Computer Physics Communications 208 (2016) 92–102

Contents lists available at ScienceDirect

Computer Physics Communications

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

A tractable prescription for large-scale free flight expansion of wavefunctions

P. Deuar

Institute of Physics, Polish Academy of Sciences, Al. Lotników 32/46, 02-668 Warsaw, Poland

ARTICLE INFO

Article history:
Received 10 February 2016
Received in revised form
6 July 2016
Accepted 6 August 2016
Available online 12 August 2016

Keywords:
Discrete Fourier transform
Ultracold atoms
Free flight evolution
Time of flight imaging
Far-field image
Solitons
Wavefunction
Classical field

ABSTRACT

A numerical recipe is given for obtaining the density image of an initially compact quantum mechanical wavefunction that has expanded by a large but finite factor under free flight. The recipe given avoids the memory storage problems that plague this type of calculation by reducing the problem to the sum of a number of fast Fourier transforms carried out on the relatively small initial lattice. The final expanded state is given exactly on a coarser magnified grid with the same number of points as the initial state. An important application of this technique is the simulation of measured time-of-flight images in ultracold atom experiments, especially when the initial clouds contain superfluid defects. It is shown that such a finite-time expansion, rather than a far-field approximation is essential to correctly predict images of defect-laden clouds, even for long flight times. Examples shown are: an expanding quasicondensate with soliton defects and a matter-wave interferometer in 3D.

1. Introduction

The free flight expansion of a quantum wavefunction, though physically very simple, is often a troublesome computational problem if the state that is required is not quite yet in the far field regime. The snag is that a computational lattice that both resolves the small initial cloud and encompasses the large expanded cloud can be prohibitively large. Here, it will be described how to overcome this while still using standard discrete fast Fourier transform (FFT) tools.

For example, this is commonly desired when simulating experiments in ultracold atoms. A ubiquitous experimental procedure in this field is the release of the atoms from a trap and the subsequent observation of the density of a strongly expanded cloud. Given that the imaged expanded cloud is usually much larger than the initial one pre-release, the observed expanded atom density corresponds approximately to the velocity distribution in the initial cloud. More precisely, it corresponds to the velocity distribution that is formed early on after release, when the interatomic interaction energy has been converted into kinetic energy. This is the picture that is often used to interpret the data.

This interpretation assumes that the detection is occurring in the far field where all structure is large compared to the initial cloud. However, in practice this is often not a good enough approximation, particularly if one is interested in fine structure inside the atomic cloud, such as defects or interparticle correlations. The reality is that the expansion is usually by a factor of tens or hundreds, so that interesting features such as defects or correlations that are of the order of 10% or 1% of the initial cloud in size have not yet attained a velocity profile at the time of detection. They are already distorted from their spatial profile in the initial cloud, but their shape has not yet stabilized to its far field form. Some examples where a long but not quite far-field expansion occurs include the interference pattern generated after release of a pair of elongated clouds [1,2], the study of Hanbury Brown–Twiss correlations in expanding clouds [3,4] and two-particle correlations in a halo of supersonically scattered atoms [5,6].

The basic numerical task here is to predict the detected density image based on whatever model we are using for the atomic field. For excited or thermal gases an ensemble of classical field [7–13] or truncated Wigner wavefunctions [14–17] are often used. The straightforward approach is to place the whole field $\Psi(\mathbf{x}, t)$ from the outset in an x-space large enough to accommodate the whole expansion. However, it is rarely technically feasible to carry out the entire expansion by this method in three dimensions despite the seemingly trivial physics. A good description of the initial state in a three-dimensional lattice can easily require $\mathcal{O}(10^5)$ lattice points or more, and an expansion by a factor of 10–100 in each direction would lead to 10^8–10^{11} lattice points. This is either intractable or

E-mail address: deuar@ifpan.edu.pl.

http://dx.doi.org/10.1016/j.cpc.2016.08.004
0010-4655/

impractical on a simple computer, and even more so for studies of defect statistics or correlations which require surveys of hundreds or thousands of realizations.

Barring access to supercomputing resources, a standard resort in this case is to make a somewhat unsavory compromise: Simulate the expansion as far in time as the computer allows, and assume that the neglected later changes are not qualitatively important. It will be shown here how to avoid this while still using standard discrete FFT tools on a simple computer.

In Section 2 the basics of the problem are described, and in Section 2.5 an estimate is made of the time regime over which a careful exact expansion of clouds with defects is necessary. Section 3 demonstrates this with an example. The numerical difficulty is examined in more detail in Section 4, and the solution derived in Section 5.1. The paper concludes in Section 7 with some discussion of practicalities and applications.

The prescription that constitutes the main result is briefly given in a stand-alone form in Section 5.2.

2. The matter at hand

The aim is to calculate what is actually measured, the spatial density distribution at the detector, $\rho(\mathbf{r}, t_{\text{final}})$. We assume that just before the trap is switched off at $t = 0$, the trapped state is described by a complex field $\Psi_0(\mathbf{r})$ that has the form of a single-particle wavefunction.

2.1. Conversion phase

Typically the expansion can be considered as consisting of two phases: An initial "conversion" phase during which the interaction energy between the atoms is converted to kinetic energy, and later free flight of the atoms. Since the interaction energy per particle is proportional to the density, an expansion in three-dimensional space by a factor of two in size will reduce this interaction energy per particle by a factor of eight. This initial expansion can be done in a straightforward way until interactions are diluted away to become negligible. One just takes a computational lattice $\bar{x}$ in x-space that is 2–4 times wider than the initial cloud Ψ_0, and evolves on that. In ultracold atoms, the workhorse Gross–Pitaevskii Equation (GPE) is typically used—see the example in Section 6.

The end result of this phase (at time t_s, say) is that we have a partly expanded wavefunction $\Psi(\mathbf{r}, t_s)$. Numerically, it is described as a table of complex numbers $\Psi_{\mathbf{n}}$ indexed by the set of integers $\mathbf{n} = \{n_1, \ldots, n_d\}$ in d-dimensional space, that enumerate the points on the numerical lattice. The lattice spacings are $\Delta x_j = L_j/M_j$ where L_j is the length of the box in the jth direction, and M_j the corresponding number of lattice points, so that $n_j = 0, 1, \ldots, (M_j - 1)$. The lattice positions are

$$x_j = a_j + \Delta x_j n_j \tag{1}$$

with offsets a_j. i.e. $\mathbf{x} = \mathbf{a} + \Delta\mathbf{x} \cdot \mathbf{n}$.

Note that the wavefunctions $\Psi(\mathbf{r})$ are not in general the complete quantum many-particle wavefunction unless the particles are non-interacting. For interacting particles, one usually works with Ψ in some kind of c-field approximation [10–13,15,16].

2.2. Lattice notation

Several numerical lattices will appear in what follows. The following notation will be used:

- Quantities with a tilde, such as $\widetilde{\Psi}(\mathbf{k})$, are in k-space.
- Bold quantities are vectors in d dimensions (usually $d = 3$), with elements indexed by j as in $\mathbf{x} = \{x_1, \ldots, x_j \ldots, x_d\}$.
- Undecorated quantities, such as $\Psi(\mathbf{x})$ denote the lattice used to represent the starting state at t_s. This has a manageable number of points, M.
- Barred quantities, such as $\overline{\Psi}(\overline{\mathbf{x}})$ will be on a magnified numerical lattice $\overline{\mathbf{x}}$ that can describe the expanded state, but is too coarse to describe the starting state at t_s.
- Underlined quantities such as $\underline{\Psi}(\underline{\mathbf{x}})$ will be used to denote a sufficiently huge lattice that both encompasses the expanded state at t_{final} and resolves the starting state at t_s, when this lattice is different from the starting undecorated one.
- The position coordinate $\mathbf{r}$ is a continuum quantity, as opposed to $\mathbf{x}$ which are corresponding lattice positions. Similarly, $\boldsymbol{\kappa}$ is a continuum wavevector, while $\mathbf{k}$ is discretized.

