

# Fermionic Quantum Turbulence: Pushing the Limits of High-Performance Computing

Gabriel Wlazłowski <sup>a,b,1</sup>, Michael M<sup>c</sup>Neil Forbes <sup>c,b,2</sup>, Saptarshi Rajan Sarkar <sup>c,3</sup>, Andreas Marek <sup>d,4</sup>, and Maciej Szpindler <sup>e,5</sup>

Received: 27 March 2024; Accepted: 29 March 2024; Advance access publication: 15 April 2024

Ultracold atoms provide a platform for analog quantum computer capable of simulating the quantum turbulence that underlies puzzling phenomena like pulsar glitches in rapidly spinning neutron stars. Unlike other platforms like liquid helium, ultracold atoms have a viable theoretical framework for dynamics, but simulations push the edge of current classical computers. We present the largest simulations of fermionic quantum turbulence to date and explain the computing technology needed, especially improvements in the Eigenvalue soLvers for Petaflop Applications (ELPA) library that enable us to diagonalize matrices of record size (millions by millions). We quantify how dissipation and thermalization proceed in fermionic quantum turbulence by using the internal structure of vortices as a new probe of the local effective temperature. All simulation data and source codes are made available to facilitate rapid scientific progress in the field of ultracold Fermi gases.

quantum turbulence | high-performance computing | ultra-cold gases | density functional theory

Computation is regarded as the third pillar of physical science, complementing theoretical and experimental physics. Each pillar has its unique methodology: theoretical physics relies on mathematical analysis, measurements are the central interest of experimental physics, and numerical modeling is the heart of computational physics. Many recent breakthroughs, like observing the Higgs boson (1, 2) or detecting gravitational waves (3), would not have been possible without advanced numerical analysis capabilities that adapt algorithmic breakthroughs to evolving hardware. Here we demonstrate the synergy between theory and computation: advances in linear algebra libraries enable Europe's fastest supercomputer (LUMI) to diagonalize matrices of record size, allowing us to simulate turbulent dynamics in quantum systems (superfluids). We use these simulations to investigate how vortices dissipate energy, driving quantum turbulence in neutron stars and ultracold-atom experiments.

As Moore's law bottoms out, using high-performance computing (HPC) effectively becomes a significant challenge. Current HPC systems consist of thousands of interconnected nodes, each comprising dozens of computing cores or multiple hardware accelerators. Specifically, accelerators like graphics processing units (GPUs) account for most of the computing power on modern platforms. Leadership supercomputers can compute from  $10^{17}$  floating point operations per second (FLOPS) for pre-exascale systems, to  $10^{18}$  FLOPS ( $=1$  EFLOPS or exa-FLOPS) for exascale systems. According to Top 500 list (June 2023) the top three supercomputers are: Frontier (Oak Ridge National Laboratory, USA) with 1.19 EFLOPS, Supercomputer Fugaku (RIKEN Center for Computational Science, Japan) with 0.44 EFLOPS, and LUMI (Euro HPC/CSC, Finland) with 0.31 EFLOPS. Here we use LUMI (Fig. 1), the fastest European system, to demonstrate some of its capabilities to advance computational physics.

While this computational potential is enormous, using these HPC capabilities requires a highly-tuned software stack capable of dealing with massive parallel and heterogeneous architectures, and core scientific libraries are constantly being adjusted to maximize performance on new hardware. These include Fast Fourier Transforms, linear algebra routines, libraries for matrix decomposition, random number generators, and solvers for algebraic and differential equations. These core libraries form the building blocks for the efficient domain-specific scientific packages that enable us to make physics breakthroughs in the domain of quantum mechanics.

Simulating quantum dynamics is one of the hardest challenges for classical computers due to the exponentially large size of a many-body wavefunction. Even storing the wavefunction for a modest nucleus like tin with  $\sim 100$  nucleons would require more bytes than there are atoms in the visible universe. The techniques of density functional theory (DFT) (4–6) and its time-dependent extension time-dependent density functional theory (TDDFT) (7, 8) have revolutionized our ability to study quantum dynamics by replacing the need to store the many-body wavefunction with an energy functional of a handful of densities. Despite needing to approximate the form of the functional, TDDFT has become one of the most successful methods for simulating dynamics in quantum systems. Its popularity is due to both its accuracy and

## Significance Statement

Accurate simulations of quantum systems are challenging for computational physics, yet essential for developing new technologies. The convergence of theory, algorithms, and supercomputers allows us to diagonalize million-by-million matrices and evolve millions of coupled partial differential equations to simulate complex phenomena in systems consisting of tens of thousands of superfluid fermions. Our results demonstrate key aspects of fermionic quantum turbulence that simpler models do not capture. We provide both data and open-source codes, establishing a benchmark for ultracold atom experiments to validate the theory. This framework will enable table-top quantum experiments to simulate complex dynamics, including compressible turbulence, and pulsar glitches in neutron stars that will enable astrophysical observations to constrain the properties of neutron-rich matter.

Author affiliations: <sup>a</sup>Warsaw University of Technology; <sup>b</sup>University of Washington, Seattle; <sup>c</sup>Washington State University, Pullman; <sup>d</sup>Max Planck Computing and Data Facility; <sup>e</sup>Academic Computer Centre CYFRONET, University of Science and Technology in Cracow

The work was conceived by G.W. and M.M.F. The TDSLDA calculations were done by G.W. The GPE calculations were done by S.S. Porting of the W-SLDA Toolkit to AMD was done by G.W., and M.S. Porting and optimizing of ELPA to AMD was done by A.M. Benchmarking ELPA on LUMI was done by A.M. Data analysis was done by G.W., S.S., and M.M.F. Support of computation processes on LUMI was done by M.S. All authors contributed to the manuscript writing.

The authors declare no competing interest.

<sup>1</sup>To whom correspondence may be addressed. Email: gabriel.wlazlowski@pw.edu.pl

versatility, providing access to static properties at zero and finite temperatures (6, 9), and real-time dynamics (8) for highly complex problems consisting of thousands of particles. For these reasons, DFT is presently the tool of choice for solid-state physics and quantum chemistry (10–12), and is used extensively in nuclear physics and astrophysics (13–15).

The discovery of high-temperature superconductivity (16) encouraged extensions of DFT to superconducting systems (17, 18). A local formulation was soon developed (19–21) called the superfluid local density approximation (SLDA), which has now been verified and validated against both quantum Monte Carlo (QMC) calculations and experiments at the few-percent level for a wide range of systems (21–23). These extensions have brought about new applications for DFT: designing superconductors (24), simulating nuclear reactions (25, 26), and benchmarking experiments with ultracold Fermi gases (27–34). Including superfluid correlations, however, comes with a cost that requires supercomputing resources.

To properly account for the Pauli exclusion principle, the SLDA variant of DFT is an orbital-based theory in the spirit of Kohn and Sham (5). When applied to a normal state, this requires one *orbital* per particle. Thus, to describe  $N_p$  particles on a spatial grid  $\mathbf{r}$  requires storing  $N_p$  functions. Extending this to superfluids requires at least twice as many functions, which we call quasiparticle wavefunctions (qpwf's)  $\varphi_n(\mathbf{r}, t) = [u_n(\mathbf{r}, t), v_n(\mathbf{r}, t)]^T$ , since we must represent both particles and holes. (Including both Hartree and Fock terms in nuclei requires four times as many functions (see e.g. (35)).) The main cost increase, however, is that superfluids allow fractional occupation of the orbitals, and to obtain convergence, one needs significantly more qpwf's than there are particles in the system. As we shall estimate below, if we represent our system on a three-dimensional spatial grid containing  $N^3$  points, then we need on the order of  $N_{\text{qpwf}} \approx 0.5N^3$  qpwf's leading to a memory cost of  $\approx 0.5N^6$  complex numbers per state. For a grid with  $N = 100$  points in each direction, a single state thus requires  $\approx 15$  TB.

We consider the simplest type of fermionic superfluid comprising equal populations of two species. (Think spin-up and spin-down, but in most cold-atom experiments these are different hyperfine states.) In these systems, the attractive interaction between these two states can be tuned at will using magnetic fields and broad Feshbach resonances to realize fermionic superfluids throughout what is called the Bardeen Cooper Schrieffer (BCS)-Bose-Einstein condensate (BEC) crossover (36). For weak inter-particle attractions, the superfluid has a BCS-like structure with large Cooper pairs providing long-range order and the associated superfluid flow while slightly modifying the dominant structure of the Fermi sea. As one increases the interaction strength, the size of these Cooper pairs gets smaller and smaller until they are best described as tightly bound dimers consisting of fermions of each type. In this BEC limit, the dimers behave as bosons, and the system is well described as a BEC with the Gross-Pitaevskii equation (GPE) Eq. (3) where  $m_B = 2m$  is now the mass of the dimer, and the dimer density  $n_B = n/2$  is half the total fermion density  $n$ . In the middle is the so-called unitary Fermi gas (UFG) where the dimers are on the threshold of being bound in the vacuum. Interestingly, there is no phase-transition separating the strongly and weakly interacting systems, hence the term crossover.

When these systems are sufficiently dilute, all the parameters of the short-range interaction can become irrelevant except for the s-wave scattering length  $a$  which describes the size of the bound dimers on the BEC side when  $a > 0$ . At unitarity  $a \rightarrow \infty$ , and it becomes negative on the BCS side where the dimers are unbound. It is thus convenient to parameterize the equation of state in terms of the dimensionless parameter  $k_F a$ , where  $k_F = (3\pi^2 n)^{1/3}$  is the so-called Fermi momentum. For reference, the total density  $n$ , Fermi energy  $\varepsilon_F(n)$ , and energy density



**Fig. 1.** LUMI, EURO HPC's pre-exascale system with AMD mi250x GPU-accelerated nodes. Each GPU node consists of the four AMD mi250x GPUS, each of which has two graphics compute dies (GCDS) being individual Heterogeneous-Compute Interface for Portability (HIP)-programmable devices. Photo copyright Fafe Creative.

$\mathcal{E}_{FG}(n)$  for this two-component free-Fermi gas are

$$n = 2 \int_{k < k_F} 1 \frac{d^3 k}{(2\pi)^3} = \frac{k_F^3}{3\pi^2}, \quad \varepsilon_F(n) = \frac{\hbar^2 k_F^2}{2m}, \quad [1a]$$

$$\mathcal{E}_{FG}(n) = 2 \int_{k < k_F} \varepsilon_F(k) \frac{d^3 k}{(2\pi)^3} = \frac{k_F^5}{10\pi^2 m} = \frac{\hbar^2 (3\pi^2)^{5/3}}{10m\pi^2} n^{5/3}. \quad [1b]$$

At zero temperature ( $T = 0$ ), for example, the equation of state throughout the crossover can be expressed as

$$\mathcal{E}_{k_F a}(n) = \xi(k_F a) \mathcal{E}_{FG}(n), \quad [2]$$

The dimensionless function  $\xi(k_F a)$  characterizes the strength of the interactions, and a major challenge for both theory and experiment has been to determine the universal value at unitarity known as the Bertsch parameter  $\xi(\infty) \approx 0.37$  (28), which combines both experiment (37, 38) and QMC (27, 39) results. (Note: computing the value of  $\xi(\infty)$  was a race between classical and analog quantum computing using cold atoms. Establishing the value was one of the first demonstrated successes of the quantum approach.) Dynamics in these systems has direct application to ultra-cold atom experiments, but also indirectly to nuclear physics. In particular, the neutron-neutron scattering length is accidentally large, making the UFG an excellent proxy for the dilute neutron matter expected to occur in the crust of neutron stars (40–42).

