http://orcid.org/0000-0001-7623-4012>

Selim Sami <http://orcid.org/0000-0002-4484-0322>

M. F. S. J. Menger <http://orcid.org/0000-0003-1442-9601>

D. Vale Cofer-Shabica <http://orcid.org/0000-0003-4623-6521>

Eric Berquist <http://orcid.org/0000-0001-8186-9522>

Shirin Faraji <http://orcid.org/0000-0002-6421-4599>

Teresa Head-Gordon <http://orcid.org/0000-0003-0025-8987>

References

- [1] R. Salomon-Ferrer, D.A. Case and R.C. Walker, *Wiley Interdiscip. Rev. Comput. Mol. Sci.* **3** (2), 198–210 (2013). doi:10.1002/wcms.1121.
- [2] B.R. Brooks, C.L. Brooks III, A.D. Mackerell Jr., L. Nilsson, R.J. Petrella, B. Roux, Y. Won, G. Archontis, C. Bartels, S. Boresch, A. Caffisch, L. Caves, Q. Cui, A.R. Dinner, M. Feig, S. Fischer, J. Gao, M. Hodoscek, W. Im, K. Kuczera, T. Lazaridis, J. Ma, V. Ovchinnikov, E. Paci, R.W. Pastor, C.B. Post, J.Z. Pu, M. Schaefer, B. Tidor, R.M. Venable, H.L. Woodcock, X. Wu, W. Yang, D.M. York and M. Karplus, *J. Comput. Chem.* **30** (10), 1545–1614 (2009). doi:10.1002/jcc.21287.
- [3] J.C. Phillips, R. Braun, W. Wang, J. Gumbart, E. Tajkhorshid, E. Villa, C. Chipot, R.D. Skeel, L. Kalvo and K. Schulten, *J. Comput. Chem.* **26** (16), 1781–1802 (2005). doi:10.1002/jcc.20289.
- [4] M.J. Abraham, T. Murtola, R. Schulz, S. Páll, J.C. Smith, B. Hess and E. Lindahl, *SoftwareX* **1–2**, 19–25 (2015). doi:10.1016/j.softx.2015.06.001.
- [5] S. Plimpton, *J. Comput. Phys.* **117** (1), 1–19 (1995). <http://lammps.sandia.gov>.
- [6] E. Epifanovsky, A.T.B. Gilbert, X. Feng, J. Lee, Y. Mao, N. Mardirossian, P. Pokhilko, A.F. White, M.P. Coons, A.L. Dempwolff, Z. Gan, D. Hait, P.R. Horn, L.D. Jacobson, I. Kaliman, J. Kusmann, A.W. Lange, K.U. Lao, D.S. Levine, J. Liu, S.C. McKenzie, A.F. Morrison, K.D. Nanda, F. Plasser, D.R. Rehn, M.L. Vidal, Z.-Q. You, Y. Zhu, B. Alam, B.J. Albrecht, A. Aldossary, E. Alguire, J.H. Andersen, V. Athavale, D. Barton, K. Begam, A. Behn, N. Bellonzi, Y.A. Bernard, E.J. Berquist, H.G.A. Burton, A. Carreras, K. Carter-Fenk, R. Chakraborty, A.D. Chien, K.D. Closser, V. Cofer-Shabica, S. Dasgupta, M. de Wergifosse, J. Deng, M. Diederhofen, H. Do, S. Ehlert, P.-T. Fang, S. Fatehi, Q. Feng, T. Friedhoff, J. Gayvert, Q. Ge, G. Gidofalvi, M. Goldey, J. Gomes, C.E. González-Espinoza, S. Gulania, A.O. Gunina, M.W.D. Hanson-Heine, P.H.P. Harbach, A. Hauser, M.F. Herbst, M.H. Vera, M. Hodecker, Z.C. Holden, S. Houch, X. Huang, K. Hui, B.C. Huynh, M. Ivanov, Á. Jász, H. Ji, H. Jiang, B. Kaduk, S. Kähler, K. Khistyayev, J. Kim, G. Kis, P. Klunzinger, Z. Koczor-Benda, J.H. Koh, D. Kosenkov, L. Koulias, T. Kowalczyk, C.M. Krauter, K. Kue, A. Kunitsa, T. Kus, I. Ladjanszki, A. Landau, K.V. Lawler, D. Lefrancois, S. Lehtola, R.R. Li, Y.