

ARTICLE OPEN



Trajectory sampling and finite-size effects in first-principles stopping power calculations

Alina Kononov¹✉, Thomas W. Hentschel², Stephanie B. Hansen³ and Andrew D. Baczewski¹✉

Real-time time-dependent density functional theory (TDDFT) is presently the most accurate available method for computing electronic stopping powers from first principles. However, obtaining application-relevant results often involves either costly averages over multiple calculations or ad hoc selection of a representative ion trajectory. We consider a broadly applicable, quantitative metric for evaluating and optimizing trajectories in this context. This methodology enables rigorous analysis of the failure modes of various common trajectory choices in crystalline materials. Although randomly selecting trajectories is common practice in stopping power calculations in solids, we show that nearly 30% of random trajectories in an FCC aluminum crystal will not representatively sample the material over the time and length scales feasibly simulated with TDDFT, and unrepresentative choices incur errors of up to 60%. We also show that finite-size effects depend on ion trajectory via “ouroboros” effects beyond the prevailing plasmon-based interpretation, and we propose a cost-reducing scheme to obtain converged results even when expensive core-electron contributions preclude large supercells. This work helps to mitigate poorly controlled approximations in first-principles stopping power calculations, allowing 1–2 order of magnitude cost reductions for obtaining representatively averaged and converged results.

npj Computational Materials (2023)9:205; <https://doi.org/10.1038/s41524-023-01157-7>

INTRODUCTION

High-performance computing has revolutionized materials science, enabling prediction, design, and unprecedented understanding of materials properties to complement and accelerate experimental efforts. Modeling dynamic, nonlinear responses to stimuli such as laser and particle irradiation falls among the most computationally demanding types of materials simulations, requiring real-time evolution of extended systems containing hundreds of atoms and thousands of quantum-mechanical electrons^{1–3}. For materials in extreme conditions, high temperatures result in orders-of-magnitude increases in the number of partially occupied electronic orbitals, either requiring additional approximations^{4–6} or further escalating computational resource requirements to millions of CPU hours or more per calculation⁷. In this context, deliberate design of simulations is crucial for maximizing insight while maintaining feasible computational costs.

Here, we focus on calculations of electronic stopping power, the rate at which a moving particle loses energy to electrons. This fundamental quantity is critically important to diverse fields. For example, radiation therapy relies on stopping powers to predict particle ranges and precisely target tumors⁸. Stopping powers also underlie radiation damage to materials in space and nuclear energy applications^{9,10}. In materials imaging and processing techniques, energy deposition by focused ion beams relates to electron emission, sample damage, and defect engineering^{11,12}. Finally, achieving ignition in fusion energy research relies on fusion products redepositing their kinetic energy into the fuel¹³.

First-principles simulations using real-time time-dependent density functional theory (TDDFT) can offer accurate predictions of electronic stopping powers and insights into underlying physical processes^{14–20}, often in more detail than possible experimentally. However, computing average stopping powers

that are comparable to experimentally observable and practically relevant values can pose a challenge. While an individual TDDFT calculation simulates a single projectile traversing a specific path, stopping power experiments measure energy loss distributions for finite-width ion beams, often incident on polycrystalline or disordered samples^{21–23}. Moreover, applications either also employ finite-width ion beams (e.g., materials imaging and processing) or involve randomly oriented radiation (e.g., equipment in space and fusion fuel). Therefore, sensitivities to the projectile’s trajectory can limit the utility of TDDFT stopping power predictions. Such sensitivities often occur for projectile velocities beyond the Bragg peak and mainly affect contributions from core electrons, or more generally, spatially non-uniform electronic orbitals^{14,17,24}.

Meaningful average stopping powers can still be obtained from first principles by averaging^{6,17,25–27} or carefully integrating^{28,29} the results of several TDDFT calculations using distinct projectile trajectories. The significant computational cost of this approach makes it tempting to select a single trajectory presumed to be representative of an ensemble average. For solids, one possible choice is the centroid trajectory, wherein the projectile travels along a crystallographic direction with a path given by the geometric centroid of a symmetry-irreducible cross-section of the crystal structure^{25,30} (see Fig. 1). This method appears adequate when core-electron contributions are small, but becomes inaccurate for fast projectiles and high-Z targets²⁹. Alternatively, a randomly chosen trajectory can achieve good agreement with empirical data even when core-electron contributions are important¹⁴. In this case, trajectory choice may constitute an uncontrolled approximation, thus necessitating quantitative methods for assessing any given choice.

Recently, Gu et al. developed an innovative pre-sampling approach³¹ that averages results from several short trajectories

¹Center for Computing Research, Sandia National Laboratories, Albuquerque, NM, USA. ²School of Applied & Engineering Physics, Cornell University, Ithaca, NY, USA. ³Pulsed Power Sciences Center, Sandia National Laboratories, Albuquerque, NM, USA. ✉email: akonono@sandia.gov; adbacz@sandia.gov

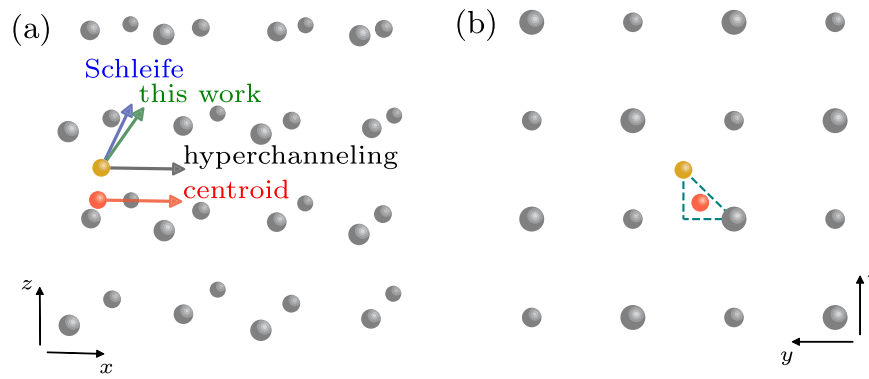


Fig. 1 Trajectories through FCC aluminum that are discussed in this work. **a** side view and **b** back view of the geometry. Red and yellow balls represent the projectile's initial position, and arrows indicate the direction of motion. The hyperchanneling and centroid trajectories are parallel to a lattice vector, unlike the off-channeling trajectories considered by Schleife et al.¹⁴ and in this work. In **b** the dashed triangle indicates the symmetry-irreducible cross-section of the lattice. The red ball lies at the centroid of this triangle.

carefully selected such that, in aggregate, they representatively sample a disordered system. Here, we present a complementary method that uses a quantitative metric to guide a priori selection of a single, representative projectile trajectory for first-principles calculations of electronic stopping power. Using proton stopping in aluminum as an exemplar, we demonstrate the utility of such an approach even in a crystalline material, as we find that achieving agreement with empirical data requires a high-quality trajectory despite widespread presumptions that a random choice suffices. While both methods can reduce computational costs associated with averaging over multiple long trajectories and help obtain reliable results even in the absence of experimental data to validate against, we expect that our approach will be more efficient than the one proposed by Gu et al.³¹ in the case of heavy ions that experience relatively long transient behavior³² before entering the steady-state stopping regime.