Simulating quantum turbulence is one of the most complex problem in quantum mechanics. The phenomenon has been studied intensively using both isotopes of superfluid helium – bosonic  $^4\text{He}$ , and fermionic  $^3\text{He}$  (see (43, 44) for reviews) – and more recently in ultracold atomic gases (see (45–50) and (51, 52) for reviews). The cold-atom platform has several advantages over liquid helium (23), including access to compressible and hypersonic regimes, superfluid mixtures for studying entrainment (53), and accurate microscopic theories in the form of the time-dependent superfluid local density approximation (TDSLDA) for fermions, and the GPE for bosons (see e.g. (54)). We demonstrate here that HPC techniques and time-dependent DFT frameworks have reached a level of maturity that allows for microscopic simulations of complex phenomena in systems consisting of tens of thousands of superfluid fermions.

## Results

Here we consider quantum turbulence in an ultra-cold atomic gas of fermions for two cases within BEC-BCS crossover:

1. The strongly coupled UFG with  $k_F a \rightarrow \infty$ . We compare this to a bosonic (GPE) theory for dimers tuned to the UFG equation of state;
2. A weakly coupled superfluid in the BCS regime with  $k_F a = -1.8$ .

Through this comparison, we demonstrate the importance of energy dissipation from heating, and its effect on the structure of quantum vortex cores in fermionic systems.

The quantum simulations we perform require two stages of computation: matrix diagonalization to obtain the initial state for evolution, and solving system of millions of coupled partial differential equations (PDES) to perform the real-time evolution. Here, we will describe them in the context of an HPC implementation on the LUMI supercomputer, and benchmark their efficiency on this platform. Before describing the technical aspects of the computation, we introduce a theoretical framework to show the source of the challenges we face.

**Theoretical Framework: TDDFT for Bosons.** We start with the simpler problem of describing a bosonic superfluid gas. If the gas is sufficiently dilute, then the superfluid state can be well described as a BEC where all the bosons (dimers in our case) occupy the same condensate wavefunction  $\psi_B(\mathbf{r}, t)$  normalized to the total boson number density  $n_B(\mathbf{r}, t) = |\psi_B(\mathbf{r}, t)|^2$ . This evolves under a non-linear Schrödinger equation called the GPE (see (55, 56)),

$$i\hbar e^{i\eta} \frac{\partial \psi_B(\mathbf{r}, t)}{\partial t} = \underbrace{\left( \frac{-\hbar^2 \nabla^2}{2m_B} + \mathcal{E}'(n_B(\mathbf{r}, t)) \right)}_{\hat{h}_B(\mathbf{r}, t)} \psi_B(\mathbf{r}, t), \quad [3]$$

where  $m_B$  is the boson mass,  $\eta$  is a small phase factor to model dissipation that we tune to better match the natural dissipation in the fermionic simulations (57), and interactions enter through the derivative of the equation of state  $\mathcal{E}(n_B)$  that characterizes the energy density as a function of the boson density  $n_B$ . This derivative is an effective mean-field chemical potential  $\mu = \mathcal{E}'(n_B)$  which repels or attracts bosons depending on the sign and strength of the interaction. One can include an external potential  $V_{\text{ext}}(\mathbf{r}, t)$  in the single-particle Hamiltonian  $\hat{h}_B(\mathbf{r}, t)$ , but we do not include one in our simulations here. Although expressed as a wavefunction, this is equivalent to an orbital-free DFT (58) of the Hohenberg-Kohn type (4), and Eq. (3) follows from a principle of stationary action for a generalized Schrödinger field  $\psi_B$  with the right-hand-side of Eq. (3) minimizing the energy functional

$$E_{\text{GPE}}[\psi_B] = \int \left( \frac{\hbar^2 |\nabla \psi_B(\mathbf{r}, t)|^2}{2m_B} + \mathcal{E}(|\psi_B(\mathbf{r}, t)|^2) \right) d^3r. \quad [4]$$

For modest system sizes, Eq. (3) can be efficiently solved on small computers, with HPC resources being required only for large simulation volumes (see e.g. (59)). These superfluids demonstrate a wide array of interesting properties, including dissipationless flow past obstacles (with  $\eta = 0$ ), and quantized vortices that mediate the energy cascades associated with quantum turbulence in spite of the lack of dissipation (44, 54).

Although the theory for fermions is much more complicated as we shall describe below, there is a limit which can be well described by a modified version of the GPE. This is the so-called BEC limit where two fermionic species have sufficiently strong attraction that they form a gas of tightly bound dimers. These dimers are bosonic in nature, and can be described by a modified GPE like Eq. (3) with  $m_B = 2m$ , total density  $n = 2n_B = 2|\psi|^2$  (i.e.  $\psi$  describes the dimers), total currents  $j = 2 \text{Im}(\psi \nabla \psi^*)$ , and a properly tuned equation of state  $\mathcal{E}$ . (See (58) for details: our modified GPE with  $\mathcal{E}(n) = \mathcal{E}_{k_{\text{fa}}=\infty}(n)$  is what they call the effective Thomas-Fermi ETF model.) To date, the majority of results for quantum turbulence in ultra-cold atomic gases have been simulated using the GPE (54).

**Theoretical Framework: TDDFT for Fermions.** Unlike bosons, fermions cannot occupy the same state due to the Pauli exclusion principle, and a density functional of the Hohenberg-Kohn type

would be highly non-local. Instead, one uses an orbital DFT of the Kohn-Sham type (5) where the functional is expressed in terms of a Slater determinant of  $N_{\text{qpwf}}$  single-particle orbitals  $\varphi_n(\mathbf{r}, t) = [u_n(\mathbf{r}, t), v_n(\mathbf{r}, t)]^T$ , which, as described above, must include at least two components to describe particle-hole excitations in a BCS type superfluid. Each of these states evolves under a single-particle Hamiltonian of the form:

$$i\hbar \frac{\partial}{\partial t} \begin{pmatrix} u_n(\mathbf{r}, t) \\ v_n(\mathbf{r}, t) \end{pmatrix} = \underbrace{\begin{pmatrix} \hat{h}(\mathbf{r}, t) & \Delta(\mathbf{r}, t) \\ \Delta^*(\mathbf{r}, t) & -\hat{h}^*(\mathbf{r}, t) \end{pmatrix}}_{\hat{\mathcal{H}}(\mathbf{r}, t)} \underbrace{\begin{pmatrix} u_n(\mathbf{r}, t) \\ v_n(\mathbf{r}, t) \end{pmatrix}}_{\varphi_n(\mathbf{r}, t)}, \quad [5]$$

where  $\hat{h}(\mathbf{r}, t)$  has a form similar to  $\hat{h}_B(\mathbf{r}, t)$  above with second-derivatives in space, and  $\Delta(\mathbf{r}, t)$  is a complex-valued function describing the superfluid correlations. Together they form the quasiparticle Hamiltonian  $\hat{\mathcal{H}}(\mathbf{r}, t)$ . The function  $\Delta$  plays the role of the order parameter, in analogy to the  $\psi_B$  function in the GPE. Unlike in the GPE, however, it no longer carries information about the density of the system. The key to a local DFT like the SLDA is that  $\hat{h}(\mathbf{r}, t)$  and  $\Delta(\mathbf{r}, t)$  depend only on a handful of local densities: the particle density  $n(\mathbf{r}, t)$ , the kinetic density  $\tau(\mathbf{r}, t)$ , the current density  $j(\mathbf{r}, t)$ , and the anomalous density  $v(\mathbf{r}, t)$ , each of which is computed from the orbitals  $\{\varphi_n(\mathbf{r}, t)\}$  via a reduction.

The precise form of equations of motion Eq. (5) will be discussed in the Methods and Materials section below, but follows from minimizing an energy functional of the form

$$E_{\text{SLDA}}[\{\varphi_n(\mathbf{r}, t)\}] = \int \mathcal{E}_{\text{SLDA}}(n(\mathbf{r}, t), \tau(\mathbf{r}, t), j(\mathbf{r}, t), v(\mathbf{r}, t)) d^3r. \quad [6]$$

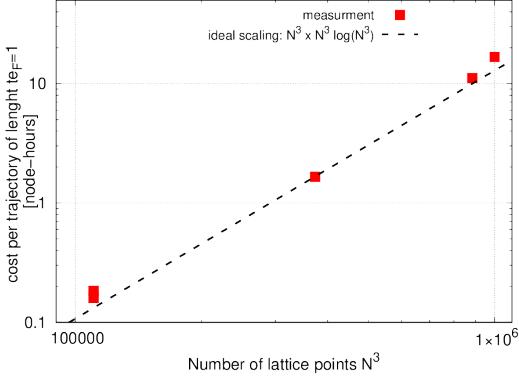
A key property of this system of PDES Eq. (5) is that the quasiparticle Hamiltonian  $\hat{\mathcal{H}}$  is the same for all quasiparticle wavefunctions. This means that at each step of evolution, one needs to communicate only the handful of local densities  $\sim 50$  MB rather than the complete state. Furthermore, the single-particle Hamiltonian is unitary, ensuring that the states remain orthonormal throughout the evolution. Thus, each node can independently perform the evolution of its quasiparticle wavefunctions using local hardware acceleration to compute the derivatives, with minimal communication that requires only an efficient message-passing interface (MPI) reduction. This is how we solve the second challenge, but to initialize this evolution, we must first obtain a good initial state.

This requires an orthonormal set of quasiparticle wavefunctions  $\{\varphi_n(\mathbf{r})\}$  which solve the self-consistent set of equations minimizing  $E_{\text{SLDA}}[\{\varphi_n\}]$ :

$$\begin{pmatrix} \hat{h}(\mathbf{r}) & \Delta(\mathbf{r}) \\ \Delta^*(\mathbf{r}) & -\hat{h}^*(\mathbf{r}) \end{pmatrix} \begin{pmatrix} u_n(\mathbf{r}) \\ v_n(\mathbf{r}) \end{pmatrix} = E_n \begin{pmatrix} u_n(\mathbf{r}) \\ v_n(\mathbf{r}) \end{pmatrix}. \quad [7]$$

While the matrix on the left-hand-side can be described efficiently in terms of the densities, the components are formally functions of the orbitals  $\hat{\varphi}_1, \hat{\varphi}_2, \dots$  and  $\Delta(\varphi_1, \varphi_2, \dots)$ . This Hermitian eigenvalue problem must be solved self-consistently, which we do iteratively through a series of diagonalizations. In the numerical implementation, the functions  $u_n(\mathbf{r})$  and  $v_n(\mathbf{r})$  are represented as vectors whose length depends strongly on the geometry and size of the problem. While some initial states can be computed efficiently – e.g. systems with high degrees of symmetry such as homogeneous matter – the maximum problem size is generally limited by the technical capabilities of the eigensolver libraries on the chosen HPC systems.

**Numerical setup and implementation.** To study turbulence, we simulate a periodic volume in space, and use a spectral representation for the quasiparticle wavefunctions on an equally-spaced  $N_x \times N_y \times N_z$



**Fig. 2.** Cost (time-to-solution  $\times$  node-count) of the parallel evolution of all quasiparticle wavefunctions for a unit time interval (expressed in dimensionless units), which corresponds to 286 integration time steps, as measured on the LUMI system. The number of nodes was adjusted to fit the problem in the memory available on the nodes, and ranges from 32 in the smallest case to 800 for the largest. The dashed line shows the expected ideal cost scaling.

Cartesian grid. In this basis, each quasiparticle wavefunction  $\varphi_n(\mathbf{r})$  is represented as a complex vector with  $2N_x N_y N_z$  components.