2.3. Free flight into the far field

The remaining evolution after t_s, the "starting time", is just free flight. Each particle of momentum $\hbar\boldsymbol{\kappa}$ flies a distance $\hbar\boldsymbol{\kappa}\, t_{\text{flight}}/m$, where the flight time is

$$t_{\text{flight}} = t_{\text{final}} - t_s. \tag{2}$$

Then with a far field assumption, i.e. that the initial starting position at t_s is irrelevant because they have flown so far, the position of a particle is

$$\mathbf{r} = \hbar t_{\text{flight}}\boldsymbol{\kappa}/m. \tag{3}$$

An estimate for the final density can then be obtained from $\widetilde{\Psi}(\boldsymbol{\kappa}, t_s)$, the momentum-space wavefunction at the end of the conversion phase. It is:

$$\rho_{\text{ff}}(\mathbf{r}) = |\Psi_{\text{ff}}(\mathbf{r})|^2 = \left(\frac{m}{\hbar t_{\text{flight}}}\right)^d \left|\widetilde{\Psi}\left(\frac{m\mathbf{r}}{\hbar t_{\text{flight}}}, t_s\right)\right|^2. \tag{4}$$

The prefactor is for normalization purposes, so that $\int d^d\boldsymbol{\kappa}\, |\widetilde{\Psi}(\boldsymbol{\kappa})|^2 = \int d^d\mathbf{r}\, |\Psi_{\text{ff}}(\mathbf{r})|^2$. Notably, this discards any phase information. However, the usual imaging in ultracold atom experiments is insensitive to that.

The starting momentum wavefunction $\widetilde{\Psi}(\boldsymbol{\kappa})$ at t_s is obtained with a norm preserving Fourier transform:

$$\widetilde{\Psi}(\boldsymbol{\kappa}) = \frac{1}{(2\pi)^{d/2}} \int d^d\mathbf{r}\, e^{-i\boldsymbol{\kappa}\cdot\mathbf{r}}\, \Psi(\mathbf{r}). \tag{5}$$

Numerically, the conversion is best made with a discrete Fourier transform (DFT). The DFT of a field Q is

$$\widetilde{Q}_{\widetilde{\mathbf{m}}} = \text{DFT}\,[Q_{\mathbf{n}}]_{\widetilde{\mathbf{m}}} = \sum_{\mathbf{n}} Q_{\mathbf{n}} \exp\left[-i\sum_{j=1}^{d} \frac{2\pi}{M_j} n_j \widetilde{m}_j\right] \tag{6}$$

with indices $\widetilde{m}_j = 0, 1, \ldots, (M_j - 1)$ labeling the position on the k-space lattice which has spacing $\Delta k_j = 2\pi/L_j$. The sum is over the whole $\mathbf{n}$ range.

In what follows, we will always be using the physical free-space wavevectors $\mathbf{k}_{\widetilde{\mathbf{m}}}$ ordered as:

$$k_j(\widetilde{m}_j) = \Delta k_j \widetilde{l}_j = \Delta k_j \times \begin{cases} \widetilde{m}_j & \widetilde{m}_j < M_j/2 \\ \widetilde{m}_j - M_j & \widetilde{m}_j \geq M_j/2. \end{cases} \tag{7}$$

The integer multipliers can also be written as $\widetilde{l}_j = \text{mod}\,[\widetilde{m}_j + \frac{1}{2}M_j, M_j] - \frac{1}{2}M_j$. For simple transformations such as (5) and (8), a set of monotonically ordered non-negative wavevectors $\widetilde{\mathbf{m}} \cdot \Delta\mathbf{k}$ is equivalent operationally to (7) because $\Delta k_j M_j (x_j - a_j)$ is an integer multiple of 2π. However, such equivalence no longer holds for calculating the kinetic energy or upon changing the lattice offset $\mathbf{a}$, both of which will be needed for expansion.

Using (5) and the DFT (6), as well as taking care of a possible offset in (1), the discrete momentum wavefunction at t_s is

$$\widetilde{\Psi}_{\tilde{\mathbf{m}}}(t_s) = \frac{\Delta V}{(2\pi)^{d/2}} e^{-i\mathbf{a}\cdot\mathbf{k}_{\tilde{\mathbf{m}}}} \,\mathrm{DFT}\left[\Psi_{\mathbf{n}}(t_s)\right]_{\tilde{\mathbf{m}}} \tag{8}$$

with lattice point volume $\Delta V = \prod_j \Delta x_j$. This is readily implemented using standard fast Fourier transform (FFT) libraries [18,19].

In most cases in the literature, the short initial expansion to t_s and conversion (4) to obtain the detected density is all that is done. This is fine provided that we are only interested in momentum differences $\delta\mathbf{k}$ much larger than those corresponding to the width W_s of the converted cloud at t_s. That is, when $|\delta\mathbf{k}| \gg mW_s/\hbar t_{\mathrm{flight}}$. Or, alternatively, that we are only interested in spatial resolutions $\gg W_s$ in the final expanded cloud.

2.4. Free flight without a far field assumption

To avoid making the rather uncontrolled far field assumption (4), and get results for a well defined final time t_{final}, consider first that in principle, the free flight evolution of the wavefunction in k-space is straightforward:

$$\widetilde{\Psi}(\mathbf{k}, t_{\mathrm{final}}) = \widetilde{\Psi}(\mathbf{k}, t_s) \exp\left[-it_{\mathrm{flight}} \frac{\hbar|\mathbf{k}|^2}{2m}\right]. \tag{9}$$

In principle, all one then needs to obtain $\rho(\mathbf{r}, t_{\mathrm{final}})$ is to inverse Fourier transform $\widetilde{\Psi}(\mathbf{k}, t_{\mathrm{final}})$ back into x-space. Generally:

$$\Psi(\mathbf{r}, t_{\mathrm{final}}) = \frac{1}{(2\pi)^{d/2}} \int d^d\boldsymbol{\kappa}\, e^{i\boldsymbol{\kappa}\cdot\mathbf{r}}\, \widetilde{\Psi}(\boldsymbol{\kappa}, t_{\mathrm{final}}). \tag{10}$$

The discrete implementation like in (6) and (8) is

$$Q_{\mathbf{n}} = \mathrm{DFT}^{-1}\left[Q_{\tilde{\mathbf{m}}}\right]_{\mathbf{n}} = \frac{1}{M} \sum_{\tilde{\mathbf{m}}} \widetilde{Q}_{\tilde{\mathbf{m}}} \exp\left[i \sum_{j=1}^{d} \frac{2\pi}{M_j} n_j \tilde{m}_j\right]. \tag{11}$$

$M = \prod_j M_j$ is the overall lattice size. With volume $V = \prod_j L_j$,

$$\Psi_{\mathbf{n}}(t_{\mathrm{final}}) = \frac{(2\pi)^{d/2}}{V} \mathrm{DFT}^{-1}\left[\widetilde{\Psi}_{\tilde{\mathbf{m}}}(t_{\mathrm{final}}) e^{i\mathbf{a}\cdot\mathbf{k}_{\tilde{\mathbf{m}}}}\right]_{\mathbf{n}}. \tag{12}$$

This step can, however, be hard on computational resources, even with an FFT because a very large lattice M is often needed to fully describe the final time state $\widetilde{\Psi}(t_{\mathrm{final}})$.

2.5. Continued defect dynamics during free-field expansion

Expansions of clouds containing narrow mobile defects are a popular experimental topic in recent years [1,20–25]. These are systems for which the transition from the starting t_s state to the far-field is nontrivial.

Let the overall width of the cloud at t_s in a chosen direction be W_s, and consider defects of width $\mathrm{w_{def}} \ll W_s$ and typical speed u. Speed differences between different defects are then also $\approx u$. The velocity distribution in the gas, however, is dominated by the shortest length scale in the system. This is usually given by the half-width of individual defects, giving a typical velocity $\sigma_v \approx 2\hbar/m\mathrm{w_{def}}$. After significant expansion, the width of the cloud will be $W_{\mathrm{final}} \approx 2t_{\mathrm{flight}}\sigma_v = 4\hbar t_{\mathrm{flight}}/m\mathrm{w_{def}}$.

