Quantum computation of stopping power for inertial fusion target design

arXiv:2308.12352

Nicholas C. Rubin,[1,*] Dominic W. Berry,[2,†] Alina Kononov,[3] Fionn D. Malone,[1] Tanuj Khattar,[1]
Alec White,[4] Joonho Lee,[1,5] Hartmut Neven,[1] Ryan Babbush,[1,‡] and Andrew D. Baczewski[3,§]

[1] *Google Quantum AI, Venice, CA, USA*
[2] *School of Mathematical and Physical Sciences, Macquarie University, Sydney, NSW, AUS*
[3] *Quantum Algorithms and Applications Collaboratory, Sandia National Laboratories, Albuquerque, NM, USA*
[4] *Quantum Simulation Technologies Inc., Boston, MA, USA*
[5] *Department of Chemistry and Chemical Biology, Harvard University, Cambridge, MA, USA*

Sandia National Laboratories

Andrew Baczewski
IPAM
October 5, 2023

QuAAC

Setting the stage

What is stopping power?

The average kinetic energy lost by a particle moving through a medium.

What does it have to do with inertial fusion?

It appears in the source term in the energy balance that defines thermonuclear ignition.

How is it calculated classically?

There are many methods, but time-dependent density functional theory is the gold standard.

What type of quantum advantage should we expect?

Systematically improvable accuracy in quantum dynamics, at polynomial cost.

How does our quantum algorithmic protocol work?

Prepare a particle in a state with prescribed velocity, evolve in time, measure kinetic energy loss.

Stopping power: context & definition

Target

Density and composition are fixed

Starts in equilibrium state

Projectile

Bare charge and mass are fixed

Initial energy prescribed

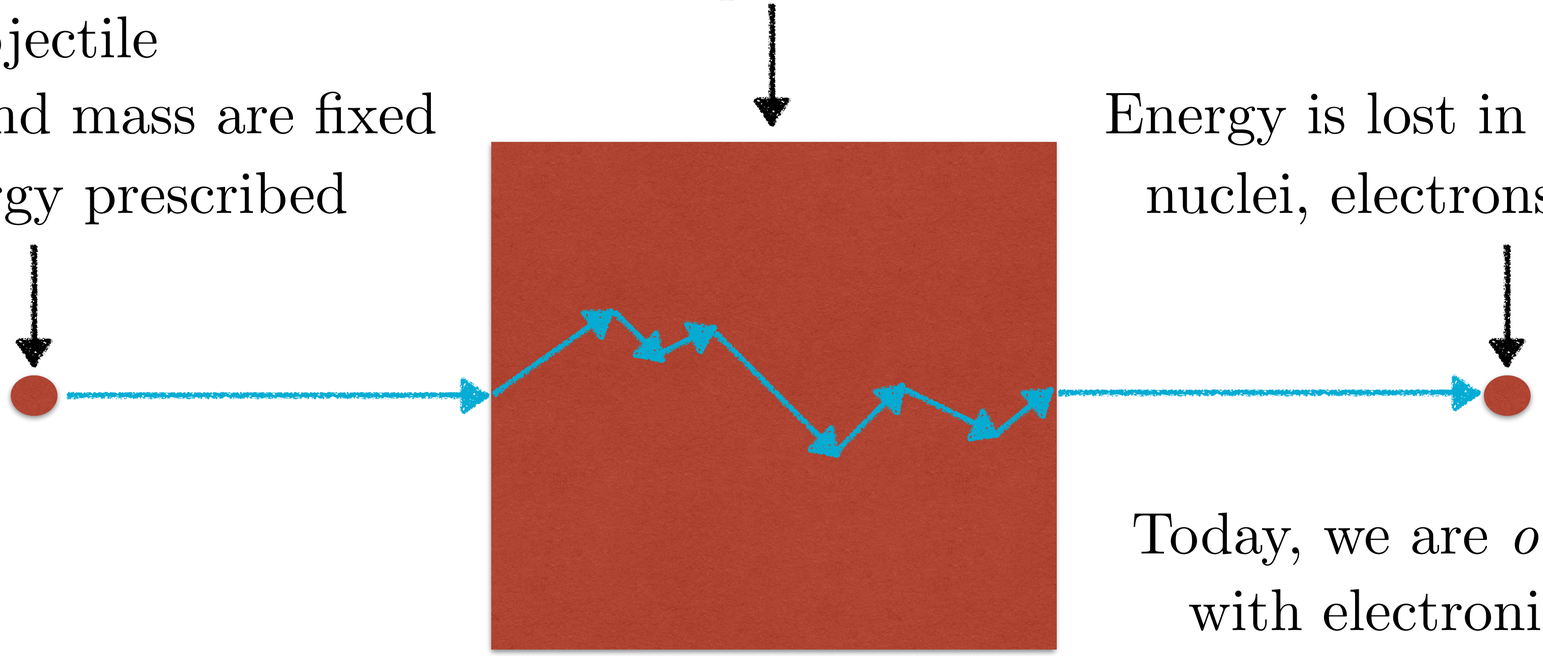

Energy is lost in collisions with nuclei, electrons, and light*.

Today, we are *only* concerned with electronic collisions.

Stopping power is the differential energy loss per unit length**, $S(E) = -\frac{dE}{dx}$

*I.e., bremsstrahlung

**This has units of force, not power.

Stopping from first principles

Simulate the dynamics of a projectile on an O(nm) length scale, O(as) time scale

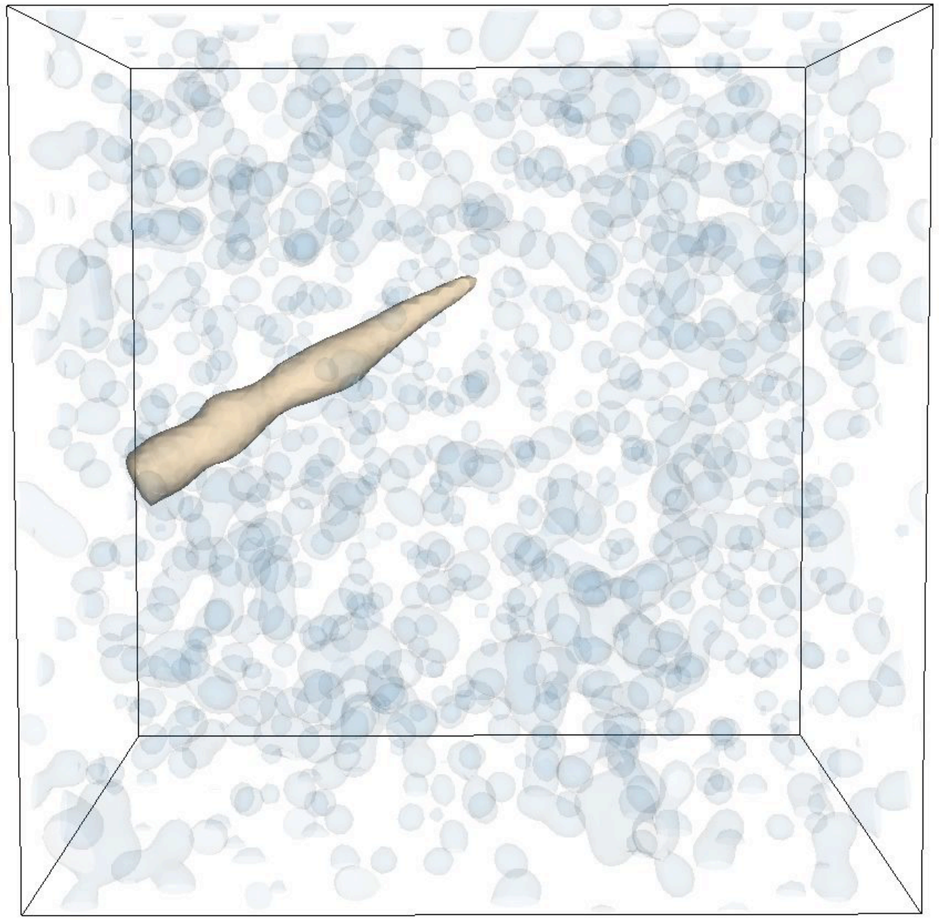

Proton stopping in warm dense deuterium

Governed by the same Hamiltonian that you've seen in chemistry/materials science

$$H_0 = \sum_{i=1}^{\eta} \frac{\nabla_i^2}{2} - \sum_{l=1}^{L} \sum_{i=1}^{\eta} \frac{\zeta_l}{||r_i - R_l||} + \sum_{i \neq j}^{\eta} \frac{1}{||r_i - r_j||}$$

$$H_{proj} = \frac{\nabla_{proj}^2}{2M_{proj}} - \sum_{i=1}^{\eta} \frac{\zeta_l}{||r_i - R_{proj}||}$$

Who cares about stopping powers?