To put the size of our problem in perspective, we focus on a cubic box with  $N = N_{x,y,z} = 100$  grid points in each direction. We define our length scale in terms of the grid-spacing  $\Delta x = \Delta y = \Delta z = 1$  so that  $V = L^3 = N^3$  and set  $\hbar = m = 1$  so that momenta  $\mathbf{p} = \hbar \mathbf{k}$  is equivalent to the wave-vector. To compute the kinetic energy in  $\hat{h}$ , we use the fast Fourier transform (FFT)  $\tilde{\varphi}(\mathbf{k}) = \mathcal{F}(\varphi(\mathbf{r}))$  where the momenta  $\mathbf{k} = 2\pi n/N$  for  $n = \{-N/2, \dots, N/2 - 1\}$  increase in steps of  $\Delta \mathbf{k} = 2\pi/N$ :

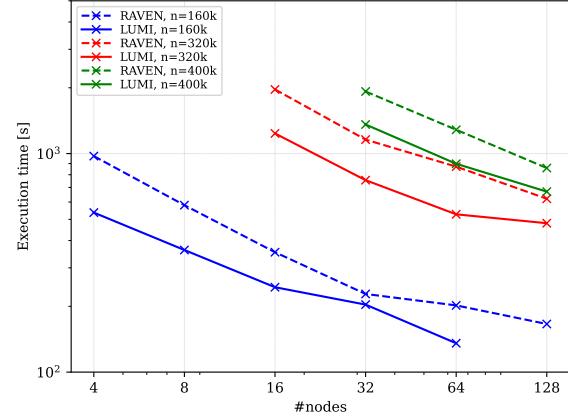
$$-\frac{1}{2} \nabla^2 \varphi_n = \mathcal{F}^{-1} \left( \frac{k^2}{2} \mathcal{F}(\varphi_n) \right). \quad [8]$$

This replaces a matrix multiplication with two FFTs and a single intermediate vector multiplication by  $k^2/2$  which is diagonal in momentum space. The remaining calculations are local in position space, simply multiplying  $\varphi_n$  by various functions of the densities.

These operations can be computed independently and locally on the computation nodes, each of which stores a small fraction of the total set of quasiparticle wavefunctions  $\varphi_n$ . The need for HPC comes from the large number of these required to adequately represent the problem. To estimate this, note that, in our units, the maximum momentum represented is  $k_{\max} = \pi$ , hence the maximum kinetic energy represented is  $E_{\max} = k_{\max}^2/2 = \pi^2/2$ . This provides a natural cutoff scale  $E_c \lesssim E_{\max}$  and we must keep those quasiparticles with energy  $E_n < E_c$ . For large  $E_n$ , the energy is dominated by the kinetic energy, and we can estimate the number of such states by considering the volume  $\frac{4}{3}\pi k_{\max}^3$  of the sphere  $E < E_c$  in momentum space in terms of the volume occupied by each quantum state ( $\Delta k$ ):

$$N_{\text{qpwf}} \approx \frac{\frac{4}{3}\pi \overbrace{\pi^3}^{k_{\max}^3}}{\underbrace{(2\pi/N)^3}_{(\Delta k)^3}} \approx 0.5N^3. \quad [9]$$

For each state we have two PDES (c.f. Eq. (5)). This demonstrates the cost of explicitly including the Pauli principle: instead of solving one PDE as for bosons, we need to solve many PDES. For the case we consider,  $N = 100$ , we will be solving in parallel the corresponding million of PDES, all of them coupled to each other!



**Fig. 3.** Strong scaling of the ELPA 1-stage solver. Shown is the time-to-solution of the ELPA 1-stage solver for real, double-precision computations in a strong-scaling setup for different matrix sizes. The solid lines show the results on LUMI with 4 AMD MI250X GPUs per node. The dashed lines show the results obtained on the RAVEN HPC system of the Max Planck Computing and Data Facility (MPCDF), with 4 NVIDIA A100 GPUs per node.

**Performance of the Time Evolution Algorithm.** The time integration is done with 5<sup>th</sup>-order multistep Adams-Basforth-Moulton (ABM) predictor-corrector method (60). The method requires the evaluation of  $\hat{h}$  and  $\Delta$  twice per time step (predictor and corrector). The cost per time step thus scales as  $N^3 \times N^3 \log N^3$  where the first factor accounts for the number of evolved states and the second one for the complexity of the FFT that we use to compute the kinetic term. (We use the `hipfft` library [hipfft.readthedocs.io](https://hipfft.readthedocs.io).) The stability of the method has been studied in (61, 62).

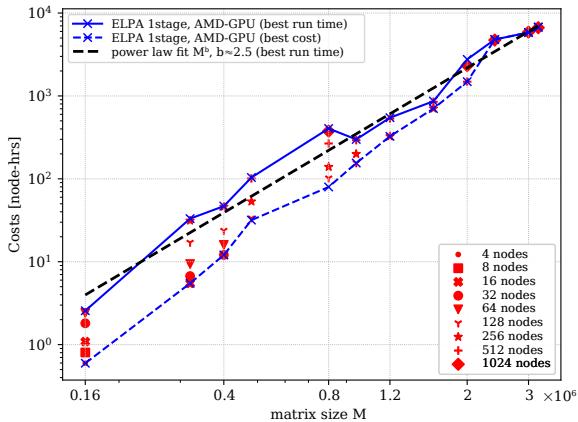
The open-source `w-SLDA` Toolkit (63) provides a parallel implementation of this time integrator, and is designed to simulate fermionic superfluids with the `TDSLDA` on modern GPU-accelerated systems. In Fig. 2, we demonstrate the measured cost  $C$  (defined as time-to-solution  $\times$  node-count) obtained on the LUMI system. Our parallel implementation of Eq. (5) exhibits the expected scaling up to the maximum problem size of  $N = 100$  corresponding to one million PDES. The main limitation is imposed by the memory requirements: the ABM method, while very accurate, require about 10 copies of the state to operate in case of our implementation. In our largest case ( $N = 100$ ) the total memory requirement is about 164 TB.

**Performance of the Matrix Diagonalization Algorithm.** Generating the initial state is even more costly, as it requires finding a self-consistent solution to Eq. (7) for the complete set of  $N_{\text{qpwf}}$  wavefunctions  $\{\varphi_n\}$ . This is done iteratively via a sequence of diagonalizations of an  $M \times M$  Hermitian matrix where  $M = 2N^3$ . We reduce the cost of iteratively solving Eq. (7) by using a multi-grid approach. We fix the domain size and solve the problem on consecutively larger lattices with  $N = 60, 80$ , and 100 points in each direction, corresponding to decreasing lattice spacings  $\Delta x = 1.67, 1.25$ , and 1.0, respectively. At each step, we interpolate the converged solution from coarser to finer lattices, providing a good initial state to accelerate the iterative algorithm. In this way, we only need a few iterations to converge on the target  $N = 100$  lattice. Even with this tremendous simplification, in the final stage of iterations requires a few dense diagonalization of two million by two million matrices.

To do this, we use the publicly available `ELPA` library (64), which was designed to efficiently solve dense symmetric or Hermitian standard or generalized eigenvalue problems (EVPS), especially with scalability to large core and/or GPU counts in mind. The `ELPA` library was first released in 2010 and has been ported and optimized for all major HPC architectures. The `ELPA` library is used in most software packages for

**Table 1. Physical parameters in the simulations**

Length scales [ $l_F$ ]:		min	core	sep.	max	method
$N$	$k_F a$	$\Delta x$	$\xi$	$l(0)$	$L_{\text{box}}$	$L(0)$
22 803	$\infty$	1.0	1.3	7.2	99	19 200
26 790	$\infty$	1.0	1.3	7.2	99	19 200
108 532	-1.8	1.5	2.6	14.5	150	15 937
						tdslda



**Fig. 4.** Scaling of the ELPA 1-stage solver with matrix size. Shown is the cost for solution of different eigenvalue problems for different matrix sizes with the ELPA 1-stage solver on the LUMI GPU system. The different red symbols indicate the costs (see also Fig. 3) for a fixed matrix size but different node counts. The solid blue line represents the fastest time-to-solution (albeit with the highest costs) achieved. Note, that the blue line was obtained by measurements on different number of GPU nodes. Further note that the blue line has been obtained by taking the arithmetic mean of the costs of at least two experiments of a specific matrix size and node count. The dashed black line shows a power-law fit to the costs for the best run time. The dashed blue line represents the lowest cost (albeit with higher time-to-solution) achieved.

electronic-structure theory\* and clearly outperforms implementations such as SCALAPACK (65). Recently, in addition to the accelerated version for NVIDIA GPUs (66, 67), a port to the AMD MI250 GPU architecture has been publicly released. Here we show the first results obtained with this port for AMD GPUs.

For the standard EVP, the ELPA library provides two solvers. The first is a 1-stage solver with three steps: i) transforming the dense matrix into a tridiagonal form, ii) diagonalizing this tridiagonal form, and iii) transforming the eigenvectors back to their original representation. Alternatively, a 2-stage solver introduces two additional steps: first transforming the dense matrix in a banded matrix, and then transforming the banded matrix into a tridiagonal form; the subsequent back transformation of the eigenvectors also requires two stages. For specific algorithmic details, see (64, 68).

To initialize the target  $N = 100$  quantum turbulence problem, we must be able to efficiently diagonalize matrices where the size is of orders of millions by millions. With access to the LUMI supercomputer, we first ensured that the new ELPA AMD GPU version works as expected on large node counts with matrices of this size. In Fig. 3 we compare the run time of the AMD GPU version of the ELPA library with the run time on NVIDIA A100 GPUs using the RAVEN system [www.mpcdf.mpg.de/services/supercomputing/raven](http://www.mpcdf.mpg.de/services/supercomputing/raven) of the MPCDF. We see that in a direct node-per-node comparison, the solutions of the eigenvalue problems on 4 MI250X GPUs are in general twice as fast as on 4 NVIDIA A100 GPUs, which is in line with the expectations. Due to limited resources, especially limitations in the maximum job run-time, we could not perform a strong-scaling analysis of the eigenvalue problem for each matrix size – especially above a linear dimension of one million. Instead, we had to rather focus on a limited number of experiments and run each eigenvalue solution for a specific matrix size on specific node counts.

This set of runs is shown in Fig. 4 which shows the costs  $C$  to solve a real double-precision dense eigenvalue problem for different matrix sizes. We include the power-law fit of the  $C = M^b$  to the data, representing the scaling behaviour for the best time-to-solution  $b \approx 2.5$ . Since this

$N$ : total number of particles.  $\xi = k_F / \pi \Delta$ : BCS coherence length (typical size of the vortex core).  $L_{\text{box}}$ : size of the simulation domain.  $a$ : scattering length.  $l_F = k_F^{-1} = [3\pi^2 n]^{-1/3}$ : inverse of Fermi momentum.  $L(0)$ : total initial length of vortices.  $l(0)$ : initial mean inter-vortex spacing. All lengths are in units of  $l_F$ . The scales are set such ratio  $l/\xi \approx 5.5$  is fixed across the runs.

power-law exponents  $b$  is still below the theoretical value of 3 (scaling of eigenvalue algorithms is  $\mathcal{O}(x^3)$ ) limit, we have not yet reached the complexity scaling limit for matrix size  $M$ , at least on this HPC system (hardware, compilers, etc.).

Nevertheless, we have successfully solved dense eigenvalue problems for real, double-precision EVPs with linear matrix sizes up to 3.2 million, which is already substantially larger than the problem size required for the target quantum turbulence simulations discussed below. To our knowledge, this is the largest dense eigenvalue problem ever solved with a direct solver. While the benchmark of the ELPA library was executed for symmetric matrices (double precision), in the production computation, we were working with Hermitian matrices (double complex precision). The scaling properties in the Hermitian mode are similar, with the complex case taking about twice as long.