It takes a time $t_v = W_s/\sigma_v$ for a rough semblance of the velocity/momentum distribution to form in real space due to expansion (this is the time for a typical particle to move across the initial cloud). Remnants of defects can continue to rearrange until a time when their relative speed would allow them to move as far as $\approx W_s$, which is a typical initial spacing between them:

$$t_{\mathrm{arrange}} \approx \frac{W_s}{u} \gg t_{\mathrm{v}}. \tag{13}$$

For clearly recognizable defects to be present, one should have defects slower than particles: $u \ll \sigma_v$. Due to this slowness, there is a period $t_v \ll t_{\mathrm{flight}} \ll t_{\mathrm{arrange}}$ during which complicated rearrangement of defect remnants can take place even though the gross shape of the cloud already resembles the far-field velocity distribution. The simple far field expression (4) is not appropriate during this time.

This is not an uncommon situation in ultracold atom experiments, and has relevance for interpretation of experimental data. For an initially trapped thermal gas in a classical field regime where its dynamics is quite well described by the Gross–Pitaevskii equation [15,26,27,10,28,29], typical defects are solitons in 1D and vortices in 2D. In this regime, when g is the s-wave scattering length and ρ the typical density, the chemical potential is $\mu \approx g\rho$, giving defect width $\mathrm{w_{def}} \approx 2\hbar/\sqrt{m\mu}$ and a speed of sound $c = \sqrt{\mu/m}$. Major defects are much slower, i.e. $u = \epsilon c$ with $\epsilon \ll 1$. With a trap frequency of ω, the initial cloud width is $W_s \approx (2/\omega)\sqrt{2\mu/m}$. This lets us estimate $t_v = \sqrt{8}/\omega$ and t_{arrange}, and one finds that the times t_{flight} during which rearrangement is still taking place in a cloud that looks to be already far-field in its gross features is

$$t_v \lesssim t_{\mathrm{flight}} \lesssim t_{\mathrm{arrange}}, \qquad \text{i.e.} \qquad 1 < \frac{\omega t_{\mathrm{flight}}}{\sqrt{8}} \lesssim \frac{1}{\epsilon}. \tag{14}$$

This can be a significant period.

3. Example: soliton dynamics during expansion

Let us consider an example of complicated free evolution even at times that would naively be considered to be in the far-field: the expansion of an elongated 1D ultracold gas in the quasi-1D regime. We take physical parameters like in a series of recent experiments [1,2], where clouds in the classical field regime were prepared. An initial c-field state of the 1d system is generated at temperature $T = 80$ nK $= 260\hbar\omega/k_B$ using the stochastic Gross–Pitaevskii equation

$$\hbar\frac{\partial\Psi(x,\tau)}{\partial\tau} = -i(1-i\gamma)\left[-\frac{\hbar^2}{2m}\nabla^2 + g_1|\Psi(x,\tau)|^2 - \mu\right]\Psi(x,\tau) + \sqrt{2\gamma\hbar k_B T}\,\eta(x,\tau) \tag{15}$$

by taking a sample of the field $\Psi_{\mathrm{ic}}(x) = \Psi(x,\tau)$ at $\tau = 10/\omega$, once the ergodic ensemble is reached. The simulation grows the field from the vacuum $\psi(x,0) = 0$. Here, $g_1 = 0.54\hbar\omega a_{\mathrm{ho}}$ is the 1D s-wave scattering length for ^{87}Rb in terms of the harmonic oscillator length $a_{\mathrm{ho}} = \sqrt{\hbar/m\omega}$. The bath coupling $\gamma = 0.02$ has a typical value, $\mu = 90\hbar\omega$ is a chemical potential chosen to give $N = 3000$ particles on average in the stationary ensemble, and $\eta(x,\tau)$ is a Gaussian complex white noise field with variance $\langle\eta(x,\tau)^*\eta(x',\tau')\rangle = \delta(x-x')\delta(\tau-\tau')$. The lattice cutoff in a plane-wave basis is chosen as $\hbar k_{\mathrm{max}} = 0.65\sqrt{2\pi mk_BT}$, according to the optimum values obtained in [29]. The generation of $\Psi_{\mathrm{ic}}(x)$ was carried out using (15) on an initial lattice with $M = 2^{11}$ points and $L = 60a_{\mathrm{ho}}$.

A proper treatment of the conversion phase will be done in the 3D example 6. Here, let us just do an immediate free-field expansion of $\Psi_{\mathrm{ic}}(x)$ from the moment the trap is switched off at $t_s = t_0 = 0$. The 1D density $\rho(x) = |\Psi(x)|^2$ approximates the marginal density of the 3D cloud when integrated over transverse directions. The fully free expansion can be quite a good approximation for a very tight initial trap that has the initial gas in a quasi-1d regime (tight transverse trap frequency $\omega_\perp$). Release causes a very rapid expansion in the transverse directions on a timescale of $1/\omega_\perp$, with width $\propto (1/\omega_\perp)\sqrt{1+\omega_\perp^2 t^2}$ [30]. Accordingly, the density drops as $\sim 1/(1+\omega_\perp^2 t^2)$, and so does the relative strength of interparticle interactions. After a time of

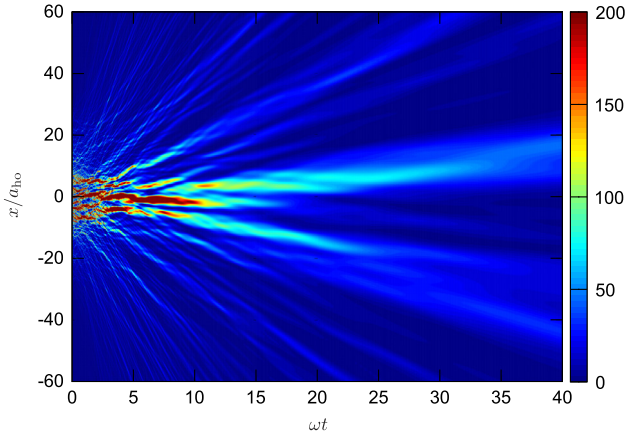

Fig. 1. Evolution of the density $\rho(x) = |\Psi(x)|^2$ (in units of $a_{\rm ho}^{-1}$) after release, calculated directly according to (16), (8) and (12). The vertical scale is narrowed down compared to the computational lattice, with $\underline{L} = 2400a_{\rm ho}$, to show the most interesting region.

several $1/\omega_\perp$ (short compared to t_v), the gas is effectively non-interacting.

The evolution of the field is shown in Fig. 1. Here in 1D, it is easily done directly using the equation

$$\frac{\partial \Psi(x,t)}{\partial t} = i\frac{\hbar}{2m}\nabla^2\Psi(x,t) \tag{16}$$

and the DFTs (8), (12). The initial state $\Psi_{\rm ic}(x)$ was padded with vacuum and evolved on a lattice of $\underline{M} = 81\,920$ points on a simulation region of length $\underline{L} = 2400a_{\rm ho}$, with $\underline{a} = -\underline{L}/2$. The purpose is to see defect evolution during expansion. Indeed, we see that appreciable defect evolution occurs until times of about $\approx 30/\omega$. This can be compared to the values of the crude estimates of (13) obtained for this system when taking $\epsilon = 0.1$: $t_v = 2.8$ for significant expansion and $t_{\rm arrange} = 28$ for end of rearrangement. A very good match.

In Fig. 2, the central section of the expanding cloud is shown for two long times $t = 10/\omega$ and $t = 40/\omega$. The figure also compares to the far-field estimate (4) given by a magnified momentum distribution. The far field estimate is found wanting even at the otherwise very long time $t = 10/\omega$, and becomes only passable at $t = 40/\omega$. For comparison, note that the detection time in the reference experiment [1] was $t_{\rm flight} = 0.65/\omega \ll t_v$, which places it even well before any significant expansion.

4. A closer look at the numerical difficulty

Consider now calculations, e.g. in 3D, for which a lattice that properly describes both the small starting and large expanded state is extremely large.