-P. Li, J. Liang, M. Liebenthal, H.-H. Lin, Y.-S. Lin, F. Liu, K.-Y. Liu, M. Loipersberger, A. Luenser, A. Manjanath, P. Manohar, E. Mansoor, S.F. Manzer, S.-P. Mao, A.V. Marenich, T. Markovich, S. Mason, S.A. Maurer, P.F. McLaughlin, M.F.S.J. Menger, J.-M. Mewes, S.A. Mewes, P. Morgante, J.W. Mullinax, K.J. Oosterbaan, G. Paran, A.C. Paul, S.K. Paul, F. Pavošević, Z. Pei, S. Prager, E.I. Proynov, Á. Rák, E. Ramos-Cordoba, B. Rana, A.E. Rask, A. Rettig, R.M. Richard, F. Rob, E. Rossomme, T. Scheele, M. Scheurer, M. Schneider, N. Sergueev, S.M. Sharada, W. Skomorowski, D.W. Small, C.J. Stein, Y.-C. Su, E.J. Sundstrom, Z. Tao, J. Thirman, G.J. Tornai, T. Tsuchimochi, N.M. Tubman, S. Prasad Veccham, O. Vydrov, J. Wenzel, J. Witte, A. Yamada, K. Yao, S. Yeganeh, S.R. Yost, A. Zech, I.Y. Zhang, X. Zhang, Y. Zhang, D. Zuev, A. Aspuru-Guzik, A.T. Bell, N.A. Besley, K.B. Bravaya, B.R. Brooks, D. Casanova, J.-D. Chai, S. Coriani, C.J. Cramer, G. Cserey, A. Eugene DePrince, R.A. DiStasio, A. Dreuw, B.D. Dunietz, T.R. Furlani, W.A. Goddard, S. Hammes-Schiffer, T. Head-Gordon, W.J. Hehre, C.-P. Hsu, T.-C. Jagau, Y. Jung, A. Klamt, J. Kong, D.S. Lambrecht, W. Liang, N.J. Mayhall, C. William McCurdy, J.B. Neaton, C. Ochsenfeld, J.A. Parkhill, R. Peverati, V.A. Ras-solov, Y. Shao, L.V. Slipchenko, T. Stauch, R.P. Steele, J.E. Subotnik, A.J.W. Thom, A. Tkatchenko, D.G. Truhlar, T. Van Voorhis, T.A. Wesolowski, K. Birgitta Whaley, H.L. Woodcock, P.M. Zimmerman, S. Faraji, P.M.W. Gill, M. Head-Gordon, J.M. Herbert and A.I. Krylov, *J. Chem. Phys.* **155** (8), 084801 (2021). doi:10.1063/5.0055522.
- [7] A. Krylov, T.L. Windus, T. Barnes, E. Marin-Rimoldi, J.A. Nash, B. Pritchard, D.G.A. Smith, D. Altarawy, P. Saxe, C. Clementi, T.D. Crawford, R.J. Harrison, S. Jha, V.S. Pande and T. Head-Gordon, *J. Chem. Phys.* **149** (18), 180901 (2018). doi:10.1063/1.5052551.
- [8] P.Y. Ren and J.W. Ponder, *J. Comp. Chem.* **23** (16), 1497–1506 (2002). doi:10.1002/jcc.10127
- [9] J.W. Ponder, C. Wu, P. Ren, V.S. Pande, J.D. Chodera, M.J. Schnieders, I. Haque, D.L. Mobley, D.S. Lambrecht, R.A. DiStasio Jr., M. Head-Gordon, G.N.I. Clark, M.E. Johnson and T. Head-Gordon, *J. Phys. Chem. B* **114** (8), 2549–2564 (2010). doi:10.1021/jp910674d
- [10] A.K. Das, L. Urban, I. Leven, M. Loipersberger, A. Aldossary, M. Head-Gordon and T. Head-Gordon, *J. Chem. Theory. Comput.* **15** (9), 5001–5013 (2019). doi:10.1021/acs.jctc.9b00478.