Applications include radiation damage in space, nuclear reactors, charged-particle microscopy...

Cancer therapies

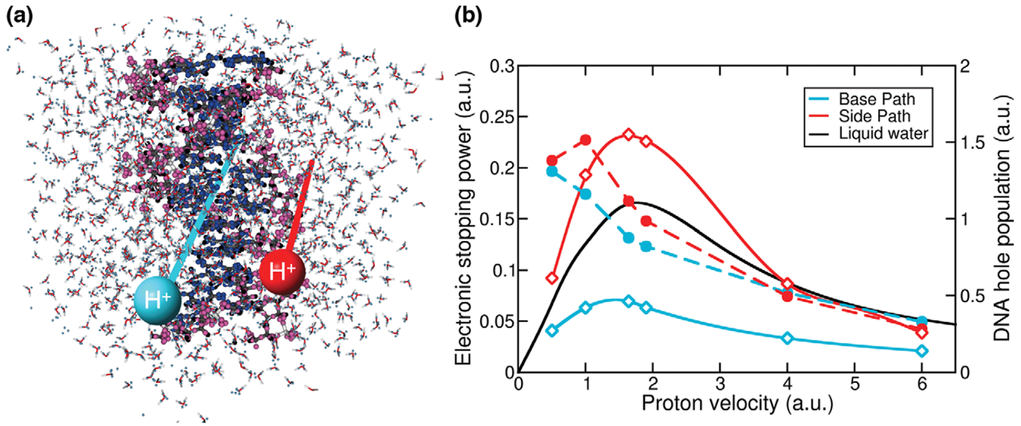

Shepard, Yost, and Kanai, PRL 130 (2023)

Nuclear spin qubits in silicon

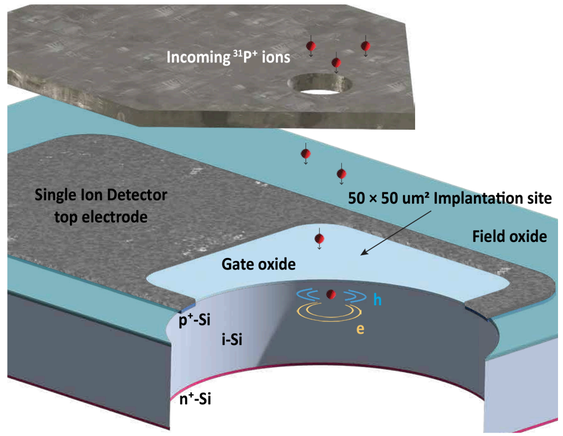

Jakob, et al., arXiv:2309.09626

It is *relatively* straightforward to measure stopping power.

One needs: (1) a charged particle source, (2) a few samples of the target*, (3) a spectrometer.

So, the value of a computational prediction is proportional to the cost of these three things.

This cost is rather high for one of the most important applications of stopping power: fusion.

*A key complication is the fact that the sample needs to be "differentially" thin.

Inertial confinement fusion

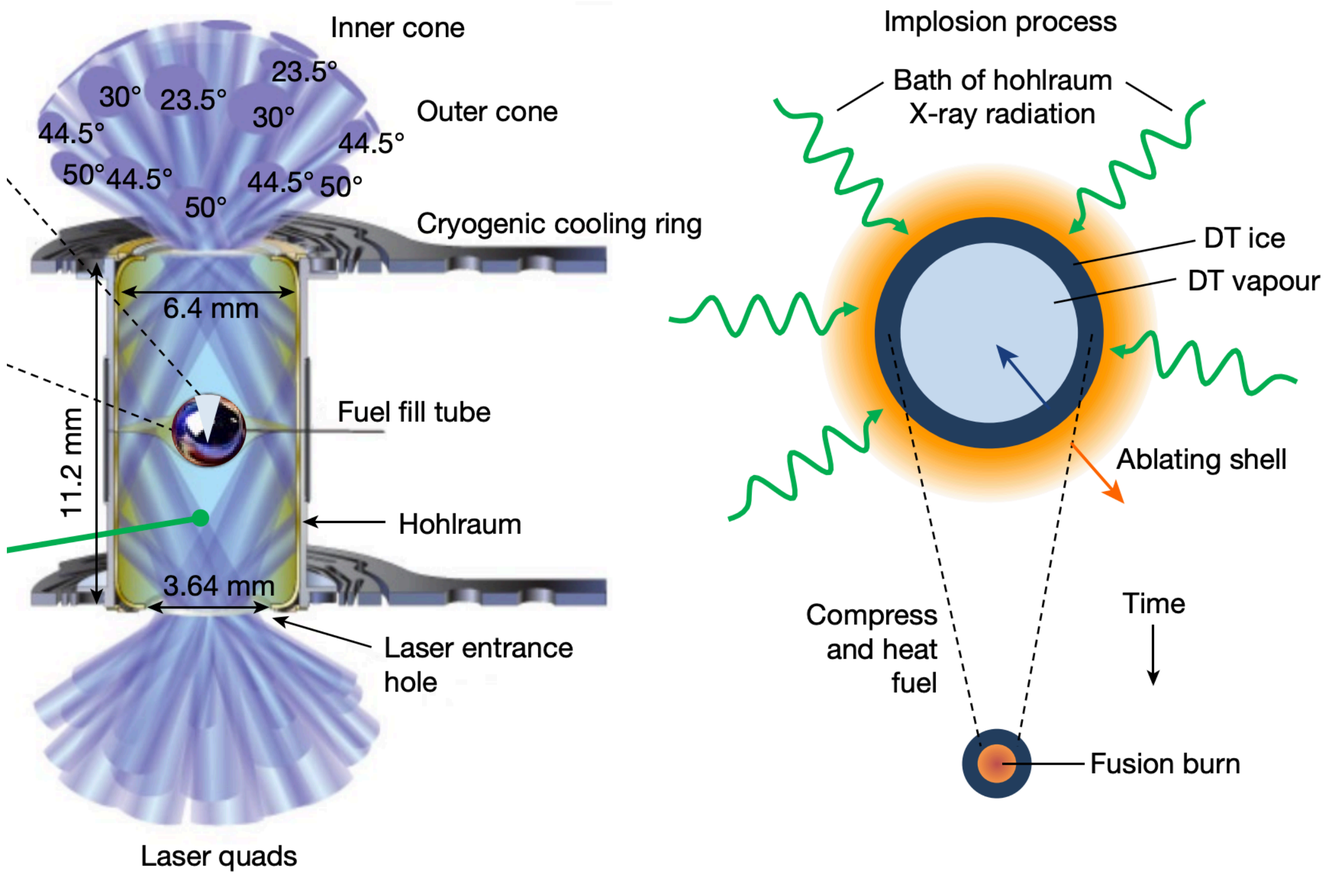

Zylstra, et al., Nature 601 (2022)

Ignition: fusion reactions are the dominant source of heat, exceeding losses.

This heating is due to stopping.