4.1. Computational effort

Let the energy per particle in the trapped state be εd, i.e. it will be ε per degree of freedom in free space. This is all converted to kinetic energy by the end of the conversion phase at t_s so that a typical wavevector is $k_{\rm typical} = \sqrt{2m\varepsilon}/\hbar$. For good measure, and particularly to allow for energy fluctuations above the mean, one needs to include higher values $k_{\rm max} \approx r_K k_{\rm typical}$ with $r_K \sim 2$. The spacing on the lattice needed to resolve the resulting wavelengths (Nyquist–Shannon theorem) is going to be $\Delta x_{\rm min} \leq \pi/k_{\rm max} = \pi\hbar/(r_K\sqrt{2m\varepsilon})$. Now, when the widths of the starting cloud in the jth direction are W_j, the widths after free flight expansion will be approximately

$$\underline{W_j} = W_j + 2r_K t_{\rm flight}\sqrt{2\varepsilon/m}, \tag{17}$$

allowing again for wavevectors of up to $k_{\rm max}$. We use the underlined notation for final quantities in anticipation of a large lattice. The minimum number of lattice points needed in each direction for the expanded cloud is $\underline{M}_j^{\rm min} = \underline{W}_j/\Delta x_{\rm min} = (r_K/\pi\hbar)[W_j\sqrt{2m\varepsilon} + 4r_K t_{\rm flight}\varepsilon]$. To have an accurate calculation extra padding and resolution usually has to be included, giving $\underline{M}_j \approx r_A \underline{M}_j^{\rm min}$ with another factor $r_A \sim 2$.

After significant expansion, when the W_j have become negligible, the required lattice size approaches $\underline{M}_j \approx 4r_K^2 r_A t_{\rm flight}\varepsilon/\pi\hbar$. Thus the overall required size $\underline{M} = \prod_j \underline{M}_j$ will be

$$\underline{M} \geq C\left(\frac{\varepsilon\, t_{\rm flight}}{\hbar}\right)^d, \tag{18}$$

where $C = (4r_K^2 r_A/\pi)^d \sim \mathcal{O}(10^d)$, which is about a thousand in 3D.

Memory requirements for double precision arithmetic will be $16 \times \underline{M}$ in bytes, while the time to carry out a FFT will scale as $\underline{M}\log\underline{M}$, and the time to do the evolution (9) is $\sim\underline{M}$. Defect experiments tend to have $\varepsilon t_{\rm flight}/\hbar \sim \mathcal{O}(100)$. For example, in the [1] experiment considered as an example here, $\varepsilon t_{\rm flight}/\hbar \approx 60$. While the time for carrying out such an FFT on desktop computers is manageable (of the order of an hour for $\underline{M} = 5 \times 10^9$ on one processor core), the real problem are memory requirements. For $\underline{M} = 5 \times 10^9$, having sufficient RAM memory (75 GB) on a desktop becomes troublesome.

4.2. Information in the wavefunction

What can be done? One can see that the effort involved in (18) is almost all wasted because there is no new information about the cloud in the final state $\underline{\Psi}(\underline{\mathbf{x}})$ that was not already there in the initial $\Psi(\mathbf{k})$. The evolution (9) and DFTs between x-space and k-space (5), (10) are deterministic and reversible. Also, we know that the final not-quite-far-field density $\underline{\rho}(\underline{\mathbf{x}})$ is going to be at least qualitatively similar to (4) which was obtained via a simple magnification of the starting momentum wavefunction (see Fig. 2). This suggests that visible structures will be much broader than in the initial state. The trouble of course is that while the density gets magnified during the free flight by factors

$$\Lambda_j = \frac{\underline{W_j}}{W_j} = 1 + \frac{2r_K t_{\rm flight}\sqrt{2\varepsilon}}{W_j\sqrt{m}}, \tag{19}$$

the velocity remains encoded in a wavelength that remains constant. As long as velocity information is kept, the size of the computational lattice must grow by these same factors Λ_j to keep resolving the largely unchanged phase oscillation. The wastefulness amounts to at least a factor of $\Lambda = \prod_j \Lambda_j$.

Clearly the thing one must do is avoid storing the entire fine lattice of size $\underline{M}$, and abandon knowledge of the properly sampled phase profile at $t_{\rm final}$, leaving only density information on a coarser lattice. The initial converted state $\Psi(\mathbf{x}, t_s)$ can be fully defined on a smaller lattice with

$$M = \prod_j M_j = \frac{\underline{M}}{\Lambda} \approx \left(\frac{r_A r_K\sqrt{2m\varepsilon}}{\pi\hbar}\right)^d \prod_j W_j \tag{20}$$

points and the usual spacings $\Delta\mathbf{x}$, which comes from (18) and (19). The right hand expression assumes $\Lambda_j \gg 1$. A magnification of the initial M lattice by a factor of Λ_j in each direction, while keeping the number of points constant, should be possible in principle without adversely affecting the quality of the final density profile.

Let us first consider an overly simple approach that tries to do this but fails in an instructive way:

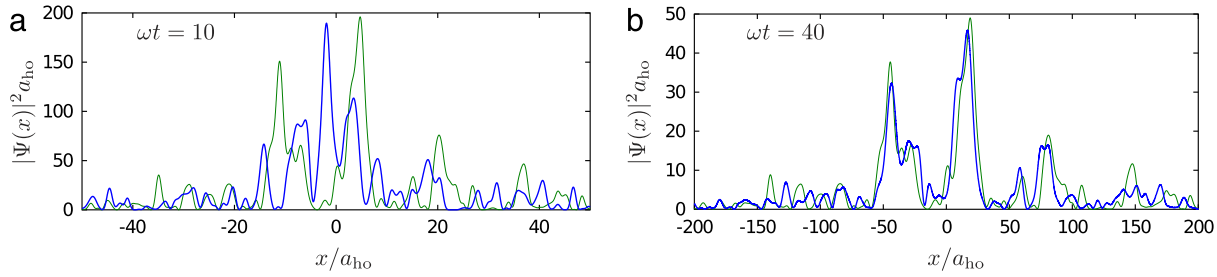

Fig. 2. True density (blue) and its far-field estimate (green) given by (4) after long times of flight. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

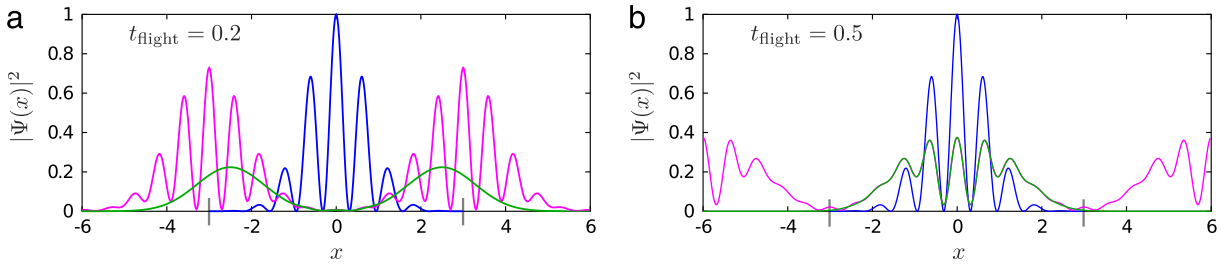

Fig. 3. Initial (blue) and final (magenta) densities after the naive attempt at expansion in 1D with prescription (24). The initial state is $\Psi(x) = \exp(-x^2/2)\cos k_0 x$ with $k_0 = 5$, defined on the range $x \in [-3, 3]$ indicated with gray bars, with $M = 600$. Expansion times are given on the plots. The true final state obtained using (12) on a larger lattice defined on $x \in [-6, 6]$ is shown in green. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

4.3. Naive DFT

Since the positions appear explicitly in (10), it is tempting to proceed as follows:

1. Obtain with the k-space wavefunction at t_s represented as $\widetilde{\Psi}_{\widetilde{\mathbf{m}}}$ on the small M lattice.
2. Apply evolution (9) to obtain $\widetilde{\Psi}(t_{\rm final})_{\widetilde{\mathbf{m}}}$ after whatever time $t_{\rm flight}$ is necessary.
3. Carry out the sum in the return transformation (10) using magnified lattice values of $\overline{\mathbf{x}}$ and automatically keeping the same relatively small number of points, $\overline{M} = M$.