- [11] A.C.T. van Duin, S. Dasgupta, F. Lorant and W.A. Goddard, *J. Phys. Chem. A* **105**, 9396–9409 (2001). doi:10.1021/jp004368u
- [12] I. Leven, H. Hao, S. Tan, X. Guan, K.A. Penrod, D. Akbarian, B. Evangelisti, M.J. Hossain, M.M. Islam, J.P. Koski, S. Moore, H.M. Aktulga, A.C.T. van Duin and T. Head-Gordon, *J. Chem. Theory. Comput.* **17** (6), 3237–3251 (2021). doi:10.1021/acs.jctc.1c00118.
- [13] I. Leven and T. Head-Gordon, *J. Phys. Chem. Lett.* **10** (21), 6820–6826 (2019). doi:10.1021/acs.jpcllett.9b02771
- [14] V. Molinero and E.B. Moore, *J. Phys. Chem. B* **113** (13), 4008–16 (2009). doi:10.1021/jp805227c.
- [15] C. Sagui, L.G. Pedersen and T.A. Darden, *J. Chem. Phys.* **120**, 73–87 (2004). doi:10.1063/1.1630791
- [16] M. Tuckerman, B.J. Berne and G.J. Martyna, *J. Chem. Phys.* **97**, 1990–2001 (1992). doi:10.1063/1.463137
- [17] X. Guan, I. Leven, F. Heidar-Zadeh and T. Head-Gordon, *J. Chem. Inf. Model.* **61**, 4357–4369 (2021). doi:10.1021/acs.jcim.1c00388
- [18] S. Sami, M.F.S.J. Menger, S. Faraji, R. Broer and R.W.A. Havenith, *J. Chem. Theory Comput.* **17** (8), 4946–4960 (2021). doi:10.1021/acs.jctc.1c00195.
- [19] W. Wang and R.D. Skeel, *J. Chem. Phys.* **123**, 164107–164107 (2005). doi:10.1063/1.2056544
- [20] A. Albaugh, O. Demerdash and T. Head-Gordon, *J. Chem. Phys.* **143** (17), 174104 (2015). doi:10.1063/1.4933375.
- [21] A. Albaugh and T. Head-Gordon, *J. Chem. Theory Comput.* **13** (11), 5207–5216 (2017). doi:10.1021/acs.jctc.7b00838.
- [22] S. Tan, I. Leven, D. An, L. Lin and T. Head-Gordon, *J. Chem. Theo. Comp.* **16** (10), 5991–5998 (2020). doi:10.1021/acs.jctc.0c00514.
- [23] R.C. Remsing, M.D. Baer, G.K. Schenter, C.J. Mundy and J.D. Weeks, *J. Phys. Chem. Lett.* **5** (16), 2767–2774 (2014).

- doi:10.1021/jz501067w < Go to ISI > ://WOS:000340807100001.
- [24] M. Vazdar, E. Pluharova, P.E. Mason, R. Vacha and P. Jungwirth, *J. Phys. Chem. Lett.* **3** (15), 2087–2091 (2012). doi:10.1021/jz300805b < Go to ISI > ://WOS:000309691500029.
- [25] P.S. Nerenberg and T. Head-Gordon, *Curr. Opin. Struct. Biol.* **49**, 129–138 (2018). doi:10.1016/j.sbi.2018.02.002.
- [26] O. Demerdash, E.H. Yap and T. Head-Gordon, *Annu. Rev. Phys. Chem.* **65**, 149–174 (2014). doi:10.1146/physchem.2014.65.issue-1
- [27] A.K. Das, M. Liu and T. Head-Gordon, *J. Chem. Theory Comput.* **18** (2), 953–967 (2022). doi:10.1021/acs.jctc.1c00537.
- [28] R.Z. Khaliullin, E.A. Cobar, R.C. Lochan, A.T. Bell and M. Head-Gordon, *J. Phys. Chem. A* **111** (36), 8753–8765 (2007). doi:10.1021/jp073685z.
- [29] P.R. Horn, Y. Mao and M. Head-Gordon, *Phys. Chem. Chem. Phys.* **18**, 23067–23079 (2016). doi:10.1039/C6CP03784D.
- [30] P.Y. Ren and J.W. Ponder, *J. Phys. Chem. B* **107** (24), 5933–5947 (2003). doi:10.1021/jp027815+
- [31] J.W. Ponder and D.A. Case, *Force Fields for Protein Simulations*, 66 vols. (Academic Press, 2003), pp. 27–85.