Fusion involves materials in extreme conditions

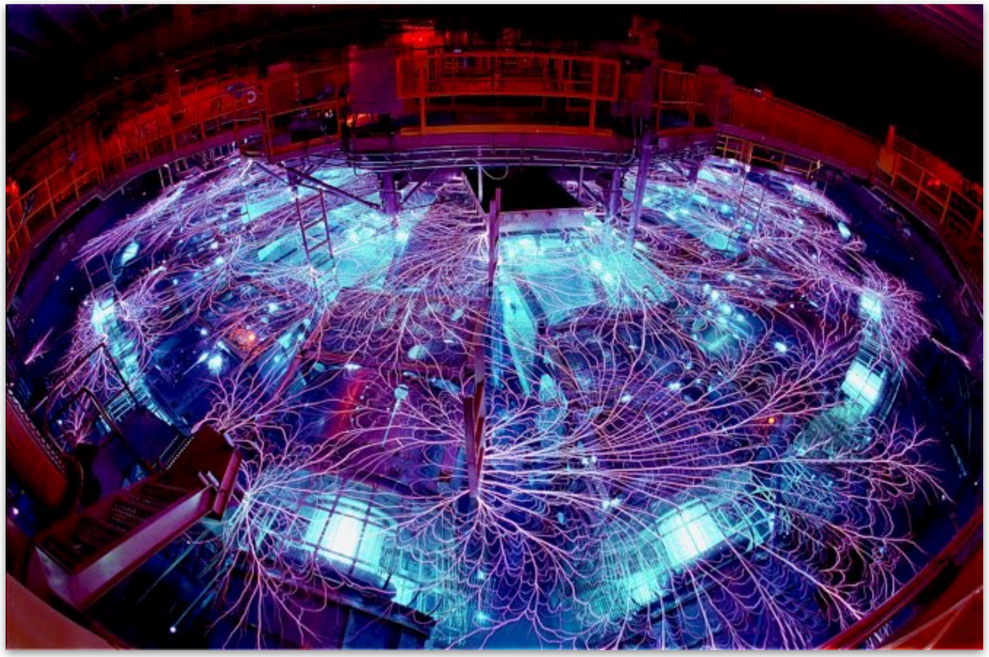

Photo credit: Randy Montoya

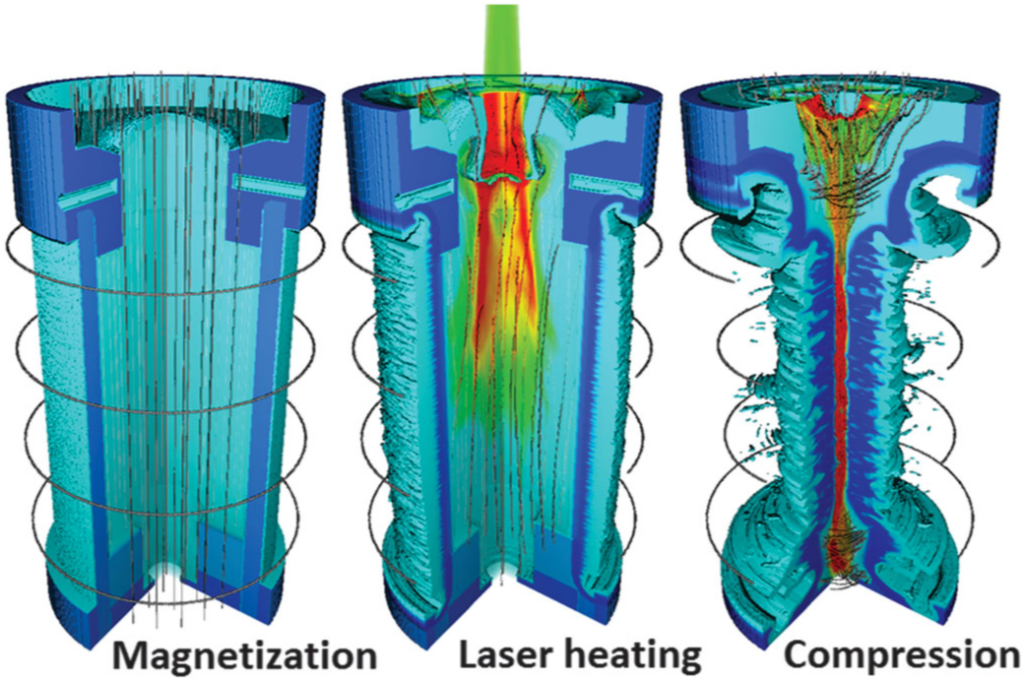

Gomez, et al., PRL 113 (2014)

Instabilities that develop on the way to ignition are part of what makes fusion hard.

This is more of a materials science problem than a nuclear physics problem.

It is compounded by the materials being in extremes of temperature/pressure.

What do I mean by "extreme"?

Fusion happens here...

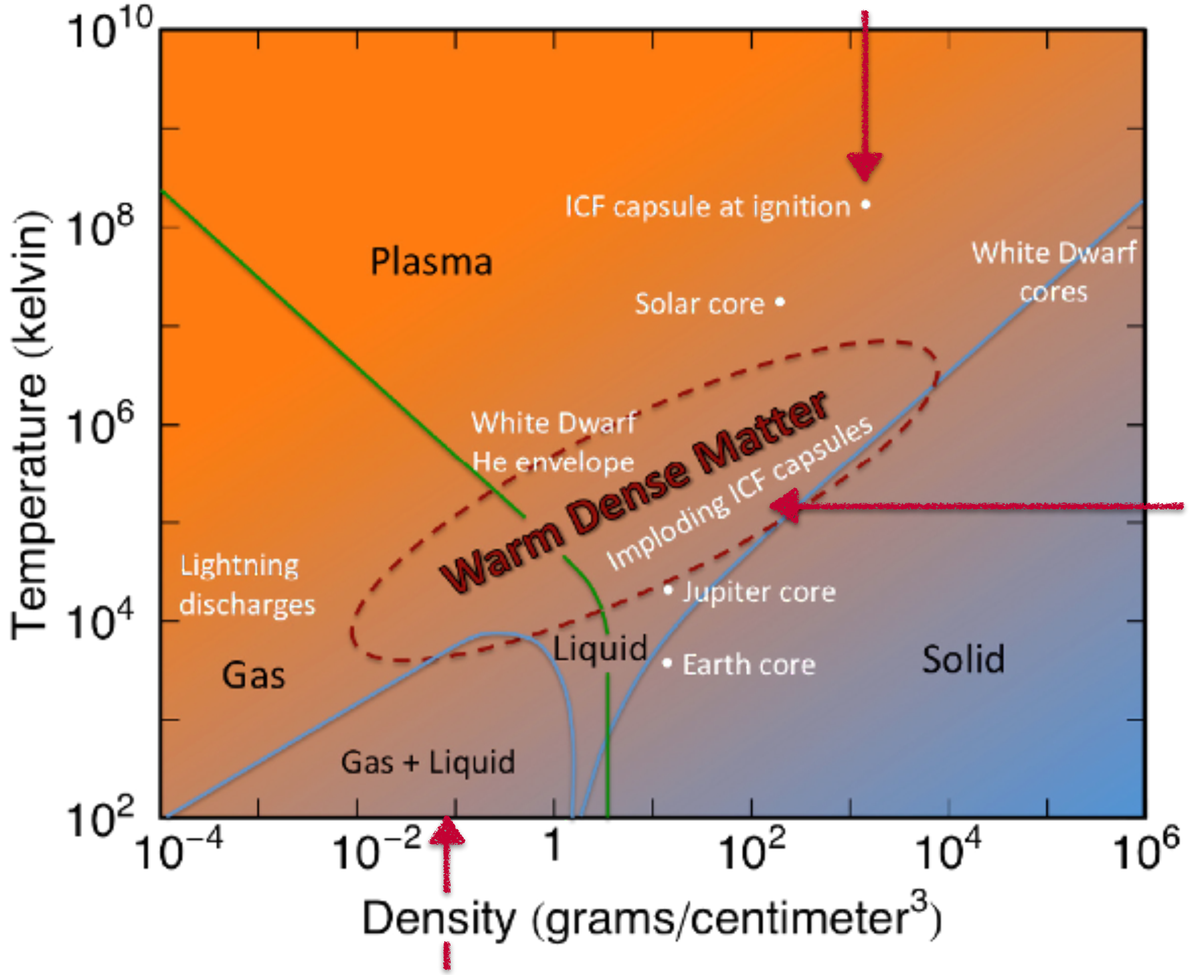

It is hard to constrain materials models here.

But the fuel starts here...

Figure credit: Mike Desjarlais

What do I mean by “extreme”?

We need to tabulate materials properties over a wide range of conditions to simulate experiments...