An appropriate magnified lattice would have the same number of points as the starting state: $\overline{M}_j = M_j$, but larger spacing $\Delta\overline{x}_j = \Lambda_j \Delta x_j$, as well as appropriately shifted zero points $\overline{a}_j$. The new positions would be

$$\overline{\mathbf{x}} = \overline{\mathbf{a}} + \boldsymbol{\Lambda} \cdot (\mathbf{x} - \mathbf{a}); \qquad \overline{x}_j = \overline{a}_j + \overline{n}_j \Delta\overline{x}_j \tag{21}$$

indexed by $\overline{n}_j = 0, \ldots, (M_j - 1)$. For step 3, the discrete exponent in (10) is

$$i\mathbf{k} \cdot \overline{\mathbf{x}} = i\mathbf{k} \cdot \overline{\mathbf{a}} + i \sum_{j=1}^{d} \frac{2\pi}{M_j} \widetilde{m}_j \Lambda_j \overline{n}_j. \tag{22}$$

To use the convenient DFT form (11), the factors $\Lambda_j \overline{n}_j$ need to be integers. Hence, the scale factor Λ_j needs to be an integer λ_j for all points on the final $\overline{\mathbf{x}}$ lattice to be calculated this way. Since phases that are a multiple of 2π are equivalent, a value of $\lambda_j \overline{n}_j > M_j$ will lead to the same $\overline{\Psi}$ as one below M_j. This makes any value of $\overline{n}_j$ correspond to an element of a DFT that sums over $\widetilde{m}_j$. Let us define an auxiliary index

$$n_j'' = \mathrm{mod}\left[\overline{n}_j \lambda_j, M_j\right] \tag{23}$$

dependent on $\overline{n}_j$, which indicates the element of the final inverse DFT that should to be used to obtain a particular point on the final $\overline{x}$ lattice. One obtains the following:

$$\overline{\Psi}^{(\rm naive)}_{\mathbf{n}}(t_{\rm final}) = \frac{\Lambda (2\pi)^{d/2}}{V} \mathrm{DFT}^{-1}\left[\widetilde{\Psi}_{\widetilde{\mathbf{m}}}(t_{\rm final}) e^{i\overline{\mathbf{a}}\cdot\mathbf{k}_{\widetilde{\mathbf{m}}}}\right]_{\mathbf{n}''} \tag{24}$$

which is very similar in form to (12), except for the indexing by $\mathbf{n}''$, shift $\overline{\mathbf{a}}$ and Λ prefactor. The last is put in by hand, to keep the amplitude of the original cloud unchanged at t_s upon magnification of the lattice. The numerical effort required by (24) is small, with the largest matrix to be stored only of size M, i.e. $\prod_j \Lambda_j$ times less than the brute force case of (18).

The results for a 1D test case can be seen in Fig. 3. They are not good. This is of course because DFTs correspond to the correct Fourier transform only for the specific relationship between x- and k-space lattices that is described in Section 2.3, and not the wishful one that step 3 implies. What is actually being carried out by (24) instead, is the free evolution for a time $t_{\rm flight}$ of an infinite number of copies of the state $\Psi(\mathbf{x}, t_s)$, repeated in a tiling pattern because of the periodic boundary conditions assumed by the DFT. As soon as the cloud begins to expand, it overlaps around its own edges and everything gets scrambled.

In k-space, the picture is that the flight time is so long that the phase winding caused by (9) advances so much that aliasing occurs. The phase difference between neighboring high-k points is $2\pi\hbar t_{\rm flight} k_{\rm max}/mL_j$ which is

$$\Delta\theta_{\rm max} = \pi\left(\frac{W_j - W_j}{L_j}\right) \tag{25}$$

after substituting for $k_{\rm max}$. As a result, the phase variation with $\mathbf{k}$ is not resolved and the state is scrambled by (24) whenever the starting box L_j is appreciably smaller than the expanded width W_j.

The moral from this naive approach is that the data in the starting $\Psi_{\mathbf{n}}$ lacks the crucial physical information that the area beyond the edges of the starting lattice is supposed to be vacuum.

5. Solving the memory problem

5.1. Derivation

To utilize the physical information about the system that the background over which the cloud expands is vacuum, let us define a buffered starting field $\Phi(\mathbf{r}, t_s)$:

$$\Phi(\mathbf{r}, t_s) = \begin{cases} \Psi(\mathbf{r}, t_s) & \text{if } \forall j \,:\, a_j \le r_j < (a_j + L_j) \\ 0 & \text{otherwise.} \end{cases} \tag{26}$$

Also, to take advantage of the highly optimized FFT algorithms, the exponent in (11) should contain only integer multiples of $2\pi i/M_j$ in each direction j. We will henceforth assume the magnification factors Λ_j to all be positive integers λ_j, as in Section 4.3, which suffices to obtain this condition. The actual cloud need not expand by an integer value, only the numerical lattice. It may also be possible, in principle, to translate cases of fractional Λ_j into an algorithm containing FFTs, when the M_j and Λ_j are appropriately matched. However, this would lead to combinatorial complications in the algorithm. We will refrain from considering this as it does not appear to offer any significant computational advantage.

As usual, the general procedure to obtain a final state is to implement the time evolution in k-space as in (9), then use a DFT to obtain the final x-space wavefunction. This final state is to be on the magnified lattice whose points are

$$\bar{x}_j = \bar{a}_j + \bar{n}_j \lambda_j \Delta x_j, \tag{27}$$

lying $\Delta\bar{x}_j = \lambda_j \Delta x_j$ apart. $\bar{n}_j = 0, \ldots (M_j - 1)$ is the index for the final "small coarse" lattice, like in (21). Its volume is $\bar{V} = \prod_j \bar{L}_j$, a factor of $\lambda = \prod_j \lambda_j$ greater than the initial volume V.

This last immediately presents a problem, because to obtain a final state on an x-space lattice of length $\bar{L}_j$ with an FFT requires transforming a k-space wavefunction that has resolution $2\pi/\bar{L}_j = \Delta k_j/\lambda_j$. This is λ_j times finer than what is available in the starting state $\Psi(\mathbf{x})$. Fortunately, the vacuum assumption (26) provides sufficient information to reconstruct the fine scale structure in k-space, as we will see below.

Time evolution must also occur on this fine lattice, and to remain exact, it must not cut off high momenta, so the huge lattice $\underline{M}$ will be required, at least formally. The required k-space wavefunction of the initial state is, generally, obtained with the transform (5). Discretizing it onto the $\underline{M}$ lattice gives

$$\underline{\widetilde{\Psi}}(\underline{\mathbf{k}}, t_s) = \frac{\Delta V}{(2\pi)^{d/2}} \sum_{\underline{\mathbf{x}}} \Phi(\underline{\mathbf{x}}, t_s)\, e^{-i\underline{\mathbf{k}}\cdot\underline{\mathbf{x}}}. \tag{28}$$

The initial points in x-space are the $\underline{\mathbf{x}}$, while the k-space lattice has fine spacing $\Delta\underline{k}_j = \Delta k_j/\lambda_j$ and values $\underline{k}_j = \widetilde{\underline{l}}_j \Delta\underline{k}_j$ indexed by $\widetilde{\underline{m}}_j = 0, \ldots, (\lambda_j M_j - 1)$ with $\widetilde{\underline{l}}_j = \mathrm{mod}\left[\widetilde{\underline{m}}_j + \frac{1}{2}\underline{M}_j,\ \underline{M}_j\right] - \frac{1}{2}\underline{M}_j$. One in every λ_j values of $\underline{k}_j$ will fall on a k_j value that is also present in the small lattice of the starting state. In particular, instead of using the large index $\underline{\widetilde{\mathbf{m}}}$, its values can be alternatively enumerated by a pair of integers in the following way:

$$\underline{\widetilde{m}}_j = \lambda_j \widetilde{m}_j + q_j, \tag{29}$$

where the coarse index $\widetilde{m}_j = 0, \ldots, (M_j - 1)$ runs over the same set of momenta as in the starting state on the small lattice M, while a fine structure index $q_j = 0, \ldots, (\lambda_j - 1)$ counts the small $\Delta\underline{k}_j$ steps within the larger momentum step Δk_j. The k values themselves are

$$\underline{k}_j = \Delta k_j \left(\widetilde{l}_j + \frac{q_j}{\lambda_j}\right) \tag{30}$$

when the small lattice size M_j is even (as is usual). The index $\widetilde{l}_j$ is as defined after Eq. (7). Odd M_j introduces a minor but distracting complication, so from here until (38) we will assume even M_j and return to this matter at the end of the section. It is convenient to define a vector of fractional momentum steps

$$\alpha_j = \frac{q_j \Delta k_j}{\lambda_j} \in [0, \Delta k_j) \tag{31}$$

so that the fine-grained momenta can be written in a concise vector notation:

$$\underline{k}_j = k_j + \alpha_j(q_j); \qquad \underline{\mathbf{k}} = \mathbf{k}_{\widetilde{\mathbf{m}}} + \boldsymbol{\alpha}_{\mathbf{q}} \tag{32}$$

in terms of the coarse starting momenta $\mathbf{k}$ and the fine grained shift $\boldsymbol{\alpha}$.

