Chemical Physics Letters 661 (2016) 42–47

Contents lists available at ScienceDirect

Chemical Physics Letters

journal homepage: www.elsevier.com/locate/cplett

Editor's choice

Accelerating *ab initio* molecular dynamics simulations by linear prediction methods

Jonathan D. Herr, Ryan P. Steele *

Department of Chemistry, University of Utah, 315 South 1400 East, Salt Lake City, UT 84112, United States
Henry Eyring Center for Theoretical Chemistry, University of Utah, 315 South 1400 East, Salt Lake City, UT 84112, United States

ARTICLE INFO

Article history:
Received 31 July 2016
In final form 19 August 2016
Available online 21 August 2016

Keywords:
Ab initio molecular dynamics
Fock extrapolation
Linear prediction

ABSTRACT

Acceleration of *ab initio* molecular dynamics (AIMD) simulations can be reliably achieved by extrapolation of electronic data from previous timesteps. Existing techniques utilize polynomial least-squares regression to fit previous steps' Fock or density matrix elements. In this work, the recursive Burg 'linear prediction' technique is shown to be a viable alternative to polynomial regression, and the extrapolation-predicted Fock matrix elements were three orders of magnitude closer to converged elements. Accelerations of 1.8–3.4× were observed in test systems, and in all cases, linear prediction outperformed polynomial extrapolation. Importantly, these accelerations were achieved without reducing the MD integration timestep.

1. Introduction

Ab initio molecular dynamics (AIMD) simulation techniques employ 'on-the-fly' generation of quantum chemistry molecular forces to propagate classical molecular dynamics trajectories on the Born-Oppenheimer potential energy surface. The ability to accurately capture strong polarization, charge transfer, and bond-breaking/-making events motivates the use of such techniques in chemical simulations of real-time dynamics, equilibrium sampling, and spectroscopic response. However, the many-orders-of-magnitude cost increase of AIMD, relative to commonly used molecular mechanics force fields, makes such simulations a daunting computational challenge for complex molecular systems.

Density-functional theory [1] (DFT) is the most commonly used quantum chemistry method for large AIMD simulations, although Møller–Plesset perturbation theory [2] (MP2)-based simulations have also been performed. For Kohn-Sham DFT or the underlying Hartree-Fock (HF) calculation for MP2, a non-linear, self-consistent field (SCF) optimization procedure is solved iteratively [3], leading to an overall computational prefactor that is determined by the number of SCF cycles required for convergence to a chosen tolerance. Although advanced integral evaluation techniques [4,5] tackle the scaling of these simulations with respect to the size of the system, motivation remains to reduce the SCF-based prefactor, particularly for molecular systems that are large, open-shell, and/or undergoing bond-rearrangement events.

The small timesteps required for AIMD trajectories ($\leqslant$0.5 fs for hydrogen-containing systems) suggest that the bulk of the electronic structure has not appreciably changed between neighboring timesteps. Restarting the SCF procedure anew from a typically poor initial guess, therefore, would appear to be computationally wasteful. This notion has motivated two diverging routes to accelerating MD simulations.

One option is the *propagation* of electronic information to the next timestep. Extended-Lagrangian MD (ELMD) methods [6–8] accelerate—or altogether avoid—the SCF procedure by propagating Fock or density matrix information. Fidelity with the true Born-Oppenheimer surface is approximated through the extended-Lagrangian formalism. One often-cited advantage of this approach is the rigorous enforcement of total energy conservation. Tradeoffs include the need for typically smaller MD timesteps (thereby reducing the effectiveness of the computational acceleration), choice of a fictitious mass, and the potential for electronic 'lag' artifacts [8–10] that can appear in observables, such as vibrational spectra. More recent implementations [11–16] successfully bridge the ELMD-BOMD divide, although stable propagation of the electronic variables for reactive systems remains an area of active research. One of the clear advantages to time-reversibly propagating electronic information has been the demonstration by Niklasson that the SCF procedure does not need to be converged tightly—sometimes requiring only a single SCF step [13,15]—in order to produce energy-conserving trajectories.

* Corresponding author.
E-mail address: ryan.steele@utah.edu (R.P. Steele).

http://dx.doi.org/10.1016/j.cplett.2016.08.050
0009-2614/

The second option involves *extrapolation* of electronic information from previous timesteps. Motivated by the observation that Fock matrix elements, for example, are simple oscillatory functions of time, these approaches attempt to fit previous timesteps' information to a simple form that can be used as an improved initial guess for the SCF procedure [17–20]. This approach allows for longer timesteps while still reducing the number of requisite SCF cycles. In its simplest version, this approach uses the converged Fock matrix from the previous timestep as this initial guess, which is consistent with the intuitive notion that the electronic structure has not changed appreciably. However, as has been noted in previous analyses [18], such a choice quickly fails when loose SCF convergence settings are employed. The use of the previous step's Fock matrix introduces a hysteresis effect that leads to rather severe non-conservation of the energy. Over only short trajectories, the energy drift swamps the inherent energy fluctuations of the chosen integration scheme. While higher-order polynomials somewhat dampen this energy-conservation concern, its effects should eventually be borne out in long MD trajectories since the history-dependent guess renders the propagator formally non-symplectic. The practical cure to this pathology, of course, is to more tightly converge the SCF procedure. As will be shown below, this approach rectifies energy-drift concerns for very long trajectories (by AIMD standards), but it also necessarily leads to additional computational cost overhead.