- [32] C. Liu, J.-P. Piquemal and P. Ren, *J. Chem. Theory Comput.* **15** (7), 4122–4139 (2019). doi:10.1021/acs.jctc.9b00261.
- [33] L.V. Slipchenko and M.S. Gordon, *Mol. Phys.* **107** (8–12, SI), 999–1016 (2009). doi:10.1080/00268970802712449
- [34] G.A. Cisneros, D. Elking, J.P. Piquemal and T.A. Darden, *J. Phys. Chem. A* **111** (47), 12049–12056 (2007). doi:10.1021/jp074817r.
- [35] G.A. Cisneros, K.T. Wikfeldt, L. Ojamae, J.B. Lu, Y. Xu, H. Torabifard, A.P. Bartok, G. Csanyi, V. Molinero and F. Paesani, *Chem. Rev.* **116** (13), 7501–7528 (2016). doi:10.1021/acs.chemrev.5b00644.
- [36] J.A. Rackers, R.R. Silva, Z. Wang and J.W. Ponder, *J. Chem. Theory Comput.* **17** (11), 7056–7084 (2021). doi:10.1021/acs.jctc.1c00628.
- [37] N. Gresh, G.A. Cisneros, T.A. Darden and J.P. Piquemal, *J. Chem. Theory Comput.* **3** (6), 1960–1986 (2007). doi:10.1021/Ct700134r.
- [38] V. Babin, C. Leforestier and F. Paesani, *J. Chem. Theory Comput.* **9** (12), 5395–5403 (2013). doi:10.1021/ct400863t < Go to ISI > ://WOS:000328437500020.
- [39] V. Babin, G.R. Medders and F. Paesani, *J. Phys. Chem. Lett.* **3** (24), 3765–3769 (2012). doi:10.1021/jz3017733 < Go to ISI > ://WOS:000312762900015.
- [40] A.K. Das, O.N. Demerdash and T. Head-Gordon, *J. Chem. Theory Comput.* **14** (12), 6722–6733 (2018). doi:10.1021/acs.jctc.8b00978.
- [41] K. Chenoweth, A.C.T. Van Duin and W.A. Goddard, *J. Phys. Chem. A* **112** (5), 1040–1053 (2008). doi:10.1021/jp709896w
- [42] S.B. Kylasa, H.M. Aktulga and A.Y. Grama, *J. Comput. Phys.* **272**, 343–359 (2014). doi:10.1016/j.jcp.2014.04.035
- [43] J. Gasteiger and M. Marsili, *Tetrahedron. Lett.* **19** (34), 3181–3184 (1978). doi:10.1016/S0040-4039(01)94977-9
- [44] W.J. Mortier, K. Van Genechten and J. Gasteiger, *J. Am. Chem. Soc.* **107** (4), 829–835 (1985). doi:10.1021/ja00290a017.
- [45] W.J. Mortier, S.K. Ghosh and S. Shankar, *J. Am. Chem. Soc.* **108** (15), 4315–4320 (1986). doi:10.1021/ja00275a013.
- [46] A.K. Rappe and W.A. Goddard III, *J. Phys. Chem.* **95** (8), 3358–3363 (1991). doi:10.1021/j100161a070
- [47] I. Leven, H. Hao, A.K. Das and T. Head-Gordon, *J. Phys. Chem. Lett.* **11**, 9240–9247 (Oct, 2020). doi:10.1021/acs.jpcllett.0c02516.
- [48] D. Furman and D.J. Wales, *J. Phys. Chem. Lett.* **10** (22), 7215–7223 (2019). doi:10.1021/acs.jpcllett.9b02810
- [49] F. Romano, E. Sanz and F. Sciortino, *J. Chem. Phys.* **134** (17), 174502 (2011). doi:10.1063/1.3578182.
- [50] E.B. Moore and V. Molinero, *Nature* **479** (7374), 506–8 (2011). doi:10.1038/nature10586.