In addition to trajectory-dependent core-electron contributions, finite-size effects limit the accuracy of first-principles stopping power calculations, particularly for fast projectiles with velocities above the Bragg peak. Typically, TDDFT calculations are expected to underestimate stopping powers by neglecting long-wavelength plasmonic excitations that a finite periodic supercell cannot support³³. We examine variations in computed stopping powers for different supercell sizes and projectile trajectories and reveal significant departures from this model of finite-size errors. Finally, we extend our trajectory optimization framework with a second quantitative metric that enables a priori selection of trajectories that minimize finite-size effects.

Together, the two trajectory metrics developed in this work allow deeper understanding and opportunities for mitigation of previously poorly controlled approximations in TDDFT simulations of energetic particles traversing matter. This contribution advances the accuracy and efficiency of first-principles stopping power calculations with wide-ranging implications for computational studies relating to radiation therapy, materials in extreme conditions, ion-beam imaging and patterning techniques, and self-heating of fusion fuel.

RESULTS

Sampling close collisions

Obtaining an application-relevant result from a single TDDFT stopping calculation requires choosing a trajectory along which the projectile experiences an environment that quantitatively resembles those likely to occur in that application. In particular, close collisions between the projectile and host nuclei can involve (semi)core-electron excitations that introduce sharp features in the stopping forces and can dominate the average stopping

power¹⁴. Thus, representatively sampling close collisions is essential for computing accurate stopping powers, and we expect that the distance between the projectile and the nearest host nuclei at all points along the trajectory provides a compact description of the environment that determines the average stopping power.

This notion previously inspired a method for assessing the quality of different trajectories in disordered systems³¹. In what follows, we present a quantitative metric that enables rigorous comparisons among different trajectories and their attendant stopping powers. Furthermore, we show that such an approach is critical for selecting trajectories even in crystalline materials because not all randomly oriented off-channeling trajectories are equally representative of an average environment over the course of a typical few-fs simulation time. Unrepresentative finite-length off-channeling trajectories lead to poor estimates of average stopping power and could skew averages based on naive trajectory sampling.

To evaluate the quality of a given trajectory over a path length x , we first calculate the distribution $P_{\text{traj}}(\delta_{\text{NN}}, x)$ of nearest-neighbor distances δ_{NN} between the host nuclei and the projectile along the trajectory. Since the ultimate TDDFT simulations only span a few fs, the projectile does not have time to significantly deflect from its initial velocity and we assume that it travels in a straight line. Thus, computing $P_{\text{traj}}(\delta_{\text{NN}}, x)$ relies only on the geometric specification of the supercell and does not require expensive TDDFT simulations. We exclude data during the first 4 Å of the projectile's path since the ultimate stopping power extraction will ignore this early transient regime (see Supplementary Note 1). This cutoff distance is specific to a proton projectile, and higher Z projectiles would have longer transient regimes³², further motivating optimization of a single representative trajectory.

We then compare $P_{\text{traj}}(\delta_{\text{NN}}, x)$ to an ideal distribution $P_{\text{ideal}}(\delta_{\text{NN}})$ generated by calculating distances to nearest host nuclei from randomly sampled points within the supercell. This choice of $P_{\text{ideal}}(\delta_{\text{NN}})$ represents the distribution experienced by randomly oriented radiation or a focused ion-beam interacting with randomly oriented grains within a polycrystalline sample. Earlier work by Gu et al. instead generated a reference distribution by sampling along a 500 μm-long trajectory³¹, a method that should lead to the same distribution provided that the selected reference trajectory representatively samples the entire supercell. Other choices may be more suitable depending on the specific application, e.g. a focused ion-beam aligned with a lattice vector of a single crystal.

The Hellinger distance D_H between the two nearest-neighbor distributions, given by

$$D_H^2(x) = 1 - \int_0^\infty \sqrt{P_{\text{traj}}(\delta_{\text{NN}}, x) P_{\text{ideal}}(\delta_{\text{NN}})} d\delta_{\text{NN}}, \quad (1)$$

provides a quantitative measure of how well a trajectory samples the simulation cell. Notably, D_H is bounded between 0 and 1, with $D_H = 1$ achieved when the two distributions have no overlap and $D_H = 0$ achieved only when the two distributions are identical. Supplementary Note 2 discusses numerical evaluation of Eq. (1).

Minimizing D_H for a target simulation length will enable selection of optimally representative trajectories. We chose the Hellinger distance over other possible ways to quantify the similarity of two distributions because it satisfies the mathematical properties of a metric. In particular, D_H is symmetric and obeys the triangle inequality, allowing sensible comparisons of two different trajectories by replacing $P_{\text{ideal}}(\delta_{\text{NN}})$ in Eq. (1) with the nearest-neighbor distribution of the second trajectory. We compare the D_H metric used in this work with the overlap index used by Gu et al.³¹ in Supplementary Note 3.

As the proton travels, $P_{\text{traj}}(\delta_{\text{NN}}, x)$ evolves and so does $D_H(x)$. All but a few pathological random trajectories will asymptotically approach $D_H = 0$ as the total distance traversed tends to infinity, but the impact of finite-size effects (e.g., spurious interactions between the projectile and its wake) will also grow with the total distance. Thus, a “good” trajectory should achieve a small D_H after the projectile travels a relatively short distance so that a single, reasonably short TDDFT calculation suffices to obtain an accurate estimate of the average stopping power. In the following, we investigate the behavior of D_H for a range of trajectories and the implications for stopping power results.

In Fig. 2a we consider the case of FCC aluminum and compare the nearest-neighbor distances occurring along several trajectories (including those illustrated in Fig. 1) and the corresponding distributions $P_{\text{traj}}(\delta_{\text{NN}}, x)$ sampled over a path length of $x = 80 \text{ \AA}$. Both the hyperchanneling and centroid trajectories lie parallel to a lattice vector and thus sample periodic environments, producing nearly static nearest-neighbor distributions and asymptotically constant D_H values of about 0.67 and 0.33, respectively (see Fig. 2b). Of course, D_H remains relatively large for the hyperchanneling trajectory because its nearest-neighbor distribution is severely skewed toward large δ_{NN} .

While the centroid trajectory significantly improves D_H over the hyperchanneling trajectory, its nearest-neighbor distribution is bimodal, oversampling near points of closest and furthest approach at $\delta_{\text{NN}} \approx 0.75$ and 1.5 \AA while lacking close collisions with $\delta_{\text{NN}} < 0.5 \text{ \AA}$. Failure to capture the ideal nearest-neighbor distribution explains the poor performance reported for the centroid trajectory in regimes where core-electron excitations contribute significantly to electronic stopping power²⁹. We also find that other channeling trajectories with different impact parameters are similarly restricted to $D_H > 0.3$ (see Fig. 3a) and therefore do not representatively sample the supercell.