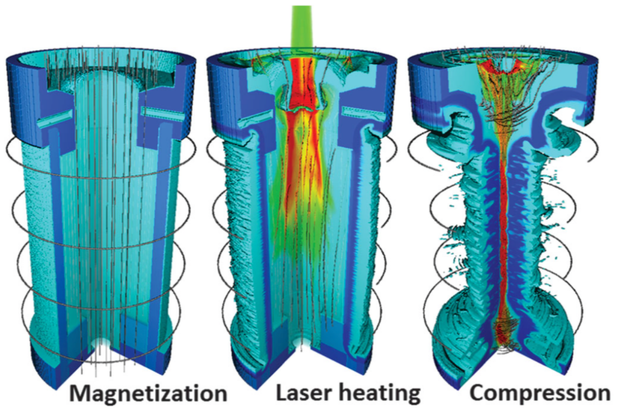

Fusion happens here...

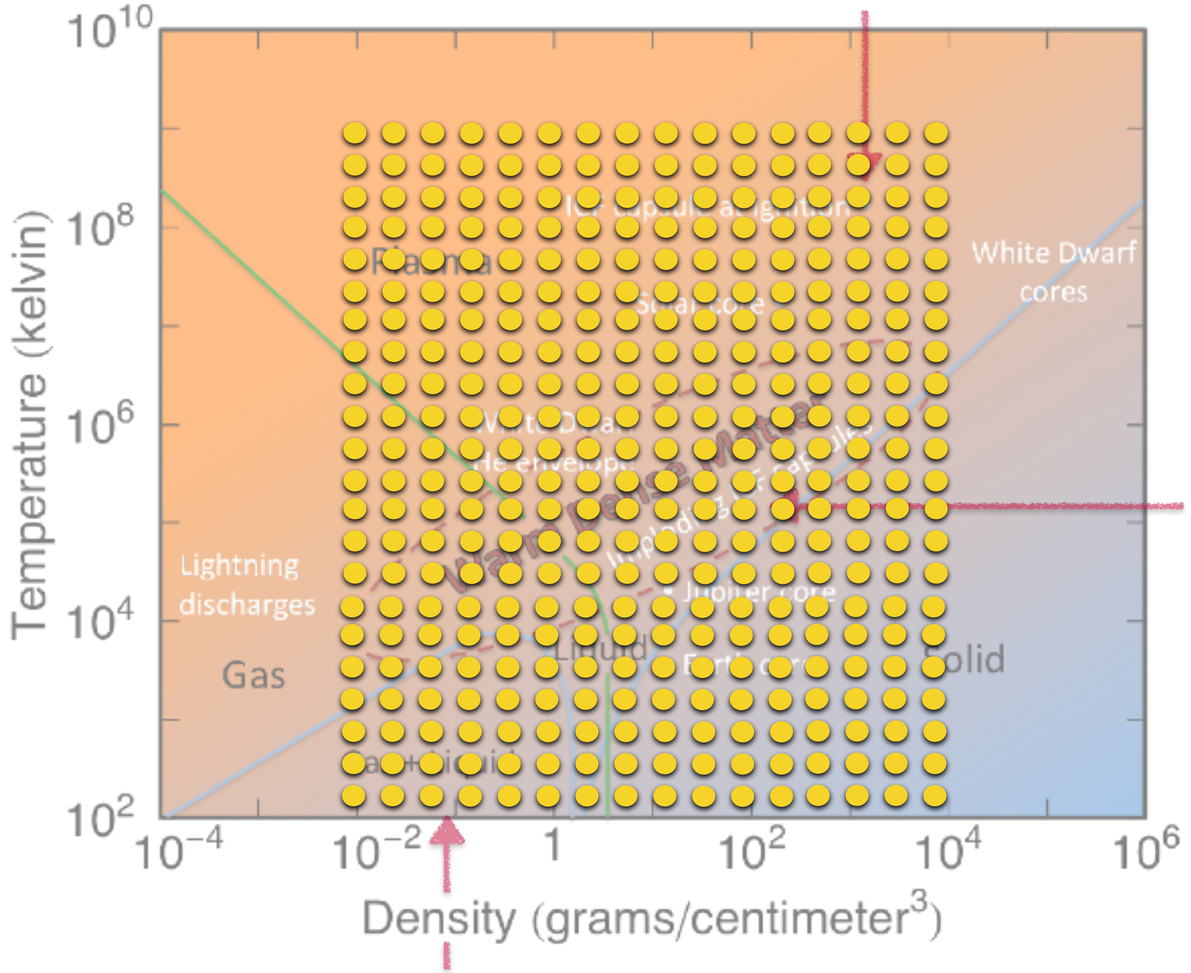

It is hard to constrain materials models here.

(Interesting basic science, too!)

But the fuel starts here...

Figure credit: Mike Desjarlais

Our best models are expensive...

Below, first principles calculations of stopping power of aluminum in the warm dense regime.

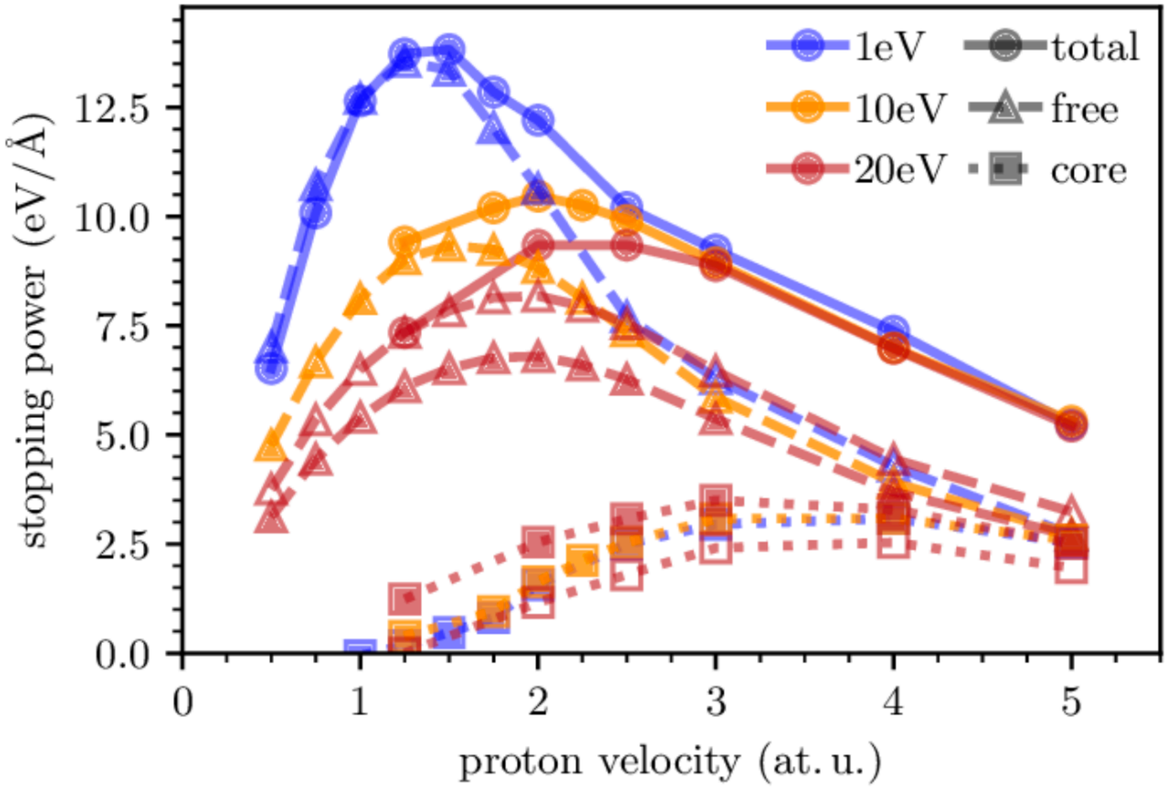

Kononov, et al., in prep (2023)

This plot took about 250 million CPU hours to produce, using a $170M machine.

We *do not* use these expensive models in hydrodynamic modeling.

We *do* use them to check the models that are/could be.

...but their accuracy is hard to assess.