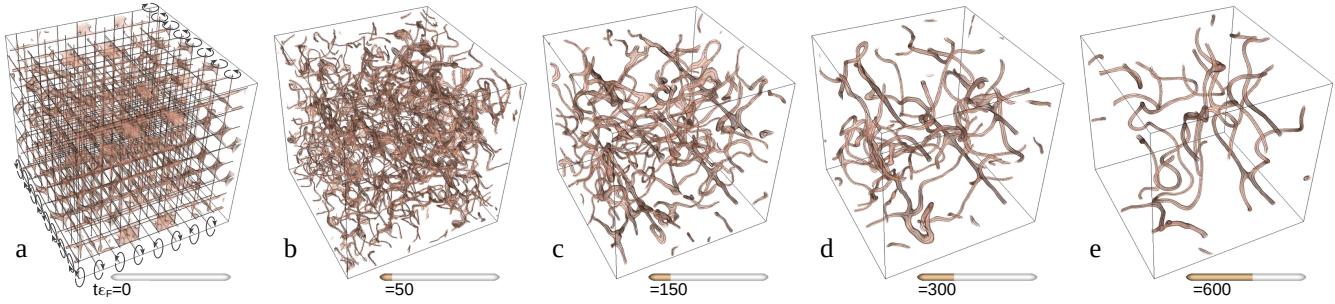
**Results for Quantum Turbulence.** We have presented the capabilities of leadership supercomputers like LUMI to deal with dense matrices and solve nonlinear PDEs. We now combine all these elements together to generate large-scale simulations of turbulent dynamics in ultra-cold Fermi gases.

We start our calculations by preparing the initial state at zero temperature ( $T = 0$ ) consisting of a regular lattice of imprinted vortices in all three directions, see Fig. 5a. The lattice consists of alternately arranged vortices and anti-vortices and the resulting state has zero total angular momentum. The generation of the initial state amounts to solving the static problem (7) with the additional constraint imposed on the phase  $\theta(\mathbf{r})$  of the order parameter  $\Delta(\mathbf{r}) = |\Delta(\mathbf{r})| e^{i\theta(\mathbf{r})}$ . The phase provides a superfluid velocity field  $v_s(\mathbf{r}) \propto \nabla \theta(\mathbf{r})$  consistent with the vortex/anti-vortex lattice with a slight long-wavelength perturbation that destabilizes the vortex lattice leading to a turbulent tangle of vortices as seen in subsequent frames of Fig. 5.

We study the strongly-interacting UFG ( $k_F a = \infty$ ) with both the full fermionic TDSLDA DFT, and a simple modified GPE-like theory Eq. (3) for dimers ( $m_B = 2m$ ) tuned to the UFG equation of state (58). To mimic the natural dissipation of the TDSLDA, we add some artificial damping to the GPE Eq. (3), considering two values:  $\eta = 0.01$  and  $\eta = 0.08$ . The lower value was found to give reasonable qualitative agreement with UFG simulations of rotating quantum turbulence (57), while the larger value better matches the flow energy decay seen in the corresponding fermionic simulation. We also study a less strongly interacting system in the BCS regime at an experimentally accessible value of  $k_F a = -1.8$ . The number of quasiparticle states extracted from the initial state preparation was  $N_{\text{qpf}} = 582\,898$  for the UFG and  $N_{\text{qpf}} = 675\,460$  for the BCS regime.

Table 1 shows some of the characteristic properties of these initial states. There are four length-scales of interest. From smallest to largest, these are: the Fermi scale  $l_F = k_F^{-1}$  is set by the density and is the smallest resolvable scale in the problem; the BCS coherence length  $\xi = k_F / \pi \Delta$  describes the size of the cores of the quantum vortices; the mean inter-

\*See e.g. [elpa.mpcdf.mpg.de/ELPA\\_USED](http://elpa.mpcdf.mpg.de/ELPA_USED).



**Fig. 5.** Selected frames of the time evolution of vortex tangle in the strongly interacting UFG ( $k_F a \rightarrow \infty$ ) superfluid gas. The lines indicate position of the vortex cores, and isosurfaces are used to visualize their sizes. For the simulation we used a periodic lattice with  $N^3 = 100^3$  points. The full movie is provided in Supplementary Material.

vortex distance  $l$  describes the vortex density and is the scale at which we initially inject energy for the turbulence; and the size of the simulation volume  $L_{\text{box}}$ , which is the largest scale in the problem.

Fig. 5 shows the evolution of the UFG in our largest TDSLDA simulation on LUMI. The initial perturbed vortex lattice (Fig. 5a) is unstable, and rapidly forms a vortex tangle (Fig. 5b) by  $t_{eF} \approx 50$ . The subsequent decay of this tangle (Figs. 5c to 5e) transfers hydrodynamic energy from the initial scale of the vortex lattice to other length scales. The bending, crossing, and reconnection of vortices seen in Figs. 5b to 5d are the primary mechanisms for quantum turbulence. Through these mechanisms, hydrodynamic energy can flow from large to small scales, resulting in the emergence of an *effective viscosity* (69, 70) even though this is a superfluid. In compressible fluids, part of this energy is converted into sound (71, 72) and further into internal excitations, making this cascade irreversible. There is also some *weak wave turbulence* (73, 74) in the phonons (sound waves), but this is a small effect, and not visible in these plots.

To quantify these, we introduce flow and condensation energies. The former is just the kinetic energy associated with the flow, while the latter estimates the energy in the condensate (Cooper pairs) which uses a simple formula derived in the BCS limit (33):

$$E_{\text{flow}}(t) = \int \frac{J^2(r, t)}{2n(r, t)} d^3r, \quad E_{\text{cond}}(t) = \frac{3}{8} \int \frac{|\Delta(r, t)|^2}{\varepsilon_F(r, t)} n(r, t) d^3r. \quad [10]$$

We compare the evolution of  $E_{\text{flow}}$  with the total length  $L(t)$  of the vortices in Figs. 6a and 6b. When the vortex core is small (we shall call these *tight* vortices) as in e.g. liquid helium, the flow energy associated with turbulence is dominated by vortices and expected to be proportional to the vortex length. Our case qualitatively differs from that of liquid helium because we have compression modes (phonons or sound), and the vortex core size is comparable to the inter-particle separation. Nevertheless, we still see quite a strong correlation between these.

The evolution of the total length  $L(t)$  during the free decay is one of the main probes to distinguish the type of turbulence (75). In an incompressible fluid, a random tangle of tight vortices with no large-scale structures in the velocity field is expected to develop what is called *Vinen* or *ultraquantum* turbulence where the total vortex length decays as  $L \propto t^{-1}$ . On the other hand, if the vortices create large-scale structures – i.e. bundles of coherent vortices – then one expects the fluid to develop eddies and dynamics that produce *Kolomogorov* or *quasiclassical* turbulence (76, 77) with a characteristic decay of  $L \propto t^{-3/2}$ . In our simulations (Fig. 6b), we see Kolomogorov-like decay in the GPE when the vortex density is high  $50 \lesssim t_{eF} \lesssim 300$ , but the TDSLDA never develops this behaviour, suggesting a fundamental difference in the decay mechanism between bosonic and fermionic simulations. In this regard, the TDSLDA appears to more closely match the decay predicted by Vinen turbulence, but we suspect this is coincidental rather than causal since,

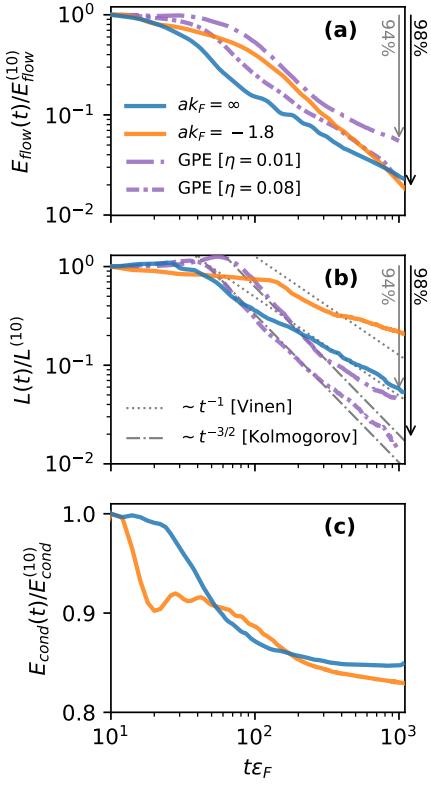
as we show below, there are additional dissipation mechanisms present in this case. Note that in this work, we do a comparative study between two methods (GPE vs TDSLDA) since we have access to data for identical setups. However, the lack of statistics does not allow us to make statements about the precise value of the decay exponent.

For compressible turbulence, one can use a Helmholtz decomposition to split  $E_{\text{flow}}$  into divergence-free (incompressible/rotational/vortices) and curl-free (compressible/irrotational/phonons) parts  $E_{\text{flow}} = E_{\text{vortices}} + E_{\text{phonons}}$ . (See e.g. (78).) The total vortex length  $L$  should be most strongly correlated with the  $E_{\text{vortices}}$  contribution, but now vortex reconnections can produce phonons (71, 72), further reducing  $E_{\text{vortices}}$ . Thus, we expect the decrease in  $L(t)$  to be more pronounced than in  $E_{\text{flow}}(t)$ .

To check if this expectation is seen in Fig. 6, we note that, by coincidence, both TDSLDA simulations have lost about 98 % of their total flow energy at  $t_{eF} \approx 1000$ , so we use this as a fiducial. (The phenomenological dissipation  $\eta = 0.08$  was chosen for one GPE to match these results; the other GPE simulation with  $\eta = 0.01$  has a 94 % loss in  $E_{\text{flow}}$  at this time.) Here, we find that the GPE demonstrates the expected trend for both cases, with a slightly greater drop in  $L$ . The TDSLDA simulations, however, have a quantitatively different behaviour with a smaller decrease in  $L$  of 94 % in the UFG, and 80 % in the BCS regime.

This result might seem to be counterintuitive: in the BCS regime we still have many vortices at the end, but they do not generate much flow. We interpret this as a strong indication that another mechanism (absent in the GPE) is responsible for reducing  $E_{\text{flow}}$  in the TDSLDA. Noting that the simple correlation  $L \sim E_{\text{vortices}}$  holds only if we do not consider corrections from the internal vortex core structure, we hypothesize that thermalization plays a significant role. To demonstrate this, in Fig. 7 we consider a vortex solution for the BCS case, as a function of temperature  $T$ , which allows us to manipulate the size and structure of the vortex core. The results are obtained by solving static Eq. (7) with the constraint that we have a single and straight vortex line. Far from the core, the density  $n(r)$  has the same behavior, but clearly, the density inside the vortex core is sensitive to the temperature (Fig. 7a). The order parameter distribution  $\Delta(r)$  also indicates that vortices get bigger with the increase of  $T$  (Fig. 7b). Accordingly, the velocity  $v(r) = j(r)/n(r)$ , which quantifies the flow energy, is suppressed by the thermal effects (Fig. 7c and inset Fig. 7d). This shows that the structure of the vortex core, which is sensitive to  $T$ , can affect  $E_{\text{flow}}$ . Specifically, in the inset Fig. 7d, we see that finite temperature significantly reduces both the flow and condensation energy.