Now luckily, the majority of the elements in the sum over points $\underline{\mathbf{x}}$ in (28) can be discarded because they are in vacuum and contribute zero. Provided we make the commonsense assumption that the large lattice includes the entire lattice $\mathbf{x}$ for the small starting cloud, then this leaves just the sum over the usual starting state indices $\mathbf{n}$ defined in (1). With this, and substituting (1), (26) and (30) into (28), one obtains:

$$\underline{\widetilde{\Psi}}(\underline{\mathbf{k}}, t_s) = \frac{\Delta V}{(2\pi)^{d/2}} \sum_{\mathbf{n}} \Psi_{\mathbf{n}}(t_s) \times \exp\left[-i \sum_{j=1}^{d} \frac{2\pi n_j}{M_j}\left(\widetilde{m}_j + \frac{q_j}{\lambda_j}\right)\right] e^{-i\underline{\mathbf{k}}\cdot\mathbf{a}} \tag{33}$$

which is just a sum over the small lattice. In fact, any element of the k-space wavefunction is given by an appropriate DFT on the small lattice:

$$\underline{\widetilde{\Psi}}(\underline{\mathbf{k}}, t_s) = \frac{\Delta V}{(2\pi)^{d/2}} e^{-i\underline{\mathbf{k}}\cdot\mathbf{a}}\, \mathrm{DFT}\left[\Psi_{\mathbf{n}}(t_s)\, e^{-i\boldsymbol{\alpha}_{\mathbf{q}}\cdot(\mathbf{x}-\mathbf{a})}\right]_{\widetilde{\mathbf{m}}} \tag{34}$$

with the help of the coarse $\widetilde{\mathbf{m}}$ and fine $\mathbf{q}$ indices. To get the entire wavefunction, a separate FFT is required for each differing value of $\mathbf{q}$.

The time evolution between the DFTs is just

$$\underline{\widetilde{\Psi}}(\underline{\mathbf{k}}, t_{\mathrm{final}}) = \underline{\widetilde{\Psi}}_{\widetilde{\mathbf{m}},\mathbf{q}}(t_{\mathrm{final}}) = \underline{\widetilde{\Psi}}(\underline{\mathbf{k}}, t_s)\ \exp\left[-i\frac{\hbar t_{\mathrm{flight}} |\underline{\mathbf{k}}|^2}{2m}\right]. \tag{35}$$

To obtain the final state in x-space, one discretizes (10) and obtains the following expression on the fine lattice:

$$\underline{\Psi}(\underline{\mathbf{x}}, t_{\mathrm{final}}) = \frac{(2\pi)^{d/2}}{\lambda V} \sum_{\widetilde{\mathbf{m}},\mathbf{q}} e^{i\underline{\mathbf{k}}\cdot\underline{\mathbf{x}}}\, \underline{\widetilde{\Psi}}_{\widetilde{\mathbf{m}},\mathbf{q}}(t_{\mathrm{final}}). \tag{36}$$

We do not need the entire huge lattice $\underline{\mathbf{x}}$, only the coarsened version with selected sparse points given by (27). Taking only the subset $\bar{\mathbf{x}}$ of points from the $\underline{\mathbf{x}}$ lattice and substituting according to (27) and (30) gives

$$\overline{\Psi}(\bar{\mathbf{x}}, t_{\mathrm{final}}) = \frac{(2\pi)^{d/2}}{\lambda V} \sum_{\widetilde{\mathbf{m}},\mathbf{q}} \underline{\widetilde{\Psi}}_{\widetilde{\mathbf{m}},\mathbf{q}}(t_{\mathrm{final}}) \times \exp\left[i \sum_{j=1}^{d} \frac{2\pi \widetilde{m}_j \bar{n}_j \lambda_j}{M_j} + i\boldsymbol{\alpha}_{\mathbf{q}}\cdot\bar{\mathbf{x}} + i\mathbf{k}_{\widetilde{\mathbf{m}}}\cdot\bar{\mathbf{a}}\right]. \tag{37}$$

This can almost be written as a sum of DFTs, except for one detail that was seen already in Section 4.3: For a DFT over the small k-space lattice $\widetilde{\mathbf{m}}$ of size M, normally the "x-space" indices should run over the range $0, \ldots (M_j - 1)$. Here instead, in the relevant part of the exponent $i\sum_{j=1}^{d} \frac{2\pi \widetilde{m}_j \bar{n}_j \lambda_j}{M_j}$, we have the quantity $n'_j = (\bar{n}_j \lambda_j)$ which increments in jumps: $n'_j = 0, \lambda_j, 2\lambda_j, \ldots, [\lambda_j(M_j - 1)]$. Fortunately, the whole exponential is unchanged upon adding M_j to n'_j. Hence we can define the auxiliary numbers n''_j like in (23), which will index the output of the DFT. Then, the final result for the coarse-grained wavefunction after flight can be written as a sum of inverse DFTs on the small lattice M:

$$\overline{\Psi}_{\bar{\mathbf{n}}}(t_{\mathrm{final}}) = \frac{(2\pi)^{d/2}}{\lambda \Delta V} \sum_{\mathbf{q}} e^{i\boldsymbol{\alpha}_{\mathbf{q}}\cdot\bar{\mathbf{x}}}\, \mathrm{DFT}^{-1}\left[\widetilde{\Psi}_{\widetilde{\mathbf{m}},\mathbf{q}}(t_{\mathrm{final}}) e^{i\mathbf{k}_{\widetilde{\mathbf{m}}}\cdot\bar{\mathbf{a}}}\right]_{\mathbf{n}''}. \tag{38}$$

Note how the proper expression (38) differs from the naive (24) by having a sum of similar DFTs indexed by $\mathbf{q}$, that account for the fine-scale fractional k-space shifts $\boldsymbol{\alpha}_{\mathbf{q}}$. To calculate these, only FFTs *on the small lattice* M are required. There is never a need to store the huge $\underline{M}$ lattice.

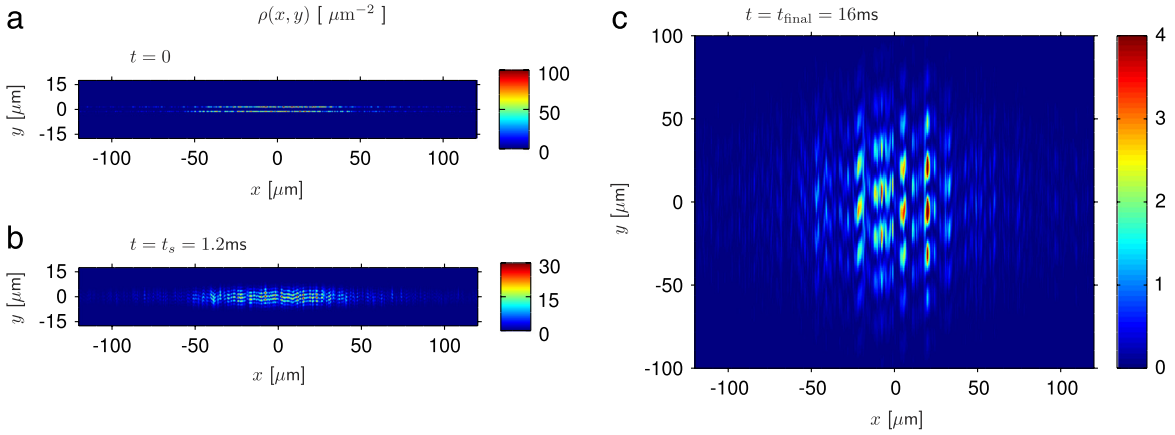

Fig. 4. Density patterns during the free flight described in Section 6. Shown is the x, y density $\rho(x, y) = \int dz |\Psi(\mathbf{r})|^2$. Panel (a): initial state released from the traps at $t = 0$. Panel (b): The "starting" state for the free expansion, at $t = t_s = 1.2$ ms, after the conversion phase. Panel (c): The state at the detector at $t = t_{\rm final} = 16$ ms.