We chose aluminum because it is an ideal system for benchmarking.

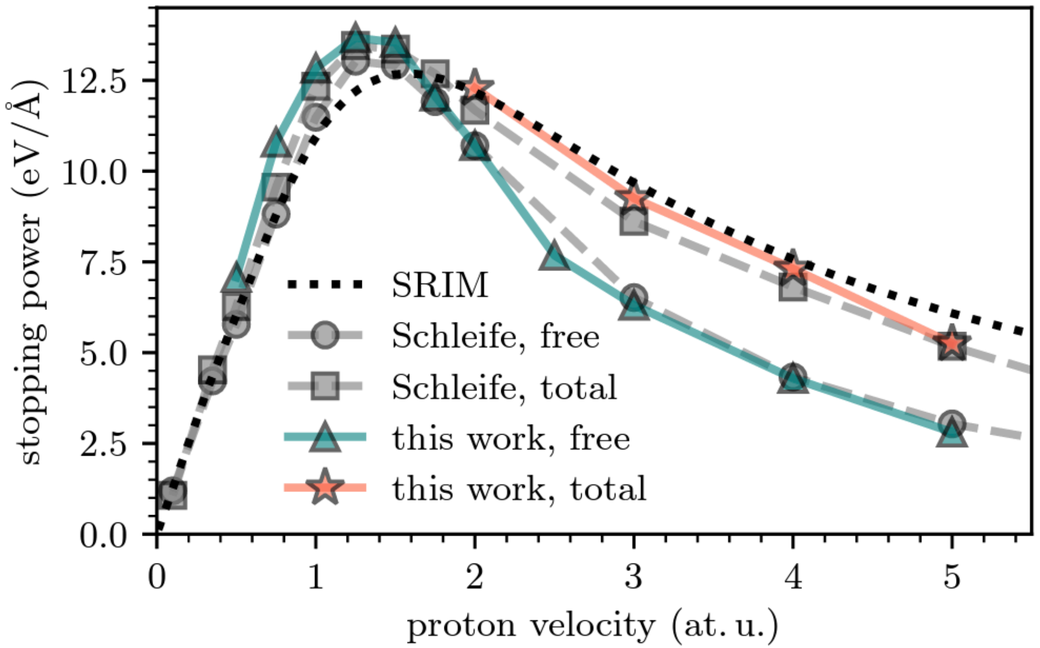

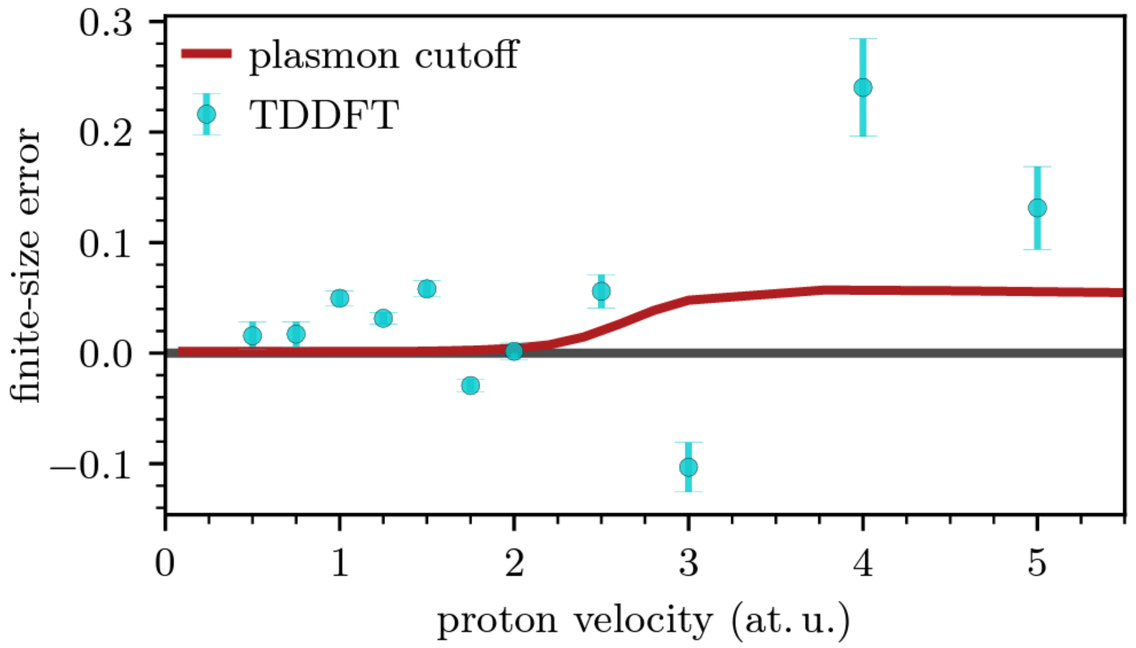

Kononov, et al., arXiv:2307.03213/accepted in npj Computational Materials

We can get results that agree well with experiment, for *some* experiments.

Even then, we don't have a systematic understanding of approximations.

Experimental data are *extremely* sparse in the warm dense matter regime.

First-principles stopping power calculations

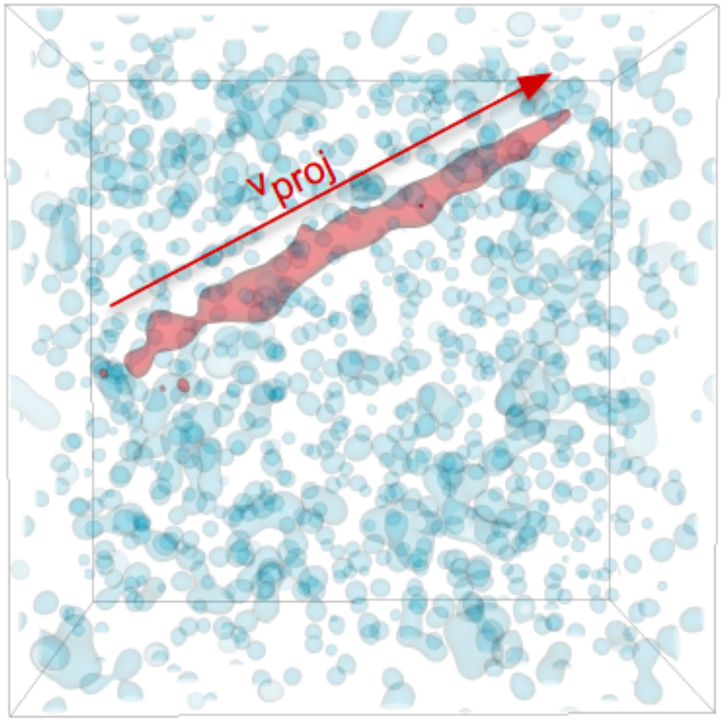

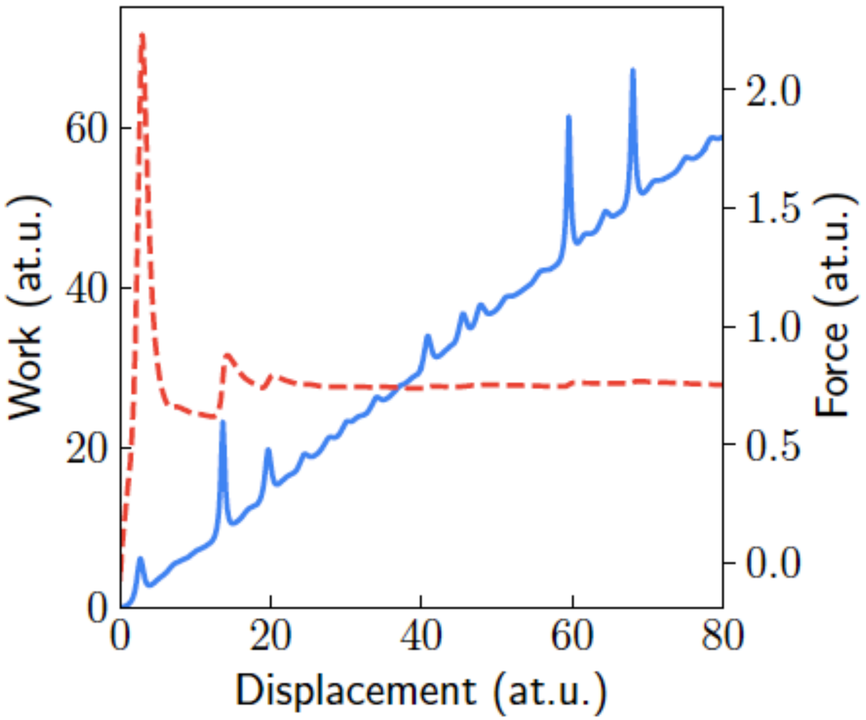

Create a representative supercell with 10s-1000s of atoms/electrons.