In contrast, both the off-channeling trajectory considered by Schleife et al.¹⁴ and the other “good” off-channeling trajectory identified in this work approximate the ideal distribution more closely and continue to reduce D_H as the proton travels farther, reaching much lower D_H values of 0.05–0.06 by 80 \AA . Even when the projectile starts at a high-symmetry point within the crystal structure, the vast majority of possible directions of motion achieve $D_H < 0.3$ by the time the projectile travels 80 \AA (see Fig. 3b). Some but not all exceptions lie at channeling directions. For a finite simulation length, D_H can be very sensitive to trajectory direction, with small changes in the trajectory angle sometimes leading to order-of-magnitude changes in D_H . Among fully random trajectories where the projectile’s initial position and direction of motion are both uniformly sampled, only 1.3%

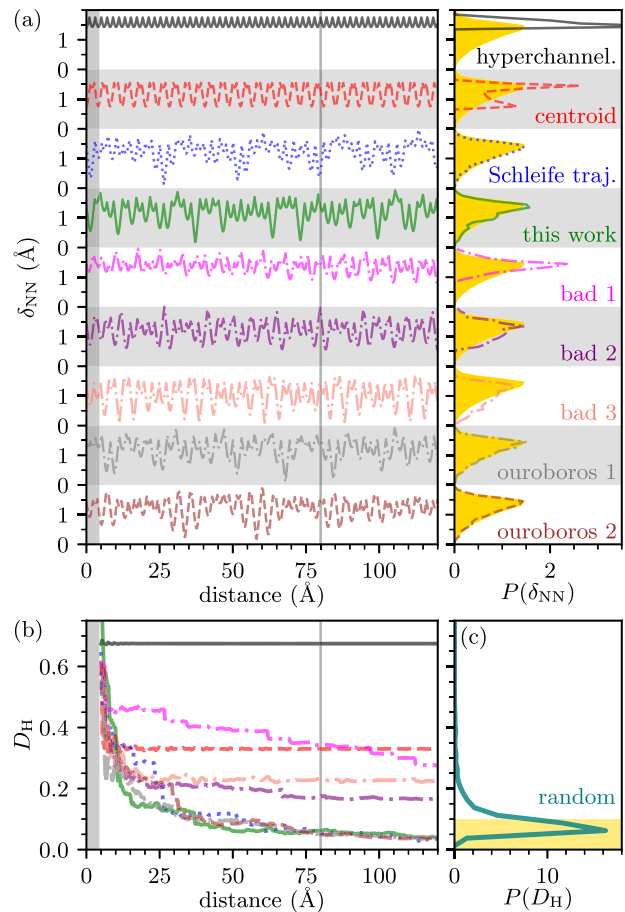


Fig. 2 Geometric analysis of various trajectories through ambient aluminum (see Figs. 1 and 3). **a** Nearest-neighbor distances (left) and distributions (right) experienced by the projectile along each trajectory. The nearest-neighbor distribution after the projectile travels 80 \AA along each trajectory is compared to the ideal distribution sampled from random points within the supercell (yellow). **b** Hellinger distance (see Eq. (1)) from the ideal distribution achieved by each trajectory as a function of total distance traveled by the projectile. The gray area indicates the 4 \AA transient regime ignored in the analysis. **c** Distribution of D_H values achieved at 80 \AA by 1440 random trajectories with uniformly sampled initial positions and directions of motion. Yellow shading indicates $D_H < 0.1$, the target identified for representative sampling of the supercell in this work.

perform as poorly as channeling trajectories with D_H still exceeding 0.3 after the proton traverses 80 \AA (see Fig. 2c).

To verify the utility of our trajectory metric and deduce a threshold D_H value for accurate stopping power predictions, we also consider three “bad” random trajectories that either undersample (bad 1 and bad 2) or oversample (bad 3) close collisions with $\delta_{\text{NN}} < 0.5 \text{ \AA}$ (see Fig. 2a). Unlike channeling trajectories, D_H continues to decrease along these bad off-channeling trajectories, but much more slowly than the good off-channeling trajectories, only achieving $D_H = 0.17$ – 0.35 by the time the proton travels 80 \AA .

Finally, we perform TDDFT stopping power calculations as described in “Methods” for the off-channeling trajectories examined above. Although the two good trajectories exhibit differing dynamical behavior, the average stopping powers extracted after the proton travels 80 \AA agree within 1% and reproduce empirical data from the SRIM database³⁴ within 3% (see Fig. 4). Meanwhile, stopping powers computed using the bad trajectories deviate from empirical data by up to 60%. As expected, trajectories that undersample close collisions (bad 1

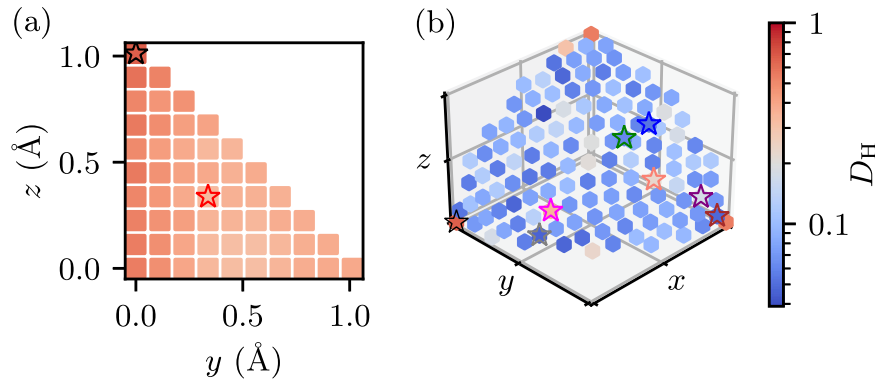


Fig. 3 The trajectory metric D_H after a proton travels 80 Å along different trajectories. **a** The proton moves in the x direction and its initial position is varied within a symmetry-irreducible cross-section of the lattice (see Fig. 1b). **b** The proton begins at a high-symmetry point within the lattice (the yellow ball in Fig. 1) and the direction of its velocity is varied over an octant of the unit sphere. The x , y , and z directions correspond to channeling trajectories. Stars outline trajectories examined in this work with colors corresponding to those in Fig. 2.