The single SCF cycle performed in the 'propagation' approach could potentially lead to unstable electronic states for slow-to-converge SCF cases, such as open-shell ions or systems near bond-arrangement events, and these cases require additional testing in the literature. Because of our interest in these difficult systems [21], the present analysis focuses solely on the 'extrapolation' approach, with the intent of improving existing acceleration techniques within tight SCF convergence settings. An investigation into whether the behavior observed in the present analysis can be combined with rigorously time-reversible propagators is reserved for future studies.

Existing Fock extrapolation techniques employ a polynomial least-squares regression (PLSR) of N previous timesteps' information, to a polynomial of a chosen order O. A recent analysis [22] has shown that an optimal N may exist, at least for a given timestep and chemical system, although no assessment of O-dependence was performed. Nonetheless, past studies have shown that with only a handful of previous steps' data, roughly 2-fold reductions in the number of Fock builds can be achieved. Extension to MP2 simulations, in both the SCF and gradient response terms, has also been demonstrated [23].

While this standard PLSR approach is appropriate for accurately fitting the data within the chosen time window, it is decidedly *not* optimal for extrapolating information outside of this window. Given that the Fock data throughout an MD trajectory is locally oscillatory (with some slight noise, stemming from incomplete convergence) and globally stochastic, techniques that are well-suited to signal prediction were conjectured to be a better choice. In particular, so-called *linear prediction* techniques—not to be confused with linear PLSR extrapolation ($O = 1$)—are designed to accurately represent exactly this type of behavior. This study presents an implementation of the Burg Linear Prediction (BLP) algorithm for Fock extrapolation, as well as rigorous tests of its efficacy and efficiency for accelerating AIMD simulations.

2. Methods

The basic components of the BLP algorithm are presented in this section, including a discussion of the distinguishing features of the BLP and PLSR approaches. This presentation is followed by a description of the techniques used to assess the methods throughout the remainder of the study.

2.1. Burg Linear Prediction

The BLP algorithm [24,25], in an extrapolated AIMD context, fits N previous timesteps' atomic-orbital (AO) basis Fock matrix elements to a *recursive series* with L coefficients forward [f] in time:

$$\tilde{F}^{[f]}_{\mu\nu}(t) = \sum_{i=1}^{L} c_i F_{\mu\nu}(t-i) \tag{1}$$

The Burg algorithm augments this simpler Levinson-Durbin recursion definition by also considering a fit to backward [b] predictions:

$$\tilde{F}^{[b]}_{\mu\nu}(t) = \sum_{i=1}^{L} c_i F_{\mu\nu}(t+i) \tag{2}$$

In this notation, the predicted element is denoted as $\tilde{F}$, and t acts as a timestep index, rather than a physical simulation time. In all analyses performed in this work, the timesteps were evenly spaced, although the LPB prescription is not restricted to this form. The coefficients are obtained using linear least-squares minimization of the *sum* of the forward and backward mean-square deviations, χ^2, defined over the fitting length N as

$$\begin{aligned}\chi^{2[f]}_{\mu\nu} &= \sum_{t=L}^{N}\left[F_{\mu\nu}(t) - \tilde{F}^{[f]}_{\mu\nu}(t)\right]^2 \\ &= \sum_{t=L}^{N}\left[F_{\mu\nu}(t) - \left(\sum_{i=1}^{L} c_i F_{\mu\nu}(t-i)\right)\right]^2\end{aligned} \tag{3}$$

$$\begin{aligned}\chi^{2[b]}_{\mu\nu} &= \sum_{t=L}^{N}\left[F_{\mu\nu}(t) - \tilde{F}^{[b]}_{\mu\nu}(t)\right]^2 \\ &= \sum_{t=L}^{N}\left[F_{\mu\nu}(t) - \left(\sum_{i=1}^{L} c_i F_{\mu\nu}(t+i)\right)\right]^2\end{aligned} \tag{4}$$

This optimization of $\chi^2 = \chi^{2[f]} + \chi^{2[b]}$, rather than $\chi^{2[f]}$ alone, distinguishes BLP from simpler recursion approaches and introduces considerable stability to the method. Further algorithmic details can be found in standard numerical texts, including the one used in the present implementation [24].

The BLP algorithm, therefore, generates a weighted linear combination of existing Fock matrix elements. In this sense, it is superficially similar to a PLSR fit,

$$\tilde{F}^{[f]}_{\mu\nu}(t) = \sum_{i=0}^{O} c_i \cdot t^i \tag{5}$$

However, as will be shown below, the BLP algorithm was found to be much more powerful for extrapolation purposes. For sums of Fourier signals, for example, preliminary tests of this approach yielded nearly quantitative reproduction of the signal for thousands of extrapolated signal cycles, rather than the 1–2 that was afforded with PLSR. Of course, in an AIMD context, the fit will be recomputed at each timestep, meaning that the algorithm only needs to be accurate to the next timestep. Even in this regime, though, the BLP approach led to several-orders-of-magnitude reductions in error for values of this subsequent step.

The use of PLSR requires a choice of N and O, compactly notated in the remainder of this work an (N,O) extrapolation. Commonly used choices in previous studies [17,18] included (4,2) to (20,10) extrapolations. Since this approach is intended as an extrapolation, rather than a fit, increases in N or O do not necessarily yield better SCF guesses, and the optimal choice is often simulation-specific