Finally, returning to the unusual case of odd M_j, (30) and (32) also apply to all points except for a few with $\widetilde{m}_j = (M_j - 1)/2$ that end up with $\underline{k}_j > k_j^{\max} = \pi M_j/L_j$. For these, one should substitute $\widetilde{l}_j \to \widetilde{l}'_j = (\widetilde{l}_j - M_j)$ in (30) and $k_j \to (k_j - M_j \Delta k_j)$ in (32) to carry out the umklapp flipping to negative wavevectors on the fine momentum grid. It turns out that the only change in the intervening expressions above is that one should replace $k_j \to (k_j - M_j \Delta k_j)$ in the $\mathbf{k}_{\widetilde{\mathbf{m}}}$ of (37) and (38) for the several points when $\widetilde{m}_j = (M_j - 1)/2$ and $q_j \geq \lambda_j/2$. This actually makes very little difference in practice provided that all $\lambda_j \ll M_j$.

5.2. The prescription

This is a summary of the main result that gathers the above results together. One starts from the initial wavefunction $\Psi(\overline{\mathbf{x}}, t_s) \equiv \Psi_{\mathbf{n}}(t_s)$ described in d dimensions on an x-space lattice with M_j points in each dimension $j = 1, \ldots, d$ of length L_j. The region outside this lattice is initially vacuum. The positions of the points are

$$x_j = a_j + \frac{n_j L_j}{M_j} \tag{39}$$

with indices $n_j = 0, \ldots, (M_j - 1)$. The momentum spacing is $\Delta k_j = 2\pi/L_j$. Free flight occurs for a time interval $t_{\rm flight} = t_{\rm final} - t_s$. Subsequently the lattice spacing on which the wavefunction is described in x-space is magnified by integer factors λ_j, giving lattice points

$$\overline{x}_j = \overline{a}_j + \frac{\overline{n}_j \lambda_j L_j}{M_j} \tag{40}$$

with indices $\overline{n}_j = 0, \ldots, (M_j - 1)$. The volume of the initial lattice is $V = \prod_j L_j$, the number of points $M = \prod_j M_j$, the volume magnification $\lambda = \prod_j \lambda_j$. The fractional momentum steps

$$\alpha_j = \frac{\Delta k_j q_j}{\lambda_j} \in [0, 1)\Delta k_j \tag{41}$$

form a vector $\boldsymbol{\alpha}_{\mathbf{q}}$ that is enumerated by the indices $q_j = 0, 1, \ldots, (\lambda_j - 1)$. Bold quantities are vectors in d-dimensional space. The final x-space wavefunction on the magnified lattice is given by:

$$\overline{\Psi}_{\overline{\mathbf{n}}}(t_{\rm final}) = \sum_{\mathbf{q}} f^{(\mathbf{q})}_{\overline{\mathbf{n}}} B^{(\mathbf{q})}_{\mathbf{n}''} \tag{42a}$$

$$f^{(\mathbf{q})}_{\overline{\mathbf{n}}} = \frac{1}{\lambda} \exp\left[i\boldsymbol{\alpha}_{\mathbf{q}} \cdot \left(\overline{\mathbf{x}}_{\overline{\mathbf{n}}} - \mathbf{a} - \boldsymbol{\alpha}_{\mathbf{q}} \frac{\hbar t_{\rm flight}}{2m} \right) \right] \tag{42b}$$

$$B^{(\mathbf{q})}_{\mathbf{p}} = {\rm FFT}^{-1} \left[\widetilde{A}^{(\mathbf{q})}_{\widetilde{\mathbf{m}}} e^{i\mathbf{k}_{\widetilde{\mathbf{m}}} \cdot \left[\overline{\mathbf{a}} - \mathbf{a} - (\mathbf{k}_{\widetilde{\mathbf{m}}} + 2\boldsymbol{\alpha}_{\mathbf{q}}) \frac{\hbar t_{\rm flight}}{2m} \right]} \right]_{\mathbf{p}} \tag{42c}$$

$$\widetilde{A}^{(\mathbf{q})}_{\widetilde{\mathbf{m}}} = {\rm FFT} \left[\Psi_{\mathbf{n}}(t_s) e^{-i\boldsymbol{\alpha}_{\mathbf{q}} \cdot [\mathbf{x} - \mathbf{a}]} \right]_{\widetilde{\mathbf{m}}} \tag{42d}$$

when all M_j are even, with the elements of the auxiliary index in (42a) being

$$n''_j = {\rm mod}\left[\overline{n}_j \lambda_j, M_j \right], \tag{42e}$$

and FFT indicating fast Fourier transforms. The wavevectors $\mathbf{k}_{\widetilde{\mathbf{m}}}$ are given by (7).

If any M_j are odd, then in the $\mathbf{k}_{\widetilde{\mathbf{m}}}$ of (42c) one should further umklapp the maximum k value in each of those dimensions j when $q_j \geq \lambda_j/2$. That is, for $\widetilde{m}_j = (M_j - 1)/2$ and $q_j \geq \lambda_j/2$, replace $k_j(\widetilde{m}_j)$ by $-\Delta k_j (M_j + 1)/2$.

6. 3D Example

As an example of a calculation that cannot be done by brute force on a PC, consider a full simulation of the relative phase measurements in the Vienna experiment [1]. Here, two neighboring ultracold atomic clouds of ^{87}Rb, $\Psi^{(\pm)}$, are initially populated in quasi-one-dimensional harmonic traps that are elongated in the x direction, as seen in Fig. 4(a). Trap frequencies are $\omega = 2\pi \times 6.4$ Hz in the x direction, and $\omega_\perp = 2\pi \times 1400$ Hz in the transverse directions. The clouds are initially separated by a small gap of $D = 2.75$ nm $= 0.645 a_{\rm ho}$ in the y direction. The traps are released at $t = 0$, and rapid expansion takes place in the y and z directions. The clouds soon interfere, forming a fringe pattern, as shown in Fig. 4(b–c). The local displacement $\delta y(x)$ of the fringes in the y direction from the middle position $y = 0$ corresponds to the phase difference that existed locally between the two initial clouds at $t = 0$.

$$\delta y(x) \propto \Delta\theta_0(x) = \angle \Psi_0^{(+)}(x) - \angle \Psi_0^{(-)}(x). \tag{43}$$

The fringe pattern is detected after 16 ms of free flight ($t_{\rm final} = 0.65/\omega$) at a detector. This is basically an integral of the final density over the z direction. In this way, a local phase difference measurement on the initial clouds can be made.

There is of course some distortion during flight, as was seen in Section 3, and a simulation of the expansion can be necessary to see quantitatively how the pattern at the detector corresponds to the initial phase profile. Here the case where the two clouds are populated by independent thermal gases will be considered. The one-dimensional c-field wavefunctions $\Psi_{\rm ic}^{(\pm)}(x)$ are generated using the same SGPE method as in Section 3 and the same parameters. Eq. (15) is run separately with independent noises for each of the two clouds. As these are quasi-1D traps, the wavefunction in the y and z directions is well approximated by just the harmonic oscillator ground state. The initial state (see Fig. 4(a)) consists of two terms:

$$\Psi_0(\mathbf{x}) = v\sqrt{\frac{m\omega_\perp}{\pi\hbar}}\sum_{\pm}\Psi_{\rm ic}^{(\pm)}(x) \times \exp\left[-\left\{\left(y\pm\frac{D}{2}\right)^2+z^2\right\}\frac{m\omega_\perp}{2\hbar}\right]. \quad (44)$$

The conversion phase is simulated with the 3D Gross–Pitaevskii equation (GPE)

$$\hbar\frac{d\Psi(\mathbf{x},t)}{\partial t} = -i\left[-\frac{\hbar^2}{2m}\nabla^2 + g\,|\Psi(\mathbf{x},t)|^2\right]\Psi(\mathbf{x},t) \quad (45)$$

using with a semi-implicit split-step algorithm [31]. It is run until 1.2 ms, which is $t_s = 0.05/\omega$. The s-wave scattering length is 5.24 nm, giving a value of $g = 0.01544\hbar\omega a_{\rm ho}^3$ for the interaction strength in 3D in terms of $a_{\rm ho} = \sqrt{\hbar/m\omega}$. The lattice used has dimensions $L_x = 67.5a_{\rm ho}$, $L_{y,z} = 8.26a_{\rm ho}$, and $M = 2304\times256\times256$ points. The calculation took 1 h 20 mins on an Intel 2.4 GHz CPU using the FFTW library [18], and used 8% of the 96 GB RAM memory on the PC. The situation at t_s is shown in Fig. 4(b).