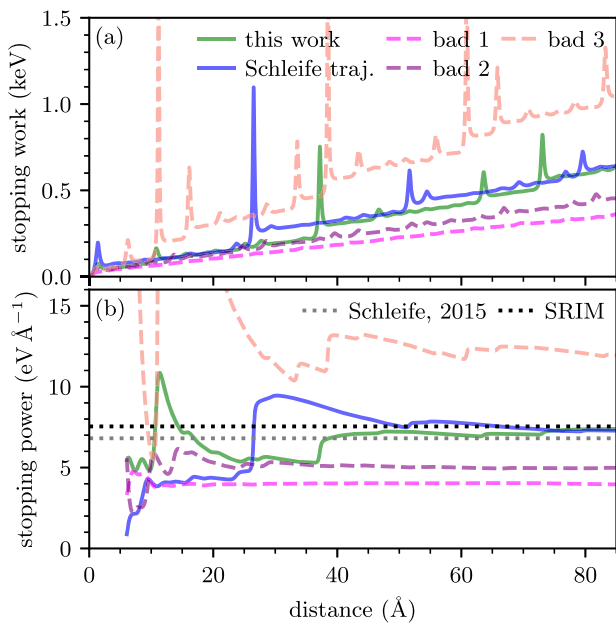


Fig. 4 TDDFT simulation results along some of the trajectories considered in Figs. 1–3. Instantaneous **a** stopping work and **b** extracted stopping power as a proton traverses aluminum with 4 atomic units of velocity. Results predicted by an earlier TDDFT study¹⁴ and the SRIM empirical model³⁴ are shown as gray and black dotted lines, respectively. Note that the blue curves represent new calculations using the same trajectory as the earlier study¹⁴.

and bad 2) underestimate stopping power while the trajectory that oversamples close collisions (bad 3) overestimates stopping power. Based on these findings, we propose that $D_H \lesssim 0.1$ suffices for representative sampling of the nearest-neighbor distribution in these stopping power calculations. Notably, 27% of random trajectories still exceed this threshold after the projectile travels 80 Å (see Fig. 2c), highlighting the importance of careful trajectory selection for accurate and efficient stopping power predictions.

Mitigating finite-size effects

Finite periodic supercells limit the wavelength of plasmonic excitations that a plane-wave TDDFT calculation can capture, often leading to underestimated electronic stopping powers at high projectile velocities³³. Finite-size errors caused by this plasmon cutoff can be estimated from linear response theory³⁵, which describes the stopping power in terms of the frequency and wave-

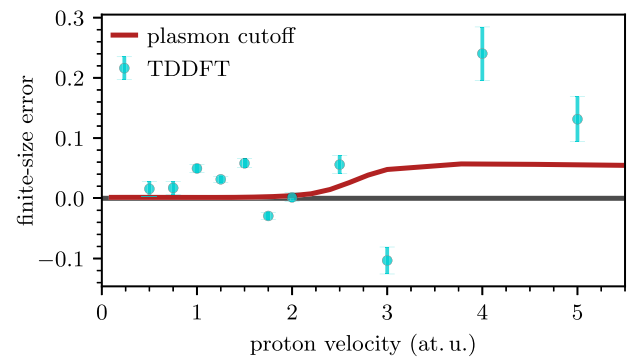


Fig. 5 Fractional finite-size errors estimated as the relative difference between valence-electron stopping powers computed with $(12.15\text{Å})^3$ and $(16.2\text{Å})^3$ cubic supercells, $(S(16.2) - S(12.15)) / S(16.2)$. TDDFT results are compared to dielectric-based predictions using corresponding plasmon wavelength cutoffs (see Eq. (2) and accompanying text). Error bars estimate uncertainties in the TDDFT data arising from variations in extracted average stopping powers as a function of total simulation time, taken as the standard deviation of results after the proton travels between 60 and 80 Å.

vector dependent dielectric function $\epsilon(k, \omega)$:

$$S(v) = \frac{2Z^2}{\pi v^2} \int_0^\infty \frac{dk}{k} \int_0^{kv} d\omega \omega \operatorname{Im} \left[\frac{-1}{\epsilon(k, \omega)} \right], \quad (2)$$

where $Z=1$ and v are the projectile charge and velocity, respectively. Imposing an upper limit of $k_{\text{cut}} = 2\pi/L$ for the k integral in Eq. (2), where L is the length of the cubic supercell, provides an estimate of the contribution from long-wavelength plasmons that the TDDFT calculations neglect.

Here, we employ the Mermin model dielectric function³⁶, which includes electron-ion collisions to improve over the uniform electron gas assumption underlying the Lindhard dielectric function and thereby achieves closer agreement with TDDFT response properties⁷. We then evaluate the integrals in Eq. (2) numerically as described in recent work⁷ with a constant electron-ion collision frequency of 0.1 atomic units (at. u.). The contribution of long-wavelength plasmons becomes significant when the integration limits contain the plasmon pole, i.e., when $k_{\text{cut}} v \gtrsim \omega_p$, where $\omega_p \approx 16$ eV is the aluminum plasma frequency. Indeed, for $L = 12.15$ Å the linear response formalism predicts sizeable finite-size errors for velocities above $\omega_p/k_{\text{cut}} \approx 2$ at. u. (see Fig. 5).

However, we find that differences between stopping powers computed using TDDFT with different size supercells only loosely correlate with finite-size errors estimated from linear response. As shown in Fig. 5, the relative difference between TDDFT results




Fig. 6 Average stopping powers as a proton with 4 atomic units of velocity traverses aluminum along four different off-channeling trajectories that achieve $D_H < 0.1$ by the time the proton travels 80 \AA . Results are compared for **a** the trajectory used throughout this work, **b** the trajectory used in an earlier study¹⁴, and **c**, **d** two other off-channeling trajectories. Total stopping powers (purple) are decomposed into contributions from valence (blue) and core (red) electrons for the 256-atom, $(16.2 \text{ \AA})^3$ supercell used throughout this work (solid) and a smaller, 108-atom, $(12.15 \text{ \AA})^3$ supercell (dashed). Valence-electron results for a larger, 500-atom, $(20.25 \text{ \AA})^3$ supercell are also shown (dotted). Horizontal lines indicate converged values for the first trajectory with a $(16.2 \text{ \AA})^3$ supercell.

calculated using $(12.15 \text{ \AA})^3$ and $(16.2 \text{ \AA})^3$ cubic supercells often significantly exceeds the values predicted using the dielectric model. This discrepancy is especially notable for slow protons with $v \lesssim 2$ at. u., where the TDDFT data still contains finite-size errors even though long-wavelength plasmons cannot be excited within linear response. Moreover, for some proton velocities, the smaller supercell actually produces greater stopping powers than the larger supercell, a negative finite-size error entirely inconsistent with the plasmon-based interpretation of finite-size effects.

As another effect not captured by the plasmonic model, the choice of trajectory influences finite-size errors independently of the fidelity of close-collision sampling. In Fig. 6, we compare stopping powers computed using three different size supercells and four different proton trajectories. In addition to the two successful trajectories examined in “Sampling close collisions” (Fig. 6a, b), we consider two other choices among those shown in Fig. 3b that achieve comparably small D_H values while traversing the simulation cell at different angles (Fig. 6c, d). Contributions from valence and core ($2s$ and $2p$) electrons were isolated through the use of different pseudopotentials (see “Methods” for more details). Close agreement of the converged core-electron stopping powers verifies that each trajectory adequately samples close collisions with aluminum ions.

While the core-electron contributions in Fig. 6 show little variation with supercell size, the valence-electron stopping powers are quite sensitive to both finite-size effects and projectile trajectory. Computed valence-electron stopping powers do not grow monotonically with increasing supercell size for the trajectories examined in Fig. 6a–c, sometimes exhibiting negative finite-size errors similar to the behavior appearing in Fig. 5. The magnitude of discrepancies between results computed with different size supercells also depends on the projectile trajectory:

for instance, Fig. 6d shows much greater differences among valence-electron contributions than the other cases.