- [51] F. Smallenburg and F. Sciortino, *Phys. Rev. Lett.* **115** (1), 015701 (2015). doi:10.1103/PhysRevLett.115.015701
- [52] F.H. Stillinger and T.A. Weber, *Phys. Rev. B Condens Matter* **31** (8), 5262–5271 (1985). doi:10.1103/physrevb.31.5262.
- [53] A. Hudait, N. Odendahl, Y. Qiu, F. Paesani and V. Molinero, *J. Am. Chem. Soc.* **140** (14), 4905–4912 (2018). doi:10.1021/jacs.8b01246.
- [54] A. Hudait, Y. Qiu, N. Odendahl and V. Molinero, *J. Am. Chem. Soc.* **141** (19), 7887–7898 (2019). doi:10.1021/jacs.9b02248.
- [55] G.S. Ayton and G.A. Voth, *J. Phys. Chem. B* **113** (13), 4413–24 (2009). doi:10.1021/jp8087868.
- [56] S. Izvekov and G.A. Voth, *J. Chem. Theory Comput.* **5** (12), 3232–44 (2009). doi:10.1021/ct900414p.
- [57] V. Krishna, W.G. Noid and G.A. Voth, *J. Chem. Phys.* **131** (2), 024103 (2009). doi:10.1063/1.3167797.
- [58] Y. Wang, W.G. Noid, P. Liu and G.A. Voth, *Phys. Chem. Chem. Phys.* **11** (12), 2002–15 (2009). doi:10.1039/b819182d.
- [59] L. Lu, S. Izvekov, A. Das, H.C. Andersen and G.A. Voth, *J. Chem. Theory Comput.* **6** (3), 954–65 (2010). doi:10.1021/ct900643r.
- [60] I.F. Thorpe, D.P. Goldenberg and G.A. Voth, *J. Phys. Chem. B* **115** (41), 11911–26 (2011). doi:10.1021/jp204455g.
- [61] A. Davtyan, J.F. Dama, G.A. Voth and H.C. Andersen, *J. Chem. Phys.* **142** (15), 154104 (2015). doi:10.1063/1.4917454.
- [62] J.W. Wagner, T. Dannenhoffer-Lafage, J. Jin and G.A. Voth, *J. Chem. Phys.* **147** (4), 044113 (2017). doi:10.1063/1.4995946.
- [63] J. Jin and G.A. Voth, *J. Chem. Theory Comput.* **14** (4), 2180–2197 (2018). doi:10.1021/acs.jctc.7b01173.
- [64] A.-P. Hynninen and M.F. Crowley, *J. Comput. Chem.* **35** (5), 406–413 (2014). doi:10.1002/jcc.23501.
- [65] L. Lagardère, L.-H. Jolly, F. Lipparini, F. Aviat, B. Stamm, Z.F. Jing, M. Harger, H. Torabifard, G.A. Cisneros, M.J. Schnieders, N. Gresh, Y. Maday, P.Y. Ren, J.W. Ponder and J.-P. Piquemal, *Chem. Sci.* **9**, 956–972 (2018). doi:10.1039/C7SC04531J
- [66] G. Schrack, *CVGIP: Image Understanding* **55**, 221–230 (1992). doi:10.1016/1049-9660(92)90022-U
- [67] M. Frigo and S.-G. Johnson, *Proce. IEEE* **93**, 216–231 (2005). Special issue on “Program Generation, Optimization, and Platform Adaptation.
- [68] M. Frigo and S.G. Johnson, *FFTW tutorial for version 3.3.8* (2018).

- [69] S. Páll and B. Hess, *Comput. Phys. Commun.* **184** (12), 2641–2650 (2013). doi:10.1016/j.cpc.2013.06.003.
- [70] J. Jung, T. Mori, C. Kobayashi, Y. Matsunaga, T. Yoda, M. Feig and Y. Sugita, *WIREs Computational Molecular Science* **5** (4), 310–323 (2015). doi:10.1002/wcms.1220.
- [71] L. Lagardere, L.H. Jolly, F. Lipparini, F. Aviat, B. Stamm, Z.F. Jing, M. Harger, H. Torabifard, G.A. Cisneros, M.J. Schnieders, N. Gresh, Y. Maday, P.Y. Ren, J.W. Ponder and J.P. Piquemal, *Chem. Sci.* **9** (4), 956–972 (2018). doi:10.1039/c7sc04531j.