Push a projectile (red) through the target (blue) with some initial velocity (v_{proj}).

The energy loss of the projectile relates to an average force - the stopping power.

Time-dependent density functional theory

The time-dependent Kohn-Sham equations[1] govern the electronic dynamics,

$$i\frac{\partial}{\partial t}\phi_n(\mathbf{r},t) = \left(-\frac{\nabla^2}{2} + v_S\left[\rho\right](\mathbf{r},t)\right)\phi_n(\mathbf{r},t),\ \forall n \in \{1,\ldots,M\}$$

where the density[2] is given as

$$\rho(\mathbf{r},t) = \sum_{n=1}^{M} f_n(T_e)|\phi_n(\mathbf{r},t)|^2$$

The density-dependent one-body potential is defined as,

$$v_S\left[\rho\right](\mathbf{r},t) = v_{ext}(\mathbf{r},t) + v_H\left[\rho\right](\mathbf{r},t) + v_{xc}\left[\rho\right](\mathbf{r},t)$$

The exchange-correlation potential, the central approximation.

The *exact* potential[3] has many features that are extremely difficult to approximate.

[1] Runge and Gross, PRL 52 (1984), [2] Mermin PR 137 (1965), and [3] Elliott, et al. PRL 109 (2012)

Approximate classical vs. exact quantum

TDDFT has the form of a mean-field theory, like TD Hartree-Fock (TDHF).

In either, we're propagating the dynamics of a *single* Slater determinant.

TDHF is *fundamentally* limited in its accuracy, TDDFT is *practically* limited in its accuracy.

Mean-field theories are generally the least expensive and least accurate.

Table 1 | Costs of exact quantum algorithms and mean-field classical algorithms for simulating fermionic dynamics

Processor	Algorithm	Observable	Space	Gate complexity
Classical	$T = 0$ mean-field with occ-RI-K/ACE[23,24]	Anything	$\widetilde{\mathcal{O}}(N\eta)$	$(N^{4/3}\eta^{7/3}t + N^{5/3}\eta^{4/3}t)(\frac{Nt}{\epsilon})^{o(1)}$
Classical	$T > 0$ mean-field (density matrix) with refs. 23,24	Anything	$\widetilde{\mathcal{O}}(NM)$	$(N^{4/3}M^{2}\eta^{1/3}t + \frac{N^{5/3}M^{2}t}{\eta^{2/3}})(\frac{Nt}{\epsilon})^{o(1)}$
Classical	$T > 0$ mean-field (sampled trajectories) with refs. 23,24	Anything	$\widetilde{\mathcal{O}}(N\eta)$	$(\frac{N^{4/3}\eta^{7/3}t}{\epsilon^{2}} + \frac{N^{5/3}\eta^{4/3}t}{\epsilon^{2}})(\frac{Nt}{\epsilon})^{o(1)}$
Quantum	Second-quantized Trotter grid algorithm[45]	Sample $\lvert\psi(t)\rangle$	$\mathcal{O}(N \log N)$	$(N^{4/3}\eta^{1/3}t + \frac{N^{5/3}t}{\eta^{2/3}})(\frac{Nt}{\epsilon})^{o(1)}$
Quantum	First-quantized Trotter grid algorithm here	Sample $\lvert\psi(t)\rangle$	$\mathcal{O}(\eta \log N)$	$(N^{1/3}\eta^{7/3}t + N^{2/3}\eta^{4/3}t)(\frac{Nt}{\epsilon})^{o(1)}$
Quantum	Interaction picture plane wave algorithm[51]	Sample $\lvert\psi(t)\rangle$	$\mathcal{O}(\eta \log N)$	$\widetilde{\mathcal{O}}(N^{1/3}\eta^{8/3}t)$
Quantum	Grid basis algorithm from Appendix K of ref. 53	Sample $\lvert\psi(t)\rangle$	$\mathcal{O}(\eta \log N)$	$\widetilde{\mathcal{O}}(N^{1/3}\eta^{8/3}t)$
Quantum	New shadows procedure here	k-RDM(t)	$\mathcal{O}(\eta \log N)$	$\widetilde{\mathcal{O}}(k^{k}\eta^{k}L\,\mathcal{C}_{\mathrm{samp}}/\epsilon^{2})$
Quantum	Gradient measurement[61]	$\langle\psi(t)\lvert O\rvert\psi(t)\rangle$	$\widetilde{\mathcal{O}}(\eta + L)$	$\widetilde{\mathcal{O}}(\sqrt{L}\,\mathcal{C}_{\mathrm{samp}}\,\lambda/\epsilon)$
Quantum	Gradient measurement[61]	$\langle\psi(t)\lvert H\rvert\psi(t)\rangle$	$\widetilde{\mathcal{O}}(\eta + L)$	$\widetilde{\mathcal{O}}(\frac{\sqrt{L}\mathcal{C}_{\mathrm{samp}}t(N^{1/3}\eta^{5/3} + N^{2/3}\eta^{1/3})}{\epsilon})$

Babbush, et al., Nature Communications 14 (2023)

Quantum algorithms for dynamics can have better asymptotics than classical algorithms for mean-field theories.

Challenges for a quantum stopping protocol

Naive approach: Do everything that we would do in TDDFT, but using our favorite quantum dynamics algorithm.

Estimate the total energy for multiple evolution times?

Energy estimation is expensive...

Estimate the projectile force for multiple evolution times?

Force operator has a large norm/variance...

Updates to projectile position -> time-dependent simulation.

Solution: make the projectile an explicit quantum degree of freedom, estimate its kinetic energy at multiple evolution times.

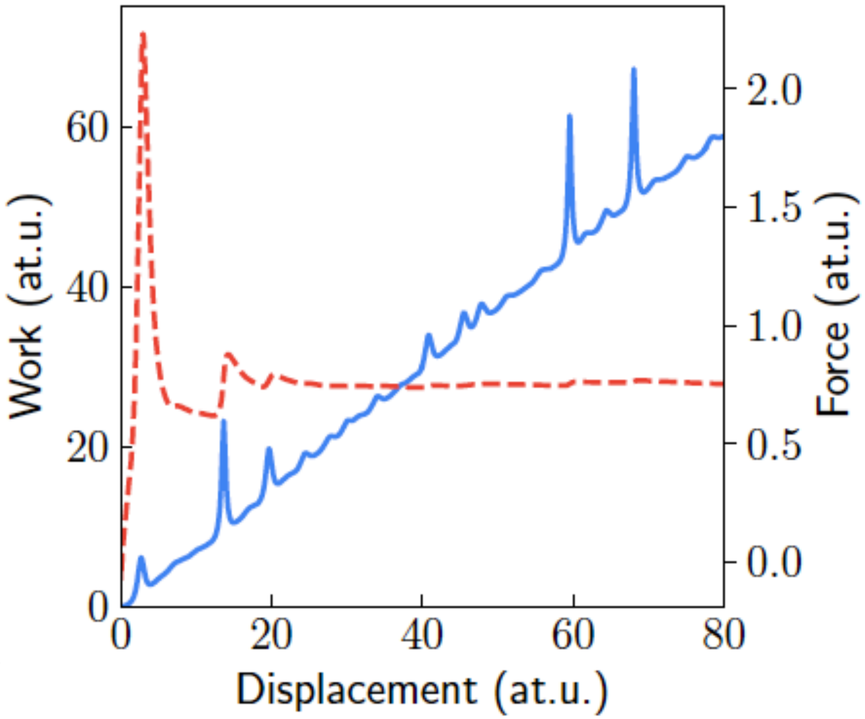

Note: we *don't* expect nuclear quantum effects in the projectile to be physically relevant. This choice is strictly in pursuit of algorithmic efficiency.