We attribute the occasionally negative and surprising trajectory-dependent nature of finite-size errors in TDDFT stopping power calculations to artificial interactions with previously excited electrons. In particular, if the projectile passes near its earlier path after re-entering a periodic supercell, then it interacts with an excited electron density rather than pristine material, distorting stopping power results. Such “ouroboros” effects are especially severe for channeling trajectories: upon reentry, the projectile traverses the exact same path, at which point the instantaneous stopping power begins to depend on supercell size¹⁴. Ouroboros effects can also pollute results for off-channeling trajectories, since the projectile has a finite interaction radius that may partially overlap with previously excited regions. Even if the projectile remains relatively far from previously traversed material, it could still interact with earlier electronic excitations that have propagated into its path. We distinguish these different types of artificial interactions as static and dynamic ouroboros effects, respectively.

The prevalence of ouroboros effects can be estimated by considering the minimum distance D_O between periodic images of the projectile’s path. This distance is upper-bounded by the simulation cell dimensions and decreases as the projectile crosses periodic boundaries. In the limit of large D_O , the projectile remains far from earlier excited material, and ouroboros effects should be minimal. In a metallic system, localized excitations only involve core atomic orbitals and are well-confined within the Wigner-Seitz radius r_{WS} of the host nuclei. Therefore, we expect that ensuring $D_O > r_{WS}$ (1.58 \AA for FCC aluminum) eliminates static ouroboros effects. Even for the lowest D_O tested, 1.3 \AA for the smallest supercell in Fig. 6d, finite-size effects do not influence the core-electron stopping power because the $2s$ and $2p$ electrons are further localized within about 0.5 \AA of aluminum nuclei.

Dynamic ouroboros effects, on the other hand, are harder to characterize and avoid. Given a Fermi velocity of $\sim 20 \text{ \AA fs}^{-1}$ in aluminum, single-particle excitations could traverse the $(16.2 \text{ \AA})^3$, 256-atom supercells used throughout this work in < 1 fs, a time scale comparable to the duration of the TDDFT simulations. First-principles calculations of plasmon dispersion in aluminum³⁷ suggest similar propagation speeds for collective excitations. Since an off-channeling projectile must travel around 80 \AA in order to adequately sample the nearest-neighbor distribution in this material (see “Sampling close collisions”), it crosses the periodic boundaries multiple times, leading to typical D_O distances of around 3 \AA that electronic excitations traverse over an even shorter time scale.

However, proton projectiles can be expected to induce fairly weak charge perturbations, as evidenced by the relative success of linear response treatments of proton stopping powers^{7,14,38}. So, dynamic ouroboros effects could be small compared to other sources of error. Furthermore, alternating artificial interactions with excitations that involve excess or reduced electron density relative to the pristine material could have partially canceling influences on the average stopping power. Indeed, the valence-electron stopping powers reported in Fig. 6 differ among each other by at most 31% of the total stopping power, whereas poor sampling of close collisions affected total stopping powers by almost a factor of 2 in Fig. 4. In Fig. 7, we show that $D_O \gtrsim 3 \text{ \AA}$ already achieves acceptable convergence of valence-electron stopping powers.

As an alternative to maximizing D_O for a single, long trajectory, ouroboros effects could be minimized by combining results computed from many short trajectories that do not cross periodic boundaries. Such a strategy could help further disentangle different sources of finite-size errors, including ouroboros effects and the lack of long-wavelength plasmons. However, since each simulation begins with a transient period, this approach would

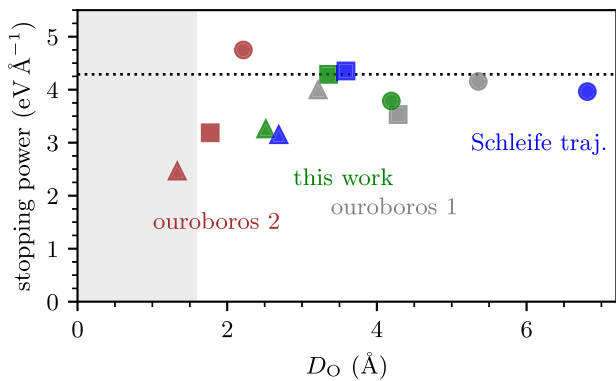


Fig. 7 Convergence of valence-electron stopping power with increasing D_O , the minimum distance between periodic images of the projectile's path, for a proton traversing aluminum with 4 atomic units of velocity. Both the stopping powers and D_O values reported were taken after the projectile traveled 80 Å through the material. Colors indicate proton trajectory, while symbols indicate supercell size: triangles, squares, and circles correspond to results from $(12.15 \text{ \AA})^3$, $(16.2 \text{ \AA})^3$, and $(20.25 \text{ \AA})^3$ supercells, respectively. The shaded area indicates $D_O < r_{WS}$, and the dotted line indicates stopping power extracted from the simulation parameters used throughout this work.

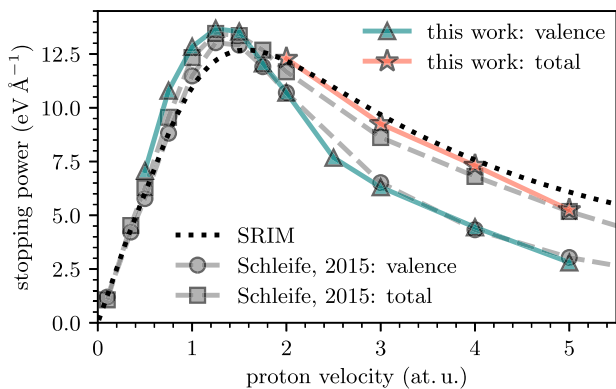


Fig. 8 Electronic stopping power of protons in ambient aluminum as predicted by this work, an earlier RT-TDDFT study¹⁴, and the SRIM empirical model³⁴. The valence-electron contribution was computed by pseudizing all but the outermost 3 electrons per aluminum atom, while the total stopping power also included excitations of explicitly modeled 2s and 2p semicore electrons.

require longer total simulation times to accumulate enough data within the steady-state regime to representatively sample close collisions in aggregate. At the same time, short trajectories may reduce the simulation cell size needed for converged results, creating a trade-off in computational cost. We expect that a hybrid strategy involving a few intermediate-length trajectories that may cross periodic boundaries but nonetheless maintain a large D_O would favorably balance these considerations, and the optimal trajectory length would depend on the stopping medium and projectile parameters.

Regardless of trajectory length, the fact that finite-size effects predominantly influence valence-electron contributions to stopping power can be exploited to reduce computational costs associated with first-principles stopping power calculations. Core contributions typically dominate computational costs because explicitly modeling core electrons dramatically increases the spectral range of the Kohn-Sham Hamiltonian, generally requiring higher plane-wave cutoff energies, smaller time steps, and/or more solver iterations in implicit time-stepping algorithms such as the one used in this work³⁹. Since core contributions are not very

sensitive to ouroboros effects or supercell size, they can be efficiently calculated using smaller supercells. Meanwhile, valence-electron contributions can be separately converged with respect to the D_O metric and supercell size while pseudizing core electrons to allow cheaper time evolution.