- [72] D. Nocito and G.J.O. Berana, *J. Chem. Phys.* **146**, 114103 (2017). doi:10.1063/1.4977981
- [73] L.P. Wang, T. Head-Gordon, J.W. Ponder, P. Ren, J.D. Chodera, P.K. Eastman, T.J. Martinez and V.S. Pande, *J. Phys. Chem. B* **117** (34), 9956–9972 (2013). doi:10.1021/jp403802c
- [74] D.T. Margul and M.E. Tuckerman, *J. Chem. Theo. Comput.* **12**, 2170–2180 (2016). doi:10.1021/acs.jctc.6b00188
- [75] A. Albaugh, M.E. Tuckerman and T. Head-Gordon, *J. Chem. Theory. Comput* **15** (4), 2195–2205 (2019). doi:10.1021/acs.jctc.9b00072.
- [76] I. Leven and T. Head-Gordon, *Phys. Chem. Chem. Phys.* **21** (34), 18652–18659 (2019). doi:10.1039/c9cp02979f.
- [77] K.T. Butler, D.W. Davies, H. Cartwright, O. Isayev and A. Walsh, *Nature* **559** (7715), 547–555 (2018). doi:10.1038/s41586-018-0337-2
- [78] M. Ceriotti, C. Clementi and O.A. von Lilienfeld, *Introduction: Machine learning at the atomic scale*, 2021.
- [79] J. Behler and M. Parrinello, *Phys. Rev. Lett.* **98** (14), 146401 (2007). doi:10.1103/PhysRevLett.98.146401
- [80] X. Gao, F. Ramezanghorbani, O. Isayev, J.S. Smith and A.E. Roitberg, *J. Chem. Inf. Model.* **60** (7), 3408–3415 (2020). doi:10.1021/acs.jcim.0c00451
- [81] K.T. Schütt, H.E. Saucedo, P.-J. Kindermans, A. Tkatchenko and K.-R. Müller, *J. Chem. Phys.* **148** (24), 241722 (2018). doi:10.1063/1.5019779
- [82] M. Haghghatlari, J. Li, X. Guan, O. Zhang, A. Das, C.J. Stein, F. Heidar-Zadeh, M. Liu, M. Head-Gordon, L. Bertels, H. Hao, I. Leven and T. Head-Gordon, *Digital Disc.* **1**, 333–343 (2022). doi:10.1039/D2DD00008C.
- [83] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga and A. Desmaison, *Adv. Neural. Inf. Process. Syst.* **32** (2019).
- [84] X. Guan, A. Das, C.J. Stein, F. Heidar-Zadeh, L. Bertels, M. Liu, M. Haghghatlari, J. Li, O. Zhang, H. Hao, I. Leven, M. Head-Gordon and T. Head-Gordon, *Sci. Data.* **9** (1), 1–7 (2022). doi:10.1038/s41597-021-01104-5
- [85] R.S. Mulliken, *J. Chem. Phys.* **23**, 1833–1840 (1955). doi:10.1063/1.1740588.
- [86] P. Bultinck, C. Van Alsenoy, P.W. Ayers and R. Carbó-Dorca, *J. Chem. Phys.* **126**, 144111 (2007). doi:10.1063/1.2715563.
- [87] T.A. Manz and D.S. Sholl, *J. Chem. Theory. Comput.* **8**, 2844–2867 (2012). doi:10.1021/ct3002199.
- [88] A.E. Reed, R.B. Weinstock and F. Weinhold, *J. Chem. Phys.* **83**, 735–746 (1985). doi:10.1063/1.449486.
- [89] A.D. Mackerell, *J. Comput. Chem.* **25** (13), 1584–1604 (2004). doi:10.1002/(ISSN)1096-987X
- [90] E. Jurrus, D. Engel, K. Star, K. Monson, J. Brandi, L.E. Felberg, D.H. Brookes, L. Wilson, J. Chen, K. Liles, M. Chun, P. Li, D.W. Gohara, T. Dolinsky, R. Konecny, D.R. Koes, J.E. Nielsen, T. Head-Gordon, W. Geng, R. Krasny, G.W. Wei, M.J. Holst, J.A. McCammon and N.A. Baker, *Protein Sci.* **27** (1), 112–128 (2018). doi:10.1002/pro.3280.