This suggests an explanation for the breakdown of the correlation between  $E_{\text{flow}}(t)$  and  $L(t)$ : the TDSLDA admits an additional dissipation mechanism whereby flow energy is “thermalized”, altering the flow structure of the vortices. To explicitly demonstrate that the vortices in the time-dependent runs get hotter, we added to Fig. 7 (thin gray lines)

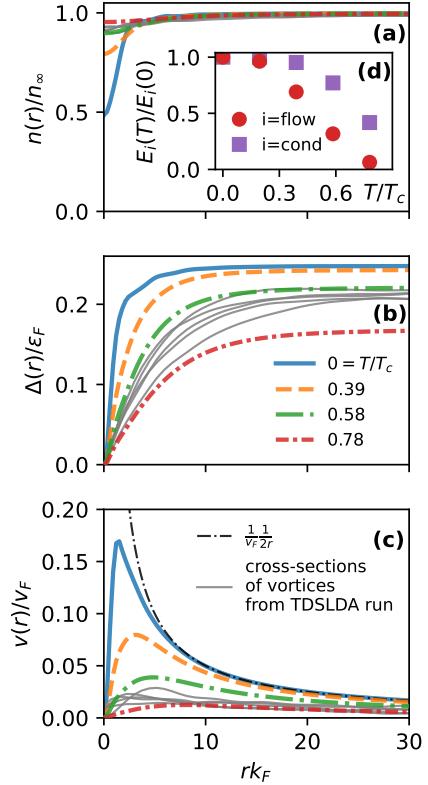


**Fig. 6.** Time evolution of: (a) the flow energy  $E_{\text{flow}}$ , (b) total vortex length  $L$ , and (c) the condensation energy  $E_{\text{cond}}$ . All quantities are normalized with respect to their initial values. The solid lines are from the TDSLDA calculations on LUMI for the UFG (dark blue) and BCS (light orange) regimes. The dashed-dotted lines correspond to the simplified GPE model with two different phenomenological dissipation coefficients  $\eta \in \{0.01, 0.08\}$ . In panel (b), the thin gray lines show the expected slope for  $L(t)$  expected for Vinen (dotted) or Kolmogorov (dashed-dotted) turbulence. The arrows indicate where the analyzed quantity drops by 94 % and 98 %.

cross-sections through five randomly selected vortices from the BCS runs at  $t\epsilon_F = 1000$ . As expected for a non-equilibrium state, the profiles of individual vortices have some variability, but all characteristics – the density profile, the order parameter profile, and the velocity field – are consistent with the static solutions obtained for a temperature  $T/T_c \approx 0.6$ .

The temperature dependence of the vortex-core density  $n_{\text{core}}$  allows us to use fermionic vortices as a local thermometer. Using static SLDA simulations of a single vortex, we calibrate  $n_{\text{core}}(T)$ , (the curve is provided in the Supplementary Material) and then use the density along the vortex lines to demonstrate the thermal evolution of the turbulence in the TDSLDA simulations. These results are presented in Fig. 8. We observe that the effective temperature of vortex lines is higher in regions of higher curvature, especially in regions where reconnections occur. This is reminiscent to the heating of wire which is sharply bent back and forth. Heating of the vortices represents an additional dissipative mechanism missing in GPE-like models.

Heating depletes the condensate, which eventually vanishes at the critical temperature  $T_c$  at the superfluid to normal phase transition. This heating is indicated by the loss of condensation energy as seen in Fig. 6c. In the UFG, the condensate is depleted only at the early stages  $t\epsilon_F \lesssim 200$ , after which it remains relatively constant. In contrast, in the BCS regime, the condensate continues to gradually deplete. In the case of the BCS vortex solutions, the depletion of the condensate by about 20 % is for  $T \approx 0.6T_c$  (Fig. 7d). This value approximately corresponds to the



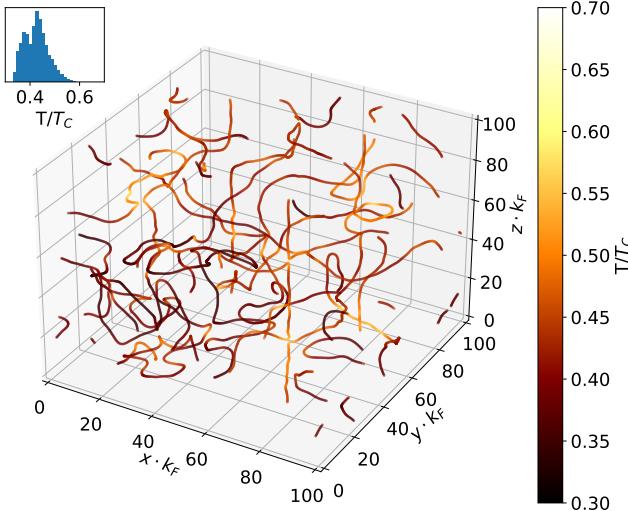
**Fig. 7.** Radial dependence of the: (a) density  $n(r)$ , (b) order parameter  $\Delta(r)$ , and (c) velocity  $v(r) = j(r)/n(r)$  for a single straight vortex line at various temperatures in the BCS regime ( $k_F a = -1.8$ ). The thin gray lines show the corresponding profiles of five randomly selected vortices from the full TDSLDA taken at time  $t\epsilon_F = 1000$ . The temperature is normalized by the critical temperature of superfluid-normal phase transition  $T_c$  within the static model: i.e. for  $T = T_c$  one would have  $\Delta(r) = 0$ . In the lower panel (c) we include the asymptotic form  $v(r) = 1/2r$  as a thin dash-dotted line. Inset (d): temperature dependence of  $E_{\text{flow}}$  and  $E_{\text{cond}}$  energies for a single vortex. Together, these suggest an effective temperature of  $T \approx 0.6T_c$  at time  $t\epsilon_F = 1000$ .

estimated effective temperature of vortices from the cross-sections. It is another signature pointing to the conversion of the flow energy into internal excitations, which effectively heat the system.

## Discussion

In the UFG, the turbulent dynamics resolved by the TDSLDA demonstrate qualitative differences when compared with the simplified approach based on a modified GPE. For instance, the GPE does not properly account for physics in the cores of vortices: Where fermionic vortices have a finite density in the core, vortices in the GPE are empty, causing them to move at different speeds. This can be somewhat compensated for by averaging procedures (58), but feeding this back into the evolution has proved tricky (see (79) for one approach). It is also insufficient to model the dissipative effects with a single phenomenological parameter  $\eta$ . Here we have adjusted  $\eta \approx 0.08$  to reasonably match the decay pattern for  $E_{\text{flow}}$  seen in the TDSLDA, but this differs from the value  $\eta \approx 0.01$  that best fits rotating turbulence (57). Nevertheless, the GPE provides a fast way of gaining some insight into the qualitative effects of superfluid dynamics.

The TDSLDA is a parameter-free, self-consistent microscopic theory that naturally captures these effects. Thus, although more costly computationally, it provides deeper insight into superfluids dynamics than any other currently available techniques. Here we have demonstrated that



**Fig. 8.** Vortex lines with the color indicating their local effective temperature. The temperature is extracted from the relation between density inside the vortex and the system's temperature for static solutions, see Fig. 7(a). This snapshot is for UFG simulation at  $t\epsilon_F = 300$ , where the reconnection event occurs and represents the same snapshot as Fig. 5d. The inset shows the histogram of temperatures along the vortex lines. The full movie is provided in the Supplementary Material. The highest temperatures are generally correlated with events like vortex crossings and reconnections at sites of high curvature, but this energy dissipates, resulting in a uniform effective temperature at later stages of evolution.

the **TDSLDA** breaks the natural correlation between the total flow energy  $E_{\text{flow}}$  and the total vortex length  $L$ . By studying the structure of vortex cores at various temperatures  $T$ , we provide evidence that this is due to additional energy dissipation and thermalization mechanisms. Directly comparing the **TDSLDA** with a **GPE**-like theory, we can distinguish this mechanism from dissipation due to vortex bending, crossing, and Kelvin modes, for example, which exist in both theories. The importance of the fermionic nature of the superfluid is further supported by the result that deviations are stronger in the **BCS** regime than in the **BEC** regime where **GPE**-like theories should be accurate.

Some caveats are in order, and results at finite  $T$  in the **SLDA** must be treated with some caution. In principle, the functional (precisely functions  $A$ ,  $B$ , and  $C$ , in Eq. (15)) should now depend on the dimensionless parameter  $k_B T / \epsilon_F(n)$ , but the form of these has not yet been fit due to lack of reliable benchmarks. Furthermore, it is an open question about how thermalization in strictly  $T = 0$  dynamical simulations (Fig. 6) should be related to the explicit thermal distribution  $f_T(E)$  used in static calculations (Fig. 7). The agreement between these strongly suggests that thermalization is the correct explanation for the break in correlation between  $E_{\text{flow}}(t)$  and  $L(t)$ , but additional analysis is needed to conclusively rule out geometric effects, to identify the importance of vortices to the thermalization process (80), etc.

This type of progress requires the development of all exascale **HPC** technology components: computing hardware, optimized scientific libraries for **FFT** (**hipFFT**) and matrix diagonalization (**ELPA**), and advanced high-level scientific application software (the **TDSLDA** via the **W-SLDA** toolkit). We have demonstrated that, with current technologies, we can diagonalize matrices of order a million by a million, and use **TDDFT** to model fermionic quantum dynamics with  $10^4$ - $10^5$  particles. This is rapidly approaching the scale of typical ultra-cold atom experiments ( $10^5$ - $10^6$  particles), allowing us to directly benchmark the **TDSLDA** against experiments. The new generation of **DFT** packages optimized for **HPC** are extremely flexible, enabling researchers to access not only superfluidity as studied here, but many other fields, including superconductivity, nuclear physics (15, 81), nuclear astrophysics, and new

opportunities for emerging fields like atomtronics and other quantum technologies (82). Simulations like these provide access to complex but important microscopic physical processes, and thus provide benchmarks to tune more economical phenomenological models like (58, 79), which may ultimately be scaled up to address problems of importance to fundamental physics in nuclear astrophysics. (E.g., the **UFG** we study here is a good model for the neutron superfluid in the crust of neutron stars, which is likely responsible for pulsar glitches.)

To maximize the scientific impact of these publicly-funded resources, we have open-sourced our codes (63), and release the raw data generated by these simulations (83) so that others can reproduce our analysis, and use these results for further benchmarking other models, thereby advancing the pace of science. We hope that other groups will follow our example, making their codes and data available for the community. This step is needed in order to maximize the amount of knowledge extracted from data obtained by expensive **HPC** systems, and share research opportunities with groups that do not have direct access to them.

## Materials and Methods

**LUMI specification.** LUMI is one of the three Euro **HPC** pre-exascale supercomputers, along with Leonardo and Marenostrum 5. It is the only major European **HPC** system equipped with AMD GPU technology. The numerical results described in this work were collected from the LUMI pilot access phase. The LUMI system has a peak performance of nearly 500 **PFLOPS**, most of which is delivered by 2560 GPU nodes. Each of these nodes includes a single AMD **epyc** CPU with 64 cores and four AMD **M1250X** GPUs. The AMD **M1250X** has a peak performance of 53 **TFLOPS** in double-precision arithmetic. The AMD **M1250X** GPU package consists of two independent devices, called graphics compute dies (**CCDs**). Each of these **CCDs** has 110 compute units, and 64 **GB** of high-bandwidth memory (**HBM2**) which can be accessed at a peak rate of 1.6 **TB/s**. The two **CCDs** in the **M1250X** package are connected by an in-package communication interface with a peak bidirectional bandwidth of up to 400 **GB/s**. Devices on different packages are linked with either a single or double communication link with a peak bidirectional bandwidth of 100 **GB/s** and 200 **GB/s**, respectively. Each GPU package is directly connected to the **Slingshot** network providing up to 2 $\times$ 50 **GB/s** peak bandwidth. The AMD **M1250** GPU family is based on the 2<sup>nd</sup> generation AMD **CDNA** ("Compute DNA") architecture which uses AMD **ROCm** (Radeon Open Compute) development stack. **ROCm** is an open-source collection of drivers, development tools and application programming interfaces (**APIs**) for GPU programming from the low-level kernel to end-user libraries. Device kernels are programmable with the **HIP** GPU programming language extension. The **HIP** extension also provides a runtime platform, numerical libraries (including the **hipFFT**), and porting tools.

**Density functional theory for superfluid fermions.** We consider here a system with equal densities of two types of fermions interacting with short-range interactions. In second-quantized notation,

$$\hat{H} = \int d^3r \frac{\hat{a}^\dagger(r) \hat{p}^2 \hat{a}(r) + \hat{b}^\dagger(r) \hat{p}^2 \hat{b}(r)}{2m} + \int d^3r d^3r' V(|r - r'|) \hat{n}_a(r) \hat{n}_b(r'), \quad [1]$$

where  $\hat{n}_a = \hat{a}^\dagger \hat{a}$  and  $\hat{n}_b = \hat{b}^\dagger \hat{b}$  are the number operators for the two species, expressed in terms of annihilation  $\hat{a}/\hat{b}$  and creation  $\hat{a}^\dagger/\hat{b}^\dagger$  operators that satisfy anti-commutation relations. We consider the limit where  $V(r)$  is short-ranged, so that the interaction can be completely described by a single parameter, the *s*-wave scattering length  $a$ .