In this work, applying this scheme to separately compute core and valence-electron contributions using $(12.15 \text{ \AA})^3$ and $(16.2 \text{ \AA})^3$ supercells, respectively, would have allowed a seven-fold speedup over using $(16.2 \text{ \AA})^3$ supercells to compute both core and valence-electron contributions simultaneously. In fact, this strategy would have still resulted in a nearly three-fold speedup if the larger, $(20.25 \text{ \AA})^3$ supercell had been used to further reduce finite-size effects in the valence-electron contribution. These savings are relative to the $\sim 5 \times 10^5$ CPU-hours used to perform each production calculation in this work, which computed proton stopping powers in aluminum by including core contributions in $(16.2 \text{ \AA})^3$ supercells (see “Methods” for more methodological details).

DISCUSSION

Figure 8 compares our final electronic stopping results as a function of proton velocity to those predicted by the SRIM empirical model³⁴ and reported in an earlier TDDFT study¹⁴ that used the same level of theory (see “Methods”). Overall, we find good quantitative agreement between the two TDDFT datasets, with modest discrepancies arising from a combination of partially canceling factors. First, the earlier study¹⁴ verified convergence with respect to plane-wave cutoff energy for a channeling trajectory, but we show in Supplementary Note 4 that converging energy transferred during close collisions occurring along off-channeling trajectories requires higher cutoffs. In particular, we find that the 680 eV cutoff used in that work¹⁴ underestimates high-velocity stopping power by $\sim 5\%$.

Another source of discrepancies arises from different pseudopotentials: this work used the PAW method⁴⁰ within an extension of VASP^{41,42}, while the earlier study¹⁴ used norm-conserving pseudopotentials^{43,44} within Qb@ll^{45,46}. In Supplementary Note 5, we show that even when all other parameters are fixed or separately converged as appropriate, the PAW and VASP methodology used in this work produces about 10% lower stopping powers than the harder, norm-conserving pseudopotentials within Qb@ll. Additionally, a 2.5% uncertainty in extracted stopping powers arises from variations with respect to the data range included in the stopping power extraction (see Supplementary Note 1). Further benchmarking of different TDDFT codes, pseudopotentials, basis sets, and other methodological details will be important for reducing uncertainties in first-principles stopping data⁴⁷.

To enable detailed interpretation of discrepancies in computed results and reduce computational costs associated with obtaining representative average stopping powers, we presented a quantitative metric to evaluate the quality of ion trajectories in first-principles electronic stopping power calculations. The approach resembles the one proposed by Gu et al.³¹ and considers the distribution of nearest-neighbor distances experienced by the ion along its path, which allows scrutiny of how representatively different trajectories sample the close collisions that determine core-electron contributions. Optimizing trajectories via this metric can help compute stopping powers more accurately and efficiently, and we expect that straightforward extensions of this metric to disordered systems at high temperatures and heterogeneous systems including compounds, alloys, and mixtures will be even more impactful.

We also identified cost-reducing strategies to systematically characterize finite-size errors. Our analysis of velocity- and trajectory-dependent finite-size effects revealed behavior inconsistent with the prevailing plasmon model³³, which we explain

through “ouroboros” effects wherein the projectile fictitiously interacts with previously excited electrons. Thus, we propose considering convergence with increasing distance between periodic images of the projectile’s path, rather than increasing supercell size alone. Accumulating results over multiple shorter trajectories can further reduce ouroboros effects and offer more information about the origin of finite-size errors, though transient behavior at the beginning of each simulation causes tradeoffs between trajectory length, finite-size errors, and computational cost. Finally, we showed that finite-size errors primarily influence valence-electron contributions to stopping power, enabling convergence without the high computational costs incurred by explicitly modeling core electrons in large supercells.

In summary, poor trajectory choices influence both core-electron contributions and finite-size errors, causing large errors of up to 60% in state-of-the-art calculations of proton stopping powers in aluminum. Similar magnitude effects can be expected in high-Z, low-density, and/or heterogeneous materials with highly non-uniform electron densities, where representatively sampling close encounters between the projectile and target ions is especially important. On the other hand, we expect particularly severe finite-size errors due to ouroboros effects for metallic systems, high-density materials, and/or trajectories closely aligned with simulation cell vectors because these situations involve factors that tend to strengthen artificial interactions between the projectile and its wake: delocalized excitations, smaller supercells, and/or close proximity between the projectile and previously disturbed material. Both close-collision sampling and ouroboros effects gain further importance for high-Z projectiles and velocities above the Bragg peak, regimes where core-electron processes and higher-energy excitations exacerbate sensitivities to projectile trajectory and simulation cell size.

Overall, our combination of approaches facilitates systematic control and analysis of two important approximations in first-principles stopping power calculations: the choice of projectile trajectory and finite supercells. These strategies will not only enhance the accuracy and efficiency of TDDFT stopping power calculations, but also prove valuable for higher levels of theory^{48,49} and quantum simulation algorithms on quantum computers^{50,51} as their viabilities improve.

METHODS

TDDFT simulations

The real-time TDDFT calculations were performed with an in-house extension^{39,41,42} of the Vienna ab initio simulation package (VASP)^{52–54}. Ground-state orbitals from density functional theory with a Fermi smearing of 100 K served as the initial condition for solving the time-dependent Kohn-Sham (KS) equations. Plane-wave cutoff energies of 750 eV achieved sufficiently converged results, and large supercells allowed reciprocal space sampling using the Γ point only. Supplementary Note 4 describes the convergence of these parameters in more detail. The adiabatic local density approximation (ALDA)^{55,56} was used for exchange and correlation (XC), and we verify the accuracy of ALDA relative to adiabatic versions of PBE⁵⁷ and SCAN⁵⁸ in Supplementary Note 6.

The electron-ion interaction was treated with the projector augmented-wave (PAW) method⁴⁰, explicitly including 3 or 11 valence electrons per aluminum atom. The two aluminum pseudopotentials allowed access to different contributions to the stopping power, with 3-electron calculations isolating the response of valence electrons, 11-electron calculations additionally including core contributions, and the difference between 11-electron and 3-electron results isolating core contributions. Supplementary Note 5 further discusses the influence of the pseudopotential approximation beyond the number of electrons explicitly modeled.

The TDDFT simulations held ionic velocities fixed while propagating the KS orbitals according to the Crank-Nicolson algorithm. For protons with velocities of 1.5 at. u. or more, the time step was chosen to scale inversely with proton velocity such that the proton traverses about 0.02 Å within each time step. For slower protons, smaller time steps of 0.3–0.4 as were needed to achieve converged results (see Supplementary Note 4 for more details).

Stopping power extraction

The ultimate stopping powers were computed from the time-dependent force on the proton, including Hellmann-Feynman and Pulay terms, but not the fully energy-conserving form ref. ⁵⁹, which is expected to be insignificant over the few-fs time scales simulated in this work. Data during the first 4 Å of the proton’s motion were excluded from the analysis to allow for dynamical ionization of the suddenly accelerated proton. The time-dependent force was integrated along the proton’s path to obtain the stopping work, or cumulative energy deposited by the proton into the electronic system. The slope of the least-squares linear fit of the stopping work then produced the average stopping power. In Supplementary Note 1, we compare different procedures for extracting an average stopping power from TDDFT data and show that the methodology described here converges to within 2% after the proton traverses 75 Å, whereas other procedures can be much more sensitive to the precise endpoint of the analysis and take longer to converge.

