Reply to Referee 2

Comment B.1: *The Kapitza-Dirac results in Figure 3c would benefit from a quantitative comparison to the standard method of integrating the Schrodinger equation in momentum space, as alluded to in the main text of the paper. This should show agreement with the simulation results for the a = 0 simulations at early times. This would provide both a quantitative check for the numerical methods, and show in what ways the simple method of integrating the Schrodinger equation breaks down at later times and for larger interaction strengths.*

Response: We appreciate the Referee's comment and agree that the manuscript would benefit from a direct comparison with analytical approaches. We implemented the Hamiltonian for Kapitza-Dirac diffraction in momentum space, diagonalized it, and time-evolved the BEC. In figure 1, we present the results for a momentum matrix with a cut-off at $\pm 4\hbar k$, plotted against the numerical results obtained with `TorchGPE`. The results show good agreement with analytical predictions, with divergence occurring as expected for longer time evolutions.

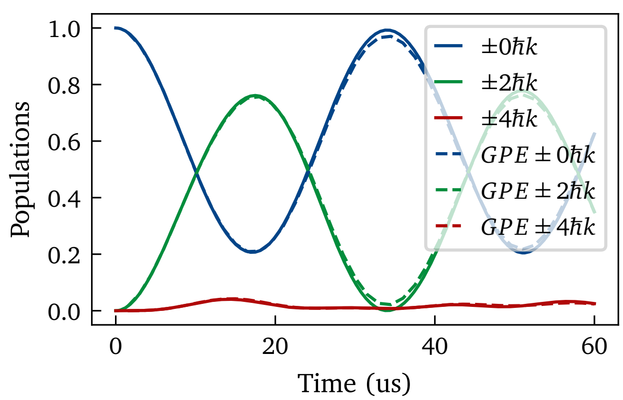

FIG. 1. Time evolution of different momentum mode populations during Kapitza-Dirac diffraction, performed by integrating the Schrödinger equation in momentum space with a cut-off at $\pm 4\hbar k$. For a comparison we add numerical results for a non-interacting gas obtained with `TorchGPE`.

Changes:

- We added a comment in the caption of figure 3.

Comment B.2: *It could be a good addition to include simulations of just a BEC in a harmonic trap before considering more complicated scenarios. This would allow one to see convergence to the analytic Thomas-Fermi distribution as one increases the atom number, for example.*

Response: We thank the Referee for this comment. We agree that – before moving to complex problems – such a framework needs to be thoroughly benchmarked against known solutions. The most straight-forward problem to apply our code to is, as the Referee suggests, to calculate the density distribution of an interacting Bose gas in a harmonic trap, since this result can be directly compared to the Thomas-Fermi approximation for the Gross-Pitaevskii equation.

The suggested comparison constitutes the first example in the "TorchGPE quickstart" section of the documentation of our software package, see TorchGPQ quickstart. The result (see also Figure 2 in this response) shows perfect agreement between the analytic solution and the result of the imaginary time evolution done by our software. We now included this example also as a new subsection in the section about benchmarking our code.

Changes:

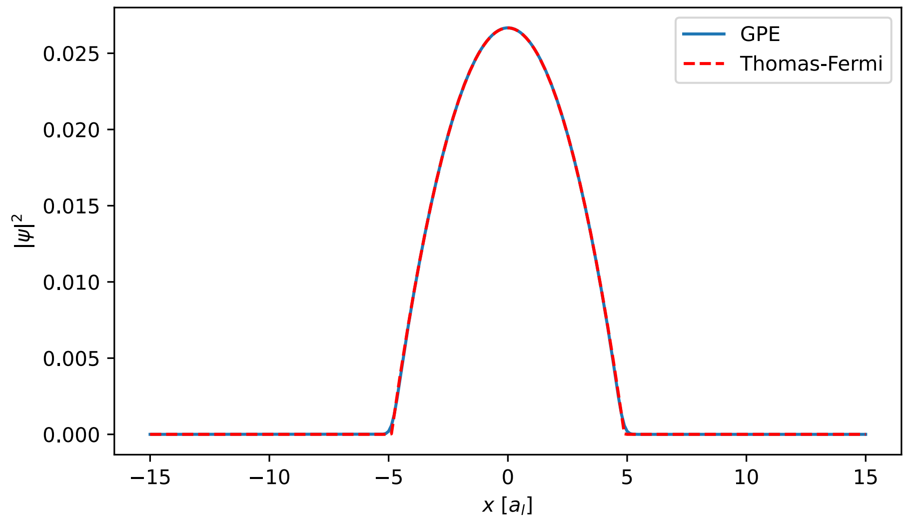

FIG. 2. Density distribution of a Bose-Einstein condensate in a harmonic trap calculated analytically using the Thomas-Fermi approximation (red dashed line) compared to the results of an imaginary time evolution using TorchGPE (blue line). We find perfect agreement between both results except for the edge of the density ditribution where only the numerical simulation is able to capture the smoothening effect of the kinetic energy which is neglected in the analytical approximation. This figure is taken from TorchGPE quickstart.

- We added a new paragraph about benchmarking the code at the example of a harmonically trapped BEC to the beginning of the section 5 "Benchmarking".

Comment B.3: *Please specify what κ is (presumably the cavity loss rate) starting in Eq. 14.*

Response: We thank the Referee for pointing us to this missing definition of κ as the cavity field decay rate.

Changes:

- We now added this definition below Equation (14).

Comment B.4: *Please specify the atom number N used for the simulations in Figure 4 so that the photon counts may be verified.*

Response: We thank the Referee for spotting the absence of this simulation parameter.

Changes:

- We added the number of particles in the caption of Figure 4.

Comment B.5: *Please comment on the differences between the tested GPU models. What is different about the hardware of these GPUs? Did the hardware specifications make any significant differences in performance, possibly beyond what is seen in the Figure 2 comparison to the CPU?*

Response: While a detailed quantitative analysis of all factors influencing a GPU model's performance is beyond the scope of this section — whose primary goal is to present typical execution times for commonly accessible hardware — we agree that it is valuable to provide some intuition regarding which GPU parameters impact the final execution time. This topic is addressed in the main text, where we attribute the change in the scaling properties of GPU runtime to memory saturation. This hypothesis is supported by the observation that when the data type is changed from Float64 to Float32, the GPU can handle grids twice as large before exhibiting a shift in runtime scaling behavior.

In response to the Referee's suggestion, we have highlighted the key characteristics of a GPU that most significantly affect the code's performance.

Changes:

- We added comments in Section 4.1 highlighting which specifications of the benchmarked processors can impact the code's performance.
- We correlated the point where the scaling behavior of the processors changes to the available VRAM memory.