In general, the state of such a system with  $N$  particles must be described by a many-body wavefunction  $\psi(r_1, r_2, \dots, r_N)$  which requires

an exponential amount of information, but DFT allows us to reduce this description to an effective theory for the total density  $n(\mathbf{r}) = \langle \hat{n}_a(\mathbf{r}) \rangle + \langle \hat{n}_b(\mathbf{r}) \rangle$  or states which can be expressed as a Slater determinant – an anti-symmetrized product of single-particle states. Treated as a variational problem of finding the best single-Slater-determinant state (with some care required to express the zero-range limit), one derives a BCS-like ansatz in what is commonly referred to as Hartree-Fock-Bogoliubov (HFB) theory or Bogoliubov-de Gennes (BdG) theory for superconductivity. (See e.g. (55, 56, 84).) These theories qualitatively capture properties of the system Eq. (11), but quantitatively fail in regions of interest like the UFG. The SLDA has a similar mathematical form, but is constructed from a different philosophy – that of density functional theory (DFT).

The TDDFT equations follow from a condition of stationary action

$$\delta S = 0, \quad S = \int_{t_0}^{t_1} \left( \left\langle \Omega(t) \left| i\hbar \frac{d}{dt} \right| \Omega(t) \right\rangle - E(t) \right) dt, \quad [12]$$

where  $|\Omega(t)\rangle$  is quasiparticle vacuum at time  $t$  and  $E(t)$  is the total energy. The key to all DFT approaches is the existence theorem due to Hohenberg and Kohn (4), with extension to the time-dependent cases by Runge and Gross (7), that for any given system,  $E(t)$  can be expressed as

$$E(t) = E_{\text{int}}[n(\mathbf{r}, t)] + \int V_{\text{ext}}(\mathbf{r}, t) n(\mathbf{r}, t) d^3r \quad [13]$$

where  $E_{\text{int}}[n]$  is a universal functional and  $V_{\text{ext}}(\mathbf{r}, t)$  is the external potential (which we set to zero here). Unfortunately, no prescription for finding  $E_{\text{int}}[n]$  is known, and it is likely extremely complicated and non-local, even for non-interacting fermions. Instead, Kohn and Sham (5) derived an equivalent but alternate formulation in terms of an energy functional of a Slater determinant  $|\Omega(t)\rangle$  of single-particle orbitals that allows for an exact local formulation for non-interacting fermions. By including both particles and holes, and with an appropriate regularization procedure, this was generalized for superfluids (19–21, 85) in a form called the SLDA that we now describe. Unlike the BdG equations, we can now tune the parameters of the theory to match experiment and ab-initio QMC calculations. We lose any notion of a variational bound, but obtain instead a theory accurate to the few-percent level for a wide range of systems (21–23).

The energy function for the SLDA can be expressed as an integral of four local “densities” – each of which is a function of the orbitals in the Slater-determinant state  $|\Omega(t)\rangle$ :

$$E(t) = \int d^3r \mathcal{E}[n(\mathbf{r}, t), \tau(\mathbf{r}, t), j(\mathbf{r}, t), v(\mathbf{r}, t)]. \quad [14]$$

At  $T = 0$ , this system has only two length-scales: the inverse Fermi momentum  $k_F^{-1}$  and the  $s$ -wave scattering length  $a$ . The SLDA functional is thus constrained by dimensional arguments (see Eq. (1)) to have the following form (in units where  $\hbar = m = k_B = 1$ ):

$$\mathcal{E} = A(k_F a) \frac{\tau}{2} + B(k_F a) \frac{3}{5} n e_F + \frac{C(k_F a)}{n^{1/3}} |v|^2 + [1 - A(k_F a)] \frac{j^2}{2n}, \quad [15]$$

where  $A$ ,  $B$ , and  $C$  are dimensionless universal functions. The first term defines the kinetic energy, the second term (missing in BdG theory) describes the Hartree energy, the third term accounts for the energy gain due to pairing correlations, and the last term is required to restore Galilean covariance.

By appropriately choosing the universal functions, one can describe the entire BEC-BCS crossover, including the weakly-interaction BCS limit  $k_F a \rightarrow 0^-$ , and the UFG  $|k_F a| \rightarrow \infty$ . The latter is especially simple because  $A$ ,  $B$ , and  $C$ , are just numbers. Precisely, the functions are constructed in such way as to ensure correct reproductions of selected properties of uniform Fermi gas at a given value of the interactions parameter  $k_F a$ . These properties are: equations of state  $\zeta(k_F a) = \mathcal{E}/\mathcal{E}_{FG}(n)$ , strength of the pairing correlations  $\Delta(k_F a)/\epsilon_F(n)$ , and the quasiparticle effective mass  $m^*(k_F a) = m/A(k_F a)$ . All of these quantities are accessible

from QMC calculations. For more details related to the construction of the functional, see (19–23, 27, 35, 85) and the source code for the reference implementation. In the calculations presented here, we have assumed that the effective mass  $m^* = m$ , so  $A(k_F a) = 1$ . This is a physically reasonable approximation that simplifies the functional slightly since the last term of Eq. (15) vanishes.

The local densities entering the functional are:

$$n(\mathbf{r}, t) = \langle \Omega(t) | \hat{a}^\dagger(\mathbf{r}) \hat{a}(\mathbf{r}) + \hat{b}^\dagger(\mathbf{r}) \hat{b}(\mathbf{r}) | \Omega(t) \rangle, \quad [\text{total}]$$

$$\tau(\mathbf{r}, t) = \langle \Omega(t) | \nabla \hat{a}^\dagger(\mathbf{r}) \cdot \nabla \hat{a}(\mathbf{r}) + \nabla \hat{b}^\dagger(\mathbf{r}) \cdot \nabla \hat{b}(\mathbf{r}) | \Omega(t) \rangle, \quad [\text{kinetic}]$$

$$j(\mathbf{r}, t) = \text{Im} \langle \Omega(t) | \hat{a}^\dagger(\mathbf{r}) \cdot \nabla \hat{a}(\mathbf{r}) + \hat{b}^\dagger(\mathbf{r}) \cdot \nabla \hat{b}(\mathbf{r}) | \Omega(t) \rangle, \quad [\text{current}]$$

$$v(\mathbf{r}, t) = \langle \Omega(t) | \hat{a}(\mathbf{r}) \hat{b}(\mathbf{r}) | \Omega(t) \rangle. \quad [\text{anomalous}]$$

After varying the functional, we will obtain a matrix equation Eq. (5) (derived below) with eigenvalues  $E_n$  and two-component eigenstates  $\phi_n(\mathbf{r}, t) = [u_n(\mathbf{r}, t), v_n(\mathbf{r}, t)]^T$ . In terms of these, the densities are:

$$n(\mathbf{r}, t) = 2 \sum_{0 < E_n < E_c} \left( |v_n|^2 f_T(-E_n) + |u_n|^2 f_T(E_n) \right), \quad [16a]$$

$$\tau(\mathbf{r}, t) = 2 \sum_{0 < E_n < E_c} \left( |\nabla v_n|^2 f_T(-E_n) + |\nabla u_n|^2 f_T(E_n) \right), \quad [16b]$$

$$j(\mathbf{r}, t) = 2 \sum_{0 < E_n < E_c} \text{Im} \left( (v_n \nabla v_n^*) f_T(-E_n) - (u_n \nabla u_n^*) f_T(E_n) \right), \quad [16c]$$

$$v(\mathbf{r}, t) = \sum_{0 < E_n < E_c} u_n v_n^* \left( f_T(-E_n) - f_T(E_n) \right), \quad [16d]$$

where we have suppressed the space-time arguments on the components  $u_n$  and  $v_n$  to save space, and we have introduced the thermal distribution function (Fermi distribution)

$$f_T(E) = \frac{1}{1 + \exp(E/T)}, \quad [17]$$

which allows us to approximate finite temperatures  $T$ . We used the finite-temperature variant only for understanding the structure of the vortex core as presented in Fig. 7; the time-dependent runs were executed at  $T = 0$ .

Note that all states up to a specified energy cutoff  $E_c = \hbar^2 k_c^2 / 2m$  contribute to the densities: This is the main difference between the DFT for superfluid systems and the original formulation of Kohn and Sham which only keeps states up to the Fermi surface. For superfluids, the single-particle state  $|\Omega(t)\rangle$  has a coherent phase and hence does not have a well-defined particle number. Instead, the BCS particle-hole correlations (Cooper pairs) allow fractional occupation of states with larger energy. The cutoff  $E_c$  is essential for a local formulation with short-range interactions because the kinetic  $\tau(E_c)$  and anomalous  $v(E_c)$  densities are linearly divergent in such a state

While both  $\tau$  and  $v$  diverge linearly  $\propto k_c$ , the combination  $\hbar^2 \tau / 2m - \Delta^* v$  is finite where  $\Delta = g_c v$  is the finite pairing gap. To improve convergence (19), we choose a scale  $k_0 \approx k_F$  and use

$$\Lambda_c = \frac{m}{\hbar^2} \frac{k_c}{2\pi^2} \left\{ 1 - \frac{k_0}{2k_c} \ln \frac{k_c + k_0}{k_c - k_0} \right\}$$

so that the functional is specified by the density dependence of the finite combination  $A/g_c + \Lambda_c$ . This finite quantity is proportional to the inverse scattering length  $1/a$  in the standard variational formulation of BdG theory, but has additional dependence on  $k_F a$  in the SLDA. The regularization scheme is constructed to assure independence of the observables with respect to the energy cutoff  $E_c$  (assuming it is large enough to encapsulate meaningful states) when considering static problems. In the context of time-dependent problems, another factor needs to be considered: the total energy must be conserved. Formally, the TDSLDA equations (5) are conservative only in  $E_c \rightarrow \infty$  limit (86). For this reason, in the computation, we use the cutoff scale  $E_c = k_{\max}^2 / 2$  defined through the maximum value of momentum  $k_{\max} = \pi$  resolved by our spatial grid. In our simulations, the relative change in of the

total energy  $|E(t) - E(0)| / E(0)$  does not exceed a fraction of a percent. See (19–23, 27, 35, 85) and the source code for further details.

The equations of motion that emerge from the stationarity condition Eq. (12) have the forms Eqs. (5) and (7) discussed in the results above. The single particle Hamiltonian  $\hat{h}$  and pairing potential  $\Delta$  are defined through functional derivatives of the energy density

$$\hat{h} = -\nabla \frac{\delta E}{\delta \tau} \nabla + \frac{\delta E}{\delta n} - \frac{i}{2} \left( \frac{\delta E}{\delta j} \nabla + \nabla \frac{\delta E}{\delta j} \right) - \mu, \quad [18a]$$

$$\Delta = -\frac{\delta E}{\delta v^*} = -\underbrace{\frac{C}{n^{1/3}}}_{g_c} v. \quad [18b]$$

They depend on densities, which in turn depend on the quasiparticle orbitals  $[u_n(r), v_n(r)]^T$ . The single particle Hamiltonian  $\hat{h}$  includes a shift by the value of the chemical potential  $\mu$ . This controls the particle number in the static solution for the initial state Eq. (7). It is formally irrelevant for the time-evolution Eq. (5), but helps improve numerical convergence by minimizing the evolution of the global phase. The regularized coupling function  $g_c$  defines the order parameter  $\Delta$ , ensuring it remains finite.