DATA AVAILABILITY

The datasets computed and analyzed during the current study are available from the corresponding author upon reasonable request.

CODE AVAILABILITY

The underlying code for the TDDFT simulations is not publicly available for proprietary reasons. The underlying code for the dielectric modeling is available on GitHub and Zenodo⁶⁰ and can be accessed via this link: <https://doi.org/10.5281/zenodo.7812518>.

Received: 20 July 2023; Accepted: 11 October 2023;

Published online: 30 October 2023

REFERENCES

- Kang, K. et al. Pushing the frontiers of modeling excited electronic states and dynamics to accelerate materials engineering and design. *Comp. Mater. Sci.* **160**, 207–216 (2019).
- Shepard, C., Zhou, R., Yost, D. C., Yao, Y. & Kanai, Y. Simulating electronic excitation and dynamics with real-time propagation approach to TDDFT within plane-wave pseudopotential formulation. *J. Chem. Phys.* **155**, 100901 (2021).
- Konoov, A. et al. Electron dynamics in extended systems within real-time time-dependent density-functional theory. *MRS Commun.* **12**, 1002–1014 (2022).
- White, A. J., Certik, O., Ding, Y. H., Hu, S. X. & Collins, L. A. Time-dependent orbital-free density functional theory for electronic stopping power: Comparison to the Mermin-Kohn-Sham theory at high temperatures. *Phys. Rev. B* **98**, 144302 (2018).
- White, A. J. & Collins, L. A. Fast and universal Kohn-Sham density functional theory algorithm for warm dense matter to hot dense plasma. *Phys. Rev. Lett.* **125**, 055002 (2020).
- White, A. J., Collins, L. A., Nichols, K. & Hu, S. X. Mixed stochastic-deterministic time-dependent density functional theory: application to stopping power of warm dense carbon. *J. Phys. Condens. Matter* **34**, 174001 (2022).
- Hentschel, T. W. et al. Improving dynamic collision frequencies: Impacts on dynamic structure factors and stopping powers in warm dense matter. *Phys. Plasmas* **30**, 062703 (2023).
- Parodi, K. & Polf, J. C. In vivo range verification in particle therapy. *Med. Phys.* **45**, e1036–e1050 (2018).
- Duzellier, S. Radiation effects on electronic devices in space. *Aerosp. Sci. Technol.* **9**, 93–99 (2005).