Summary of our protocol

Step 0: Choose a representation for the system (target + projectile).

Step 1: Prepare the initial state for the electron-projectile dynamics.

Step 2: Time evolution, using qubitization/interaction picture/Trotter.

Step 3: Measure the projectile's kinetic energy loss along its trajectory.

Step 4: Postprocess the sampled outcomes to estimate stopping power.

Step 0: Representing the system

The system is described in first quantization:

Basis: Babbush, et al., npj QI 5 (2019)
Block encodings: Su, et al., PRX Quantum 2 (2021)

Each quantum particle is described by a register of qubits that encode a plane-wave basis.

$$k_p = 2\pi p/\Omega^{1/3}$$

$$p \in G = \left[-(N^{1/3}-1)/2, (N^{1/3}-1)/2\right]^{\otimes 3}$$

Electrons are all treated as quantum, only the projectile nucleus is treated as quantum.

$$3\eta\lceil\log(N^{1/3})\rceil + 3\lceil\log(N_{proj}^{1/3})\rceil \approx 3(\eta+1)\lceil\log(N^{1/3})\rceil + 9$$

The non-projectile nuclei will not move appreciably over the fast timescale associated with electronic stopping. The projectile is moving about as fast as the electrons!

There are several non-trivial extensions of the block encodings in Su, et al., that account for details of incorporating the projectile into the typical electronic structure Hamiltonian.

Step 1: Prepare the initial state

The joint electron-projectile system is initialized as:

$$\exp(-\beta H_0)/\mathrm{Tr}\left[\exp(-\beta H_0)\right] \otimes |\psi_{\mathrm{proj}}(t=0)\rangle\langle\psi_{\mathrm{proj}}(t=0)|$$

Generically hard to prepare/sample from.

We use a mean-field initial state from Mermin-Kohn-Sham density functional theory.

Initial electronic state drawn from the canonical ensemble associated w/mean-field initial state— probabilities easy to compute.

Relies on efficiently preparing Slater determinants - Babbush, et al., Nature Communications 14 (2023).

Gaussian wave packet in momentum space, sharply peaked near $\mathrm{v}_{\mathrm{proj}}$.

The standard deviation of the wave packet is a free parameter.

Ultimately we can make it 10^4 times smaller than a physical proton!

Uses Bagherimehrab, et al., PRX Quantum 3 (2022).

Step 2: Time evolve the system

We considered time evolution using qubitization, the interaction picture, and Trotterization.

(Easy to bound) (Not competitive) (Hard to bound)

Qubitization counts: QSP to implement Jacobi-Anger[1,2,3], block encoding includes the projectile.

Trotterization counts: QROM interpolation[4] + Newton-Raphson[5] for inverse square root.

Significant numerical testing went into estimating constants for tighter Trotter bounds based on Low, et al., PRX Quantum (2023).

And the Trotter numerics still give worst-case (state-independent) bounds...

Note: Trotter estimates are for a real-space grid representation with comparable resolution.

[1] Low & Chuang, PRL 118 (2017), [2] Low & Chuang, Quantum 3 (2019), [3] Babbush, Berry & Neven, PRA 99 (2019),
[4] Sanders, et al., PRX Quantum 1 (2020), [5] Jones, et al., New J Phys (2012)

“Profiling” our protocol

Cirq-FT (Google tool) was used to build and profile a model of the entire protocol.

The dominant Toffoli cost (C4) is in a subroutine that involves applying controlled SWAPs to move each electron into a working register during SELECT.

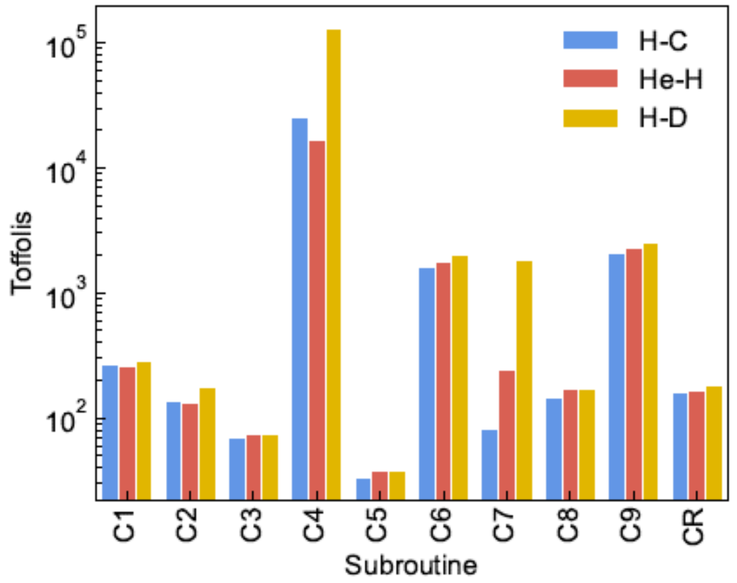

Step 3: Measure the kinetic energy

We compared estimates of the kinetic energy loss at the standard quantum limit (SQL) to a recent Heisenberg-scaling approach: Kothari & O'Donnell, SODA (2023)

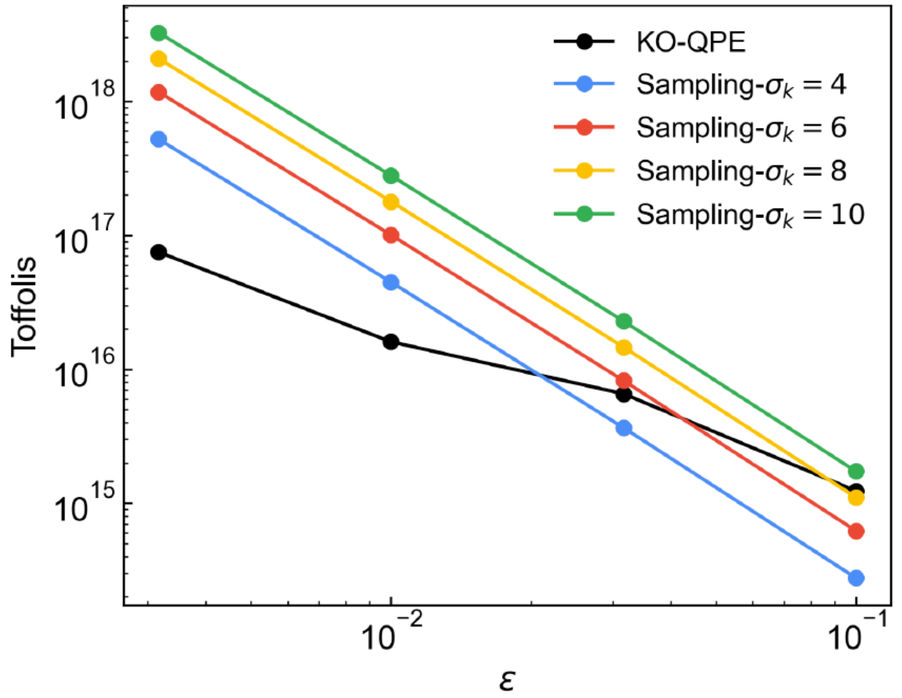

We find that the aggregate Toffoli count for the approach at the SQL is lower for the target precision. Higher precision might be required for non-fusion applications.

Step 4: Postprocess sampled outcomes

To assess how the sampling requirements impact estimates of the stopping power, we test our classical Monte Carlo estimate on a TDDFT trajectory w/fixed projectile variance*.