**Initial state preparation and vortex lines detection.** To generate the quantum turbulence, we start with a regular alternating array of interleaved vortices and anti-vortices in all three directions. (See Fig. 5a.) Using the Biot-Savart law, including image vortices from the periodic box, we obtain the phase profile  $\theta_0(r)$  for this periodic array with a superfluid velocity field  $v_s(r) \propto \nabla \theta_0(r)$  consistent with the vortex/anti-vortex lattice. On this periodic phase profile, we add a few small low-frequency Fourier components  $\theta(r) = \theta_0(r) + a_0 \sum_{n=0}^{N_v} c_n \cos(k_n \cdot r)$ , where  $N_v$  is the number of vortices in each direction. The coefficients and frequencies are  $c_n = (-1)^n / (2n+1)^2$ , and  $k_n^i = (2n+1)(2\pi) / L_{\text{box}}^i$ . The magnitude of the perturbation  $a_0$  is adjusted so that the additional large-scale flow increases  $E_{\text{flow}}$  by 5% compared with the vortex lattice (as computed within the GPE). This phase profile is then held fixed in  $\Delta(r) = |\Delta(r)| e^{i\theta(r)}$ , and the iterative solution for the magnitude of the gap  $|\Delta|$ , density, etc. is found using Eq. (7) as described above, which requires the improved ELPA diagonalization routines.

To compute the total length of  $L$ , we need first to identify the position of vortices from numerical data. For this, we search in  $\Delta(r)$  fields for points around which its phase winds. This also implies that, at this point, the order parameter vanishes. We interpolate grid data using discrete variable representation (DVR) to identify such points with subgrid resolution (87). To connect points into lines we use pseudovorticity  $\omega(r) = \nabla \times j(r)$ , which should point along the line. A detailed description of the algorithm is given in (57).

## Supplementary Materials

Supplementary material is available at PNAS Nexus online.

## Acknowledgments

The MPCDF, and especially A. Marek, is grateful for the fruitful discussions and support of AMD for optimizing the ELPA library on AMD GPUs. The authors are grateful to the support team of the CSC data center and their help with using the LUMI HPC system. We thank P. Magierski for inspiring discussions, and C. Barenghi for suggesting additional tests. This work was supported by the Polish National Science Center (NCN) under Contract No. UMO-2017/26/E/ST3/00428, and later stages of this work were supported under Contract No. UMO-2022/45/B/ST2/00358. The development of the ELPA library is financially supported by the NOMAD Centre of Excellence (European Commission grant agreement ID 951786). S.S. and M.M.F. acknowledge support from the National Science Foundation (NSF) through Grant No. PHY-2012190.

## Author contributions

The work was conceived by G.W., and M.M.F. The TDSLDA calculations were done by G.W. The GPE calculations were done by S.S. Porting of the w-SLDA Toolkit to AMD was done by G.W., and M.S. Porting and optimizing of ELPA to AMD was done by A.M. Benchmarking ELPA on LUMI was done by A.M. Data analysis was done by G.W., S.S., and M.M.F. Support of computation processes on LUMI was done by M.S. All authors contributed to the manuscript writing.

## Data and materials availability

Data underlying the study is available via Zenodo repository (83).

## References

1. S Chatrchyan, , et al., Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC. *Phys. Lett. B* **716**, 30–61 (2012).
2. G Aad, , et al., Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC. *Phys. Lett. B* **716**, 1–29 (2012).
3. BP Abbott, , et al., Observation of gravitational waves from a binary black hole merger. *Phys. Rev. Lett.* **116**, 061102 (2016).
4. P Hohenberg, W Kohn, Inhomogeneous electron gas. *Phys. Rev.* **136**, B864–B871 (1964).
5. W Kohn, LJ Sham, Self-consistent equations including exchange and correlation effects. *Phys. Rev.* **140**, A1133–A1138 (1965).
6. R Dreizler, E Gross, *Density Functional Theory: An Approach to the Quantum Many-body Problem*. (Springer), (1990).
7. E Runge, EKU Gross, Density-functional theory for time-dependent systems. *Phys. Rev. Lett.* **52**, 997 (1984).
8. MAL Marques, NT Maitra, FMS Nogueira, EKU Gross, A Rubio, eds., *Fundamentals of Time-Dependent Density Functional Theory*. Lecture Notes in Physics. (Springer, Heidelberg) Vol. 837, (2012).
9. C Folhais, F Nogueira, MAL Marques, eds., *A Primer in Density Functional Theory*. (Springer Berlin Heidelberg), (2003).
10. A Zangwill, The education of Walter Kohn and the creation of density functional theory. *Arch. Hist. Exact Sci.* **68**, 775–848 (2014).
11. Boosting materials modelling. *Nat. Mater.* **15**, 365–365 (2016).
12. N Mardirossian, M Head-Gordon, Thirty years of density functional theory in computational chemistry: an overview and extensive assessment of 200 density functionals. *Mol. Phys.* **115**, 2315–2372 (2017).
13. G Colò, Nuclear density functional theory. *Adv. Phys.* **X** **5**, 1740061 (2020).
14. J Yang, J Piekarewicz, Covariant density functional theory in nuclear physics and astrophysics. *Annu. Rev. Nucl. Part.* **70**, 21–41 (2020).
15. T Nakatsukasa, K Matsuyanagi, M Matsuo, K Yabana, Time-dependent density-functional description of nuclear dynamics. *Rev. Mod. Phys.* **88**, 045004 (2016).
16. JG Bednorz, KA Müller, Possible high-Tc superconductivity in the Ba-La-Cu-O system. *Zeitschrift für Physik B Condens. Matter* **64**, 189 (1986).
17. LN Oliveira, EKU Gross, W Kohn, Density-functional theory for superconductors. *Phys. Rev. Lett.* **60**, 2430–2433 (1988).
18. OJ Wacker, R Kümmel, EKU Gross, Time-dependent density-functional theory for superconductors. *Phys. Rev. Lett.* **73**, 2915–2918 (1994).
19. A Bulgac, Local density approximation for systems with pairing correlations. *Phys. Rev. C* **65**, 051305 (2002).
20. A Bulgac, Local-density-functional theory for superfluid Fermionic systems: The unitary gas. *Phys. Rev. A* **76**, 040502 (2007).
21. A Bulgac, MM Forbes, P Magierski, *The Unitary Fermi Gas: From Monte Carlo to Density Functionals*, Lecture Notes in Physics, chapter 9, pp. 305 – 373. Vol. 836 of Zwerger (36), pp. 305 – 373 (2012).
22. A Bulgac, Time-dependent density functional theory and the real-time dynamics of Fermi superfluids. *Annu. Rev. Nucl. Part. Sci.* **63**, 97–121 (2013).
23. A Bulgac, MM Forbes, G Włazłowski, Towards quantum turbulence in cold atomic fermionic superfluids. *J. Phys. B* **50**, 014001 (2017).
24. K Choudhary, K Garry, Designing high-Tc superconductors with BCS-inspired screening, density functional theory, and deep-learning. *Npj Comput. Mater.* **8**, 244 (2022).
25. A Bulgac, P Magierski, KJ Roche, I Stetcu, Induced Fission of  $^{240}\text{Pu}$  within a Real-Time Microscopic Framework. *Phys. Rev. Lett.* **116**, 122504 (2016).
26. P Magierski, K Sekizawa, G Włazłowski, Novel role of superfluidity in low-energy nuclear reactions. *Phys. Rev. Lett.* **119**, 042501 (2017).
27. MM Forbes, S Gandolfi, A Gezerlis, Resonantly interacting fermions in a box. *Phys. Rev. Lett.* **106**, 235303 (2011).
28. MM Forbes, S Gandolfi, A Gezerlis, Effective-range dependence of resonantly interacting fermions. *Phys. Rev. A* **86**, 053603 (2012).
29. A Bulgac, Time-Dependent Density Functional Theory and the Real-Time Dynamics of Fermi Superfluids. *Annu. Rev. Nucl. Part. Sci.* **63**, 97–121 (2013).
30. A Bulgac, MM Forbes, MM Kelley, KJ Roche, G Włazłowski, Quantized Superfluid Vortex Rings in the Unitary Fermi Gas. *Phys. Rev. Lett.* **112**, 025301 (2014).
31. G Włazłowski, A Bulgac, MM Forbes, KJ Roche, Life cycle of superfluid vortices and quantum turbulence in the unitary Fermi gas. *Phys. Rev. A* **91**, 031602(R) (2015).
32. G Włazłowski, K Sekizawa, M Marchwińska, P Magierski, Suppressed solitonic cascade in spin-imbalanced superfluid Fermi gas. *Phys. Rev. Lett.* **120**, 253002 (2018).
33. G Włazłowski, K Xhani, M Tylutki, NP Proukakis, P Magierski, Dissipation mechanisms in Fermionic Josephson junction. *Phys. Rev. Lett.* **130**, 023003 (2023).
34. A Barresi, A Boulet, P Magierski, G Włazłowski, Dissipative dynamics of quantum vortices in Fermionic superfluid. *Phys. Rev. Lett.* **130**, 043001 (2023).

35. A Bulgac, MM Forbes, *Time-dependent density functional theory*, chapter 4, pp. 4-1-4-34. In Schunck (84), pp. 4-1-4-34 (2019).

36. W Zwerger, ed., *The BCS-BEC Crossover and the Unitary Fermi Gas*, Lecture Notes in Physics. (Springer-Verlag, Berlin Heidelberg) Vol. 836, (2012).

37. MJH Ku, AT Sommer, LW Cheuk, MW Zwierlein, Revealing the Superfluid Lambda Transition in the Universal Thermodynamics of a Unitary Fermi Gas. *Science* **335**, 563-567 (2012).

38. G Zürn, et al., Precise characterization of  $^6\text{Li}$  Feshbach resonances using trap-sideband resolved RF spectroscopy of weakly bound molecules. *Phys. Rev. Lett.* **110**, 135301 (2013).

39. J Carlson, S Gandolfi, KE Schmidt, S Zhang, Auxiliary field quantum Monte Carlo for strongly paired fermions. *Phys. Rev. A* **84**, 061602 (2011).

40. A Gezerlis, J Carlson, Strongly paired fermions: Cold atoms and neutron matter. *Phys. Rev. C* **77**, 032801 (2008).

41. A Lopez, P Kelly, K Dauer, E Vitali, Fermionic superfluidity: from cold atoms to neutron stars. *Eur. J. Phys.* **43**, 065801 (2022).

42. A Bulgac, Time-dependent density functional theory for fermionic superfluids: From cold atomic gases - to nuclei and neutron stars crust. *physica status solidi (b)* **256**, 1800592 (2019).

43. CF Barenghi, RJ Donnelly, WF Vinen, eds., *Quantized Vortex Dynamics and Superfluid Turbulence*. (Springer Berlin Heidelberg), (2001).

44. M Tsubota, M Kobayashi, H Takeuchi, Quantum hydrodynamics. *Phys. Rep.* **522**, 191-238 (2013) Quantum hydrodynamics.

45. EAL Henn, JA Seman, G Roati, KMF Magalhães, VS Bagnato, Emergence of turbulence in an oscillating bose-einstein condensate. *Phys. Rev. Lett.* **103**, 045301 (2009).

46. EAL Henn, et al., Observation of vortex formation in an oscillating trapped bose-einstein condensate. *Phys. Rev. A* **79**, 043618 (2009).

47. M Caracanhas, et al., Self-similar expansion of the density profile in a turbulent bose-einstein condensate. *J. Low Temp. Phys.* **166**, 49-58 (2011).

48. N Navon, AL Gaunt, RP Smith, Z Hadzibabic, Emergence of a turbulent cascade in a quantum gas. *Nature* **539**, 72-75 (2016).

49. PES Tavares, et al., Matter wave speckle observed in an out-of-equilibrium quantum fluid. *Proc. Natl. Acad. Sci.* **114**, 12691-12695 (2017).