10. Allen, T., Busby, J., Meyer, M. & Petti, D. Materials challenges for nuclear systems. *Mater. Today* **13**, 14–23 (2010).
11. Hlawacek, G., Veligura, V., van Gastel, R. & Poelsema, B. Helium ion microscopy. *J. Vac. Sci. Technol. B* **32**, 020801 (2014).
12. Li, Z. & Chen, F. Ion beam modification of two-dimensional materials: characterization, properties, and applications. *Appl. Phys. Rev.* **4**, 011103 (2017).
13. Zylstra, A. B. & Hurricane, O. A. On alpha-particle transport in inertial fusion. *Phys. Plasmas* **26**, 062701 (2019).
14. Schleife, A., Kanai, Y. & Correa, A. A. Accurate atomistic first-principles calculations of electronic stopping. *Phys. Rev. B* **91**, 014306 (2015).
15. Lim, A. et al. Electron elevator: excitations across the band gap via a dynamical gap state. *Phys. Rev. Lett.* **116**, 043201 (2016).
16. Quashie, E. E., Saha, B. C. & Correa, A. A. Electronic band structure effects in the stopping of protons in copper. *Phys. Rev. B* **94**, 155403 (2016).
17. Yao, Y., Yost, D. C. & Kanai, Y. K-shell core-electron excitations in electronic stopping of protons in water from first principles. *Phys. Rev. Lett.* **123**, 066401 (2019).
18. Kononov, A. & Schleife, A. Pre-equilibrium stopping and charge capture in proton-irradiated aluminum sheets. *Phys. Rev. B* **102**, 165401 (2020).
19. Kononov, A. & Schleife, A. Anomalous stopping and charge transfer in proton-irradiated graphene. *Nano Lett.* **21**, 4816–4822 (2021).
20. Shepard, C., Yost, D. C. & Kanai, Y. Electronic excitation response of DNA to high-energy proton radiation in water. *Phys. Rev. Lett.* **130**, 118401 (2023).
21. Moro, M. V. et al. Traceable stopping cross sections of Al and Mo elemental targets for 0.9–3.6-MeV protons. *Phys. Rev. A* **93**, 022704 (2016).
22. Roth, D. et al. Electronic stopping of slow protons in transition and rare earth metals: breakdown of the free electron gas concept. *Phys. Rev. Lett.* **118**, 103401 (2017).
23. Malko, S. et al. Proton stopping measurements at low velocity in warm dense carbon. *Nat. Commun.* **13**, 2893 (2022).
24. Ullah, R., Artacho, E. & Correa, A. A. Core electrons in the electronic stopping of heavy ions. *Phys. Rev. Lett.* **121**, 116401 (2018).
25. Yost, D. C. & Kanai, Y. Electronic stopping for protons and α particles from first-principles electron dynamics: the case of silicon carbide. *Phys. Rev. B* **94**, 115107 (2016).
26. Yost, D. C. & Kanai, Y. Electronic excitation dynamics in DNA under proton and α -particle irradiation. *J. Am. Chem. Soc.* **141**, 5241–5251 (2019).
27. Li, S.-M. et al. First-principles study of the electronic stopping power of indium for protons and He ions. *Phys. Rev. B* **104**, 214104 (2021).
28. Maliyov, I., Crocombette, J.-P. & Bruneval, F. Electronic stopping power from time-dependent density-functional theory in gaussian basis. *Eur. Phys. J. B* **91**, 172 (2018).
29. Maliyov, I., Crocombette, J.-P. & Bruneval, F. Quantitative electronic stopping power from localized basis set. *Phys. Rev. B* **101**, 035136 (2020).
30. Ojanperä, A., Krasheninnikov, A. V. & Puska, M. Electronic stopping power from first-principles calculations with account for core electron excitations and projectile ionization. *Phys. Rev. B* **89**, 035120 (2014).
31. Gu, B. et al. Efficient ab initio calculation of electronic stopping in disordered systems via geometry pre-sampling: Application to liquid water. *J. Chem. Phys.* **153**, 034113 (2020).
32. Lee, C.-W., Stewart, J. A., Dingreville, R., Foiles, S. M. & Schleife, A. Multiscale simulations of electron and ion dynamics in self-irradiated silicon. *Phys. Rev. B* **102**, 024107 (2020).
33. Correa, A. A. Calculating electronic stopping power in materials from first principles. *Comp. Mater. Sci.* **150**, 291–303 (2018).
34. Ziegler, J. F., Ziegler, M. D. & Biersack, J. P. SRIM - the stopping and range of ions in matter (2010). *Nucl. Instrum. Meth. B* **268**, 1818–1823 (2010).
35. Lindhard, J. & Winther, A. Stopping power of electron gas and equipartition rule. *Matemat.-Fysis. Meddel.* **34**, 1–24 (1964).
36. Mermin, N. D. Lindhard dielectric function in the relaxation-time approximation. *Phys. Rev. B* **1**, 2362–2363 (1970).
37. Ramakrishna, K., Cangi, A., Dornheim, T., Baczewski, A. & Vorberger, J. First-principles modeling of plasmons in aluminum under ambient and extreme conditions. *Phys. Rev. B* **103**, 125118 (2021).
38. Peralta, J. P., Fiori, M., Mendez, A. M. P. & Montanari, C. C. Stopping-power calculations and the Levine-Mermin dielectric function for inner shells. *Phys. Rev. A* **105**, 062814 (2022).
39. Baczewski, A. D., Shulenburg, L., Desjarlais, M. P. & Magyar, R. J. Numerical implementation of time-dependent density functional theory for extended systems in extreme environments. Sandia report SAND2014-0597, Unlimited Release (2014).
40. Blöchl, P. E. Projector augmented-wave method. *Phys. Rev. B* **50**, 17953 (1994).
41. Baczewski, A., Shulenburg, L., Desjarlais, M., Hansen, S. & Magyar, R. X-ray thomson scattering in warm dense matter without the chihara decomposition. *Phys. Rev. Lett.* **116**, 115004 (2016).
42. Magyar, R. J., Shulenburg, L. & Baczewski, A. D. Stopping of deuterium in warm dense deuterium from ehrenfest time-dependent density functional theory. *Contrib. Plasm. Phys.* **56**, 459–466 (2016).
43. Vanderbilt, D. Optimally smooth norm-conserving pseudopotentials. *Phys. Rev. B* **32**, 8412–8415 (1985).
44. Troullier, N. & Martins, J. L. Efficient pseudopotentials for plane-wave calculations. *Phys. Rev. B* **43**, 1993–2006 (1991).
45. Schleife, A., Draeger, E. W., Kanai, Y. & Correa, A. A. Plane-wave pseudopotential implementation of explicit integrators for time-dependent Kohn-Sham equations in large-scale simulations. *J. Chem. Phys.* **137**, 22A546 (2012).
46. Draeger, E. W. et al. Massively parallel first-principles simulation of electron dynamics in materials. *J. Parallel Distr. Com.* **106**, 205–214 (2017).
47. Malko, S. et al. Importance of ion stopping power research for IFE. *IFE Science & Technology Community Strategic Planning Workshop*. <https://lasers.llnl.gov/content/assets/docs/nif-workshops/ife-workshop-2021/white-papers/malko-PPPL-IFE-workshop-2022.pdf> (2022).
48. Balzer, K., Schlünzen, N. & Bonitz, M. Stopping dynamics of ions passing through correlated honeycomb clusters. *Phys. Rev. B* **94**, 245118 (2016).
49. Moldabekov, Z. A., Dornheim, T., Bonitz, M. & Ramazanov, T. Ion energy-loss characteristics and friction in a free-electron gas at warm dense matter and nonideal dense plasma conditions. *Phys. Rev. E* **101**, 053203 (2020).
50. Babbush, R. et al. Quantum simulation of exact electron dynamics can be more efficient than classical mean-field methods. *Nat. Commun.* **14**, 4058 (2023).
51. Rubin, N. C. et al. Quantum computation of stopping power for inertial fusion target design. Preprint at <https://arxiv.org/abs/2308.12352> (2023).
52. Kresse, G. & Furthmüller, J. Efficient iterative schemes for ab initio total-energy calculations using a plane-wave basis set. *Phys. Rev. B* **54**, 11169 (1996).
53. Kresse, G. & Furthmüller, J. Efficiency of ab-initio total energy calculations for metals and semiconductors using a plane-wave basis set. *Comp. Mater. Sci.* **6**, 15–50 (1996).
54. Kresse, G. & Joubert, D. From ultrasoft pseudopotentials to the projector augmented-wave method. *Phys. Rev. B* **59**, 1758 (1999).
55. Zangwill, A. & Soven, P. Resonant photoemission in barium and cerium. *Phys. Rev. Lett.* **45**, 204–207 (1980).
56. Zangwill, A. & Soven, P. Resonant two-electron excitation in copper. *Phys. Rev. B* **24**, 4121–4127 (1980).
57. Perdew, J. P., Burke, K. & Ernzerhof, M. Generalized gradient approximation made simple. *Phys. Rev. Lett.* **77**, 3865–3868 (1996).
58. Sun, J., Ruzsinszky, A. & Perdew, J. P. Strongly constrained and appropriately normed semilocal density functional. *Phys. Rev. Lett.* **115**, 036402 (2015).
59. Ojanperä, A., Havu, V., Lehtovaara, L. & Puska, M. Nonadiabatic ehrenfest molecular dynamics within the projector augmented-wave method. *J. Chem. Phys.* **136**, 144103 (2012).
60. Hentschel, T. twhentschel/ueg-dielectric: UEG dielectric v0.2.0. <https://doi.org/10.5281/zenodo.8029770> (2023).

ACKNOWLEDGEMENTS

We thank André Schleife, Cheng-Wei Lee, Susan Atlas, Alexandra Olmstead, and Alexander White for helpful technical discussions. We are also grateful to Joel Stevenson for technical support, Raymond Clay for feedback on the manuscript, and Heath Hanshaw for pre-publication review. AK, ADB, and SBH were partially supported by the US Department of Energy Science Campaign 1. SBH and TWH were partially supported by the US Department of Energy, Office of Science Early Career Research Program, Office of Fusion Energy Sciences under Grant No. FWP-14-017426. All authors were partially supported by Sandia National Laboratories' Laboratory Directed Research and Development (LDRD) Project No. 218456. This article has been co-authored by employees of National Technology & Engineering Solutions of Sandia, LLC under Contract No. DE-NA0003525 with the U.S. Department of Energy (DOE). The authors own all right, title, and interests in and to the article and are solely responsible for its contents. The United States Government retains and the publisher, by accepting the article for publication, acknowledges that the United States Government retains a non-exclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this article or allow others to do so, for United States Government purposes. The DOE will provide public access to these results of federally sponsored research in accordance with the DOE Public Access Plan <https://www.energy.gov/downloads/doe-public-access-plan>.

AUTHOR CONTRIBUTIONS

AK performed TDDFT simulations, analyzed data, and drafted the manuscript. TWH contributed dielectric modeling for the analysis of finite-size errors. All authors interpreted the results, revised the manuscript, and approved the final version.

COMPETING INTERESTS

The authors declare no competing interests.

ADDITIONAL INFORMATION

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s41524-023-01157-7>.

Correspondence and requests for materials should be addressed to Alina Kononov or Andrew D. Baczewski.

Reprints and permission information is available at <http://www.nature.com/reprints>

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2023