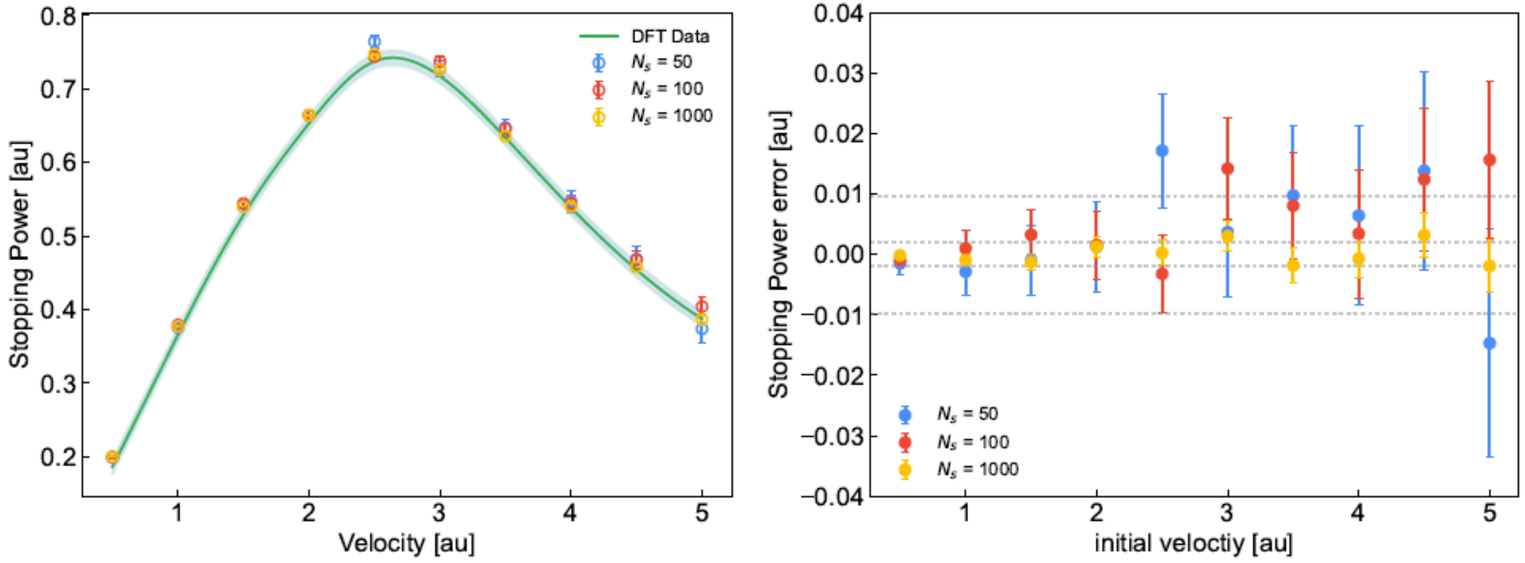

However, better accounting for variance due to sampling the thermal distribution over initial states might require more care.

*We do not expect the effective variance to change appreciably over the timescales of our simulations.

Resource estimates for fusion-relevant systems

Projectile + Host	η	QSP Toffoli	Product Formula Toffoli	QSP Qubits	Product Formula Qubits
Alpha + Hydrogen (50%)	28	5.593×10^{14}	1.124×10^{13}	1749	2666
Alpha + Hydrogen (75%)	92	2.033×10^{16}	3.069×10^{14}	3309	3902
Alpha + Hydrogen	218	1.992×10^{17}	1.399×10^{15}	5650	6170
Proton + Deuterium	1729	2.121×10^{20}	2.079×10^{17}	33038	33368
Proton + Carbon	391	2.225×10^{18}	1.074×10^{16}	8841	9284

For the smallest instance, the Toffoli count is ~100x the state-of-the-art for FeMoco[1].

For larger instances, counts are closer to FeMoco in 2016[2] - cause for optimism!

Strongly non-equilibrium dynamics of *even* 28 electrons in a large basis is classically challenging.

This would still tell us interesting things about how mean-field approximations fail...

Classical perspective: 10% of 40 PFlop/s for a week = 2.5 x 10^{21} floating point ops...

[1] Lee, PRX Quantum 2 (2021) and [2] Reiher, et al., PNAS 114 (2017)

Conclusions

I've shown you a quantum algorithmic protocol for estimating stopping powers in a classically challenging thermodynamic regime, relevant to inertial fusion.

The smallest "useful" instances have Toffoli counts that are 100x the state of the art for sampling the eigenspectra of industrially relevant molecules.

...but/and...

These are still among the first end-to-end estimates for implementing a practically relevant quantum dynamics calculation...

The competing classical resources are orders of magnitude larger than those typical of a lot of ground state quantum chemistry.

Questions? Comments? My email address is adbacze@sandia.gov.

Sandia National Laboratories

Sandia National Laboratories is a multi-missions laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for DOE's National Nuclear Security Administration under contract DE-NA0003525.

QuAAC

Our best experiments are also expensive

Measurement of Charged-Particle Stopping in Warm Dense Plasma

A. B. Zylstra, J. A. Frenje, P. E. Grabowski, C. K. Li, G. W. Collins, P. Fitzsimmons, S. Glenzer, F. Graziani, S. B. Hansen, S. X. Hu, M. Gatu Johnson, P. Keiter, H. Reynolds, J. R. Rygg, F. H. Séguin, and R. D. Petrasso

Phys. Rev. Lett. **114**, 215002 – Published 27 May 2015

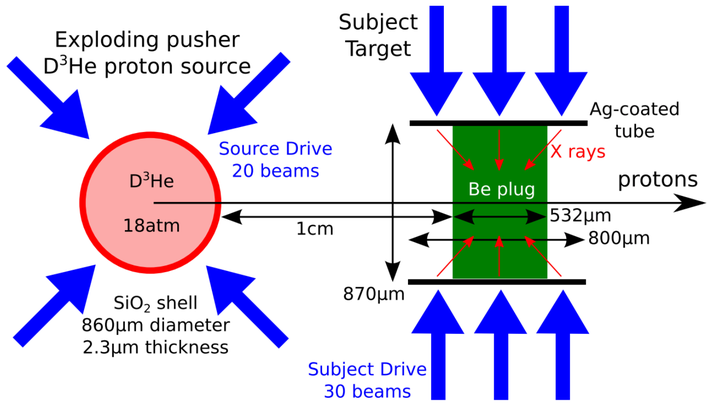

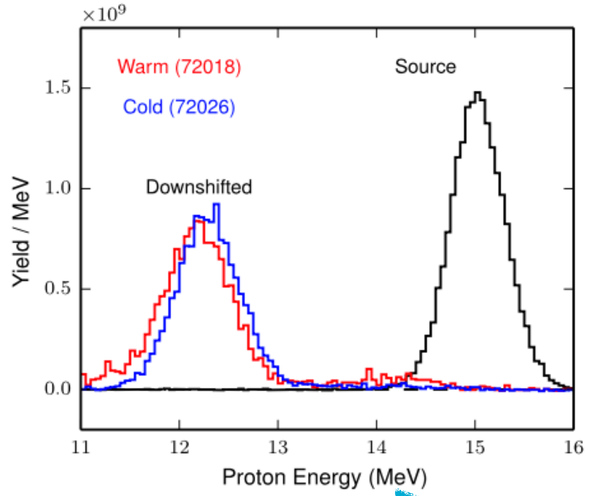

$$\mathrm{D} + {}^3\mathrm{He} \rightarrow {}^4\mathrm{He} + p(14.63\ \mathrm{MeV})$$

What is the energy-dependent force on a proton traversing an isochorically heated Be plasma near 32 eV?

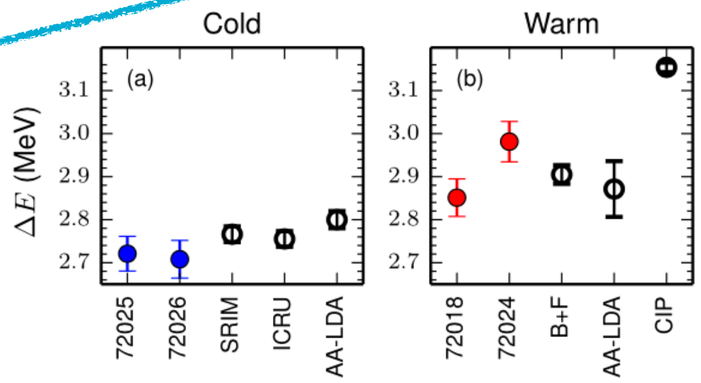

There are fewer than half a dozen data sets that constrain stopping in this regime.