50. HAJ Middleton-Spencer, et al., Strong quantum turbulence in bose-einstein condensates. *Phys. Rev. Res.* **5**, 043081 (2023).

51. MC Tsatsos, et al., Quantum turbulence in trapped atomic Bose-Einstein condensates. *Phys. Rep.* **622**, 1-52 (2016).

52. L Madeira, M Caracanhas, F dos Santos, V Bagnato, Quantum turbulence in quantum gases. *Annu. Rev. Condens. Matter Phys.* **11**, 37-56 (2020).

53. K Hossain, S Gupta, MM Forbes, Detecting entrainment in Fermi-Bose mixtures. *Phys. Rev. A* **105**, 063315 (2022).

54. M Tsubota, K Fujimoto, S Yui, Numerical studies of quantum turbulence. *J. Low Temp. Phys.* **188**, 119-189 (2017).

55. C Pethick, H Smith, *Bose-Einstein Condensation in Dilute Gases*. (Cambridge University Press), (2002).

56. LP Pitaevskii, S Stringari, *Bose-Einstein Condensation*, International Series of Monographs on Physics. (Oxford University Press, Oxford) Vol. 116, (2003).

57. K Hossain, et al., Rotating quantum turbulence in the unitary Fermi gas. *Phys. Rev. A* **105**, 013304 (2022).

58. MM Forbes, R Sharma, Validating simple dynamical simulations of the unitary Fermi gas. *Phys. Rev. A* **90**, 043638 (2014).

59. M Kobayashi, et al., Quantum turbulence simulations using the Gross-Pitaevskii equation: High-performance computing and new numerical benchmarks. *Comp. Phys. Comm.* **258**, 107579 (2021).

60. RW Hamming, *Numerical Methods for Scientists and Engineers*. (McGraw-Hill, Inc., New York, NY, USA), (1973).

61. S Jin, KJ Roche, I Stetcu, I Abdurrahman, A Bulgac, The LISE package: Solvers for static and time-dependent superfluid local density approximation equations in three dimensions. *Comput. Phys. Commun.* **269**, 108130 (2021).

62. A Bulgac, I Abdurrahman, G Włazłowski, Sensitivity of time-dependent density functional theory to initial conditions. *Phys. Rev. C* **105**, 044601 (2022).

63. W-SLDA Toolkit ([wslda.fizyka.pw.edu.pl/](http://wslda.fizyka.pw.edu.pl/)) (2023).

64. A Marek, et al., The elpa library: scalable parallel eigenvalue solutions for electronic structure theory and computational science. *J. Physics: Condens. Matter* **26**, 213201 (2014) see also [elpa.mpcdf.mpg.de](http://elpa.mpcdf.mpg.de).

65. P Kus, et al., Optimizations of the eigensolvers in the elpa library. *Parallel Comput.* **85**, 167-177 (2019).

66. P Kus, H Lederer, A Marek, GPU optimization of large-scale eigenvalue solver. *Lect. Notes Comput. Sci. Eng.* (2017).

67. VWz Yu, et al., GPU-acceleration of the ELPA2 distributed eigensolver for dense symmetric and hermitian eigenproblems. *Comput. Phys. Commun.* **262**, 107808 (2021).

68. T Auckenthaler, et al., Parallel solution of partial symmetric eigenvalue problems from electronic structure calculations. *Parallel Comput.* **37**, 783-794 (2011) 6th International Workshop on Parallel Matrix Algorithms and Applications (PMAA'10).

69. S Babuin, E Varga, I Skrbek, E Lévéque, PE Roche, Effective viscosity in quantum turbulence: A steady-state approach. *EPL (Europhysics Lett.)* **106**, 24006 (2014).

70. V Shukla, PD Mininni, G Krstulovic, PC di Leoni, ME Brachet, Quantitative estimation of effective viscosity in quantum turbulence. *Phys. Rev. A* **99**, 043605 (2019).

71. M Leadbatter, T Winiecki, DC Samuels, CF Barenghi, CS Adams, Sound emission due to superfluid vortex reconnections. *Phys. Rev. Lett.* **86**, 1410-1413 (2001).

72. A Villois, D Proment, G Krstulovic, Irreversible dynamics of vortex reconnections in quantum fluids. *Phys. Rev. Lett.* **125**, 164501 (2020).

73. S Nazarenko, *Wave Turbulence*, Lecture Notes in Physics. (Springer Berlin Heidelberg), (2011).

74. V Zakharov, V L'vov, G Falkovich, *Kolmogorov Spectra of Turbulence I: Wave Turbulence*, Springer Series in Nonlinear Dynamics. (Springer Berlin Heidelberg), (2012).

75. P Walmsley, D Zmeev, F Pakpour, A Golov, Dynamics of quantum turbulence of different spectra. *Proc. Natl. Acad. Sci.* **111**, 4691-4698 (2014).

76. AW Baggaley, J Laurie, CF Barenghi, Vortex-density fluctuations, energy spectra, and vortical regions in superfluid turbulence. *Phys. Rev. Lett.* **109**, 205304 (2012).

77. AW Baggaley, CF Barenghi, A Shukurov, YA Sergeev, Coherent vortex structures in quantum turbulence. *EPL (Europhysics Lett.)* **98**, 26002 (2012).

78. C Nore, M Abid, ME Brachet, Kolmogorov turbulence in low-temperature superflows. *Phys. Rev. Lett.* **78**, 3896-3899 (1997).

79. SN Klimin, J Tempere, G Lombardi, JT Devreese, Finite temperature effective field theory and two-band superfluidity in fermi gases. *Eur. Phys. J. B* **88**, 122 (2015).

80. MA Silaev, Universal mechanism of dissipation in Fermi superfluids at ultralow temperatures. *Phys. Rev. Lett.* **108**, 045303 (2012).

81. G Colò, Nuclear density functional theory. *Adv. Physics: X* **5**, 1740061 (2020).

82. L Amico, et al., Colloquium: Atomtronic circuits: From many-body physics to quantum technologies. *Rev. Mod. Phys.* **94**, 041001 (2022).

83. G Włazłowski, MM Forbes, SR Sarkar, A Marek, M Szpindler, Quantum turbulence in superfluid Fermi gas: results of numerical simulation (2023) Available at [doi.org/10.5281/zenodo.8355244](https://doi.org/10.5281/zenodo.8355244).

84. N Schunck, ed., *Energy Density Functional Methods for Atomic Nuclei*. (IOP Publishing), (2019).

85. A Boulet, G Włazłowski, P Magierski, Local energy density functional for superfluid Fermi gases from effective field theory. *Phys. Rev. A* **106**, 013306 (2022).

86. P Magierski, J Grineviciute, K Sekizawa, Pairing dynamics and time-dependent density functional theory. *Acta Phys. Pol. B* **49**, 281 (2018).

87. A Bulgac, MM Forbes, Use of the discrete variable representation basis in nuclear physics. *Phys. Rev. C* **87**, 051301 (2013).

88. H Childs, et al., Visit: An end-user tool for visualizing and analyzing very large data in *High Performance Visualization-Enabling Extreme-Scale Scientific Insight*. pp. 357-372 (2012).

**Supplementary Material for:****Fermionic Quantum Turbulence: Pushing the Limits of High-Performance Computing****1. Raw Data**

We provide raw data through Zenodo repository (83). The datasets include:

1. dataset with results for BCS run ( $ak_F = -1.8$ ) obtained with TDSLDA method (`bcs` folder).
2. results of static calculations for single vortex line for BCS run ( $ak_F = -1.8$ ) obtained with SLDA method (`bcs_static_vortex` folder).
3. dataset with results for UFG run ( $ak_F = \infty$ ) obtained with TDSLDA method (`ufg` folder).
4. dataset with results for UFG run ( $ak_F = \infty$ ) obtained with GPE method. The results are provided for two dissipation coefficients  $\eta = 0.01$  and  $0.08$  (`ufg_gpe` folder)

The files contain the following information:

- Time evolution of the density  $n(r, t)$ , the current  $j(r, t)$  and the order parameter  $\Delta(r, t)$ .
- Positions of extracted vortex lines from time-dependent runs.
- Reproducibility packs for TDSLDA runs. They contain the full information about the settings and the computation process.
- Static solutions for the single vortex in the BCS regime ( $ak_F = -1.8$ ).

In the repository, there is `README.md` file that provides information about used data formats. There are also provided example scripts/codes demonstrating how to read data in python and C.

**2. List of Movies**

Below, we provide the list of accompanying movies (in mp4 format). 3D views were created by VisIt software (88). The visualizations presented the volume distribution of the order parameter and lines indicating the vortex cores' location. All movies are also accessible on YouTube.

**Supplementary movie 1** Run for the unitary Fermi gas ( $ak_F = \infty$ ) by means of TDSLDA.

YouTube: [youtu.be/hOLPmPVQ4xo](https://youtu.be/hOLPmPVQ4xo)

**Supplementary movie 2** Run for the BCS regime ( $ak_F = -1.8$ ) by means of TDSLDA.

YouTube: [youtu.be/OplDOPdaPKM](https://youtu.be/OplDOPdaPKM)

**Supplementary movie 3** Run for the unitary Fermi gas ( $ak_F = \infty$ ) by means of GPE, with the dissipation coefficient  $\eta = 0.01$ .

YouTube: [youtu.be/3Dux6PHD4e4](https://youtu.be/3Dux6PHD4e4)

**Supplementary movie 4** Run for the unitary Fermi gas ( $ak_F = \infty$ ) by means of GPE, with the dissipation coefficient  $\eta = 0.081$ .

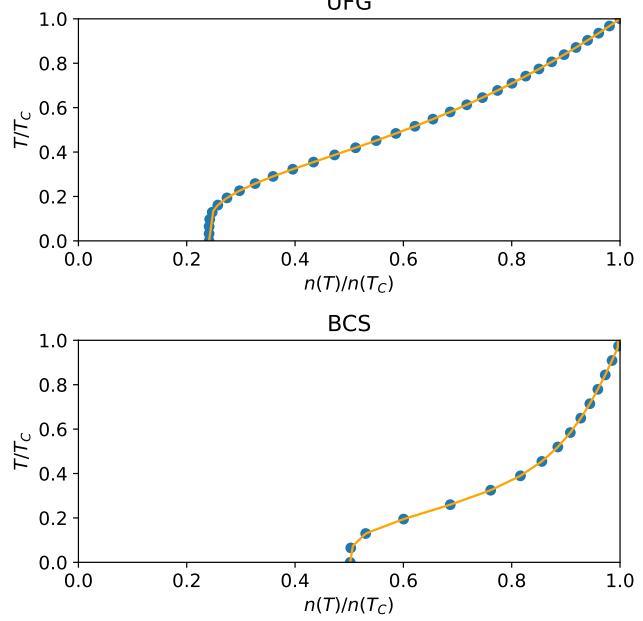
YouTube: [youtu.be/7FXI-NIAq20](https://youtu.be/7FXI-NIAq20)

**Supplementary movie 5** Dynamics of vortex lines and their local effective temperatures for the unitary Fermi gas ( $ak_F = \infty$ ).

YouTube: [youtu.be/Wd\\_gEUZcbu4](https://youtu.be/Wd_gEUZcbu4)

**Supplementary movie 6** Dynamics of vortex lines and their local effective temperatures for the BCS regime ( $ak_F = -1.8$ ).

YouTube: [youtu.be/zGSaVhT3e74](https://youtu.be/zGSaVhT3e74)



**Fig. 9.** The relation between the vortex core density and the gas temperature for UFG (top) and BCS (bottom) regimes. These calibration curves were used to assign the local effective temperatures, as presented in Fig. 8.

**3. Thermometer calibration curve**

In Fig. 9, we provide the relation between vortex core density (normalized to the bulk density) and gas temperature. These curves were used for assign the local temperature of vortex cores, as shown in movies **Supplementary movie 5** and **Supplementary movie 6**.