Subsequently, the $\Psi(\mathbf{x},t_s)$ were fed into the free flight prescription (42) developed here, for expansion out to the detector at 16ms. The lattice expansion factors were $\lambda_y = \lambda_z = 8$ and $\lambda_x = 1$, i.e. no expansion in the x direction. However, the initial $\bar{x}$ lattice was slightly buffered with vacuum at the edges with respect to the one used for the conversion phase, having a length $\bar{L}_x = 94.92a_{\rm ho}$. This was to allow some natural spreading, which was too small to make a lattice expansion by $\lambda_x = 2$ worthwhile. The final coarse lattice had $\underline{M} = 3240\times256\times256$ points. The calculation took 64 min on the same PC and used 14% of RAM. The resulting predicted detector image is shown in Fig. 4(c).

For comparison, a brute force calculation using the plain (12) was not able to reach the detection time. The best that was obtained on the aforementioned PC without going into swap space was expansion out to $t = 0.40/\omega$, corresponding to 10 ms, or 62% of the flight. This took 3 h on a $\underline{M} = 2880\times1350\times1350$ lattice, and used 88% of RAM, as well as requiring special additional work with the code to pass 64 bit pointers into the FFTW library. 64 bit pointers were required when $\underline{M} \geq 2^{31}$.

Fig. 5 shows the apparent phase differences that can be inferred from the free-expansion at different times, and the true initial phase difference. The fringe shift $\delta y(x)$ was estimated from the y position of the maximum density peak for a given x. The proportionality constant in (43) can be estimated by considering the free flight evolution of (44) in the y direction only, ignoring other effects. One finds

$$\psi(y,t) = \frac{1}{\sqrt{1+i\omega_\perp t}}\sum_{\pm}A_\pm\exp\left[-\frac{m\omega_\perp(2y\mp D)^2}{8\hbar(1+i\omega_\perp t)}\right] \quad (46)$$

where $A_\pm = \left(\frac{m\omega_\perp}{\pi\hbar}\right)^{1/4}\Psi_0^{(\pm)}(x)$. We are most interested in the limit $\omega_\perp t \gg 1$ and $y \gg D$. We expand the exponents in the density $|\psi(y)|^2$ to lowest nontrivial orders in the small quantities

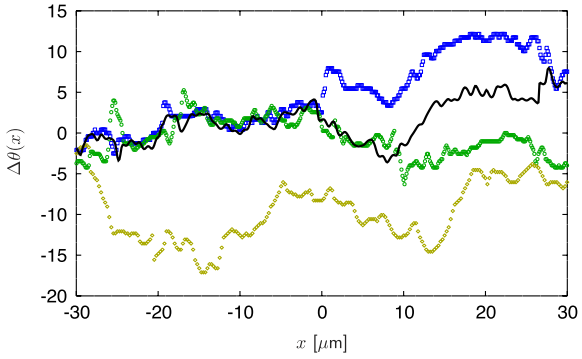

Fig. 5. Phase differences along x in the central region of the cloud, after unwrapping modulo 2π to remove sudden jumps greater then π. The black line shows the actual phase difference $\angle\Phi_+(x) - \angle\Phi_-(x)$ between the initial clouds. The blue squares and green circles show the apparent phase differences inferred from the fringe shifts $\delta y(x)$ in the density images of Fig. 4(b, c), respectively, using (48). The yellow diamonds show what would be inferred from the very late density image of Fig. 6 at 62 ms. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

$\eta_t = 1/\omega_\perp t$ and $\eta_y = D/y$: that is, $\mathcal{O}(\eta_t^2, \eta_y^{-2})$ for the amplitude and $\mathcal{O}(\eta_t, \eta_y^{-1})$ for the phases. This gives

$$|\psi(y,t)|^2 \approx \frac{\exp\left[-\frac{2my^2}{\hbar\omega_\perp t^2}\right]}{\sqrt{1+(\omega_\perp t)^2}} \times\left\{|A_+|^2+|A_-|^2+2|A_+||A_-|\cos\left[\frac{myD}{\hbar t}-\Delta\theta_0(x)\right]\right\}. \quad (47)$$

Hence, the phase difference estimate (modulo 2π) is

$$\Delta\theta_0(x) \approx y_{\rm peak}(x)\,\frac{mD}{\hbar t} \quad (48)$$

where $y_{\rm peak}$ is the location of the peak nearest to $y = 0$. What we see in Fig. 5 is that the global long-wavelength behavior of the phase difference is generally predicted well by the fringes in the expanded cloud. This is apart from some remnant localized shifts of modulo 2π that move an entire segment by 2π without affecting the long-wavelength phase trend. These shifts are at 0 μm and 28 μm for the early cloud at t_s, and near 10 μm for the cloud at the detector. However, one can also see that true to the behavior seen in the 1D case of Section 3, the prediction of local details in the phase difference is largely scrambled during the time of flight.

The effort required by (42) for even very long expansions scales relatively graciously. For example, continuing the expansion out to $t = 62$ ms $= 2.5/\omega$ is also possible. This is shown in Fig. 6 and the yellow plot in 5, and one sees continuing change in the fringe profile. In particular, the phase difference estimate is starting to become bad, with large long-scale discrepancies. This is because now the expansion has lasted long enough that a lot of movement of the defects in the x direction has occurred. This is expected, since the estimate for formation time of a momentum distribution from Section 2.5 is $t_v = \sqrt{8}/\omega = 70$ ms here. The final cloud is now a relatively huge $0.3\times0.6\times0.6$ mm in size, having expanded by a factor of about 10^4 in the y and z directions since its release from the trap. The final lattice is scaled by $\lambda_x = 2$, and $\lambda_y = \lambda_z = 30$ from that at t_s. This calculation took 18 h on the reference PC, and used 7 GB of memory. The direct approach with the $\underline{M}$ lattice would have needed 4000 GB.

7. Discussion

7.1. Efficiency

To take advantage of the memory savings in (42), one should evaluate each term in the sum (42a) labeled by $\mathbf{q}$ sequentially,

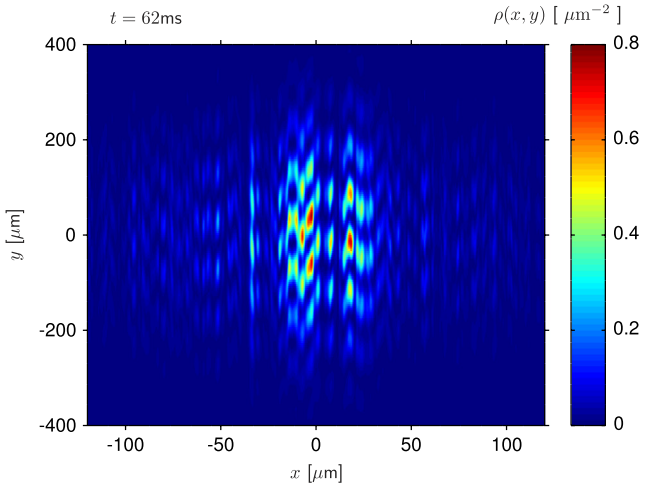

Fig. 6. The density pattern that would be seen after $t = 62$ ms of free flight, details as in Fig. 4, except for a 1:4 aspect ratio between the axes.

and accumulate its contribution to $\overline{\Psi}_{\bar{\mathbf{n}}}(t_{\text{final}})$. This way, memory requirements will be $\approx 2M$ complex numbers [$32 \times M$ bytes for the usual double precision]—one array of size M for carrying out the FFTs, and one to store the accumulated sum $\overline{\Psi}_{\bar{\mathbf{n}}}(t_{\text{final}})$. Some time efficiency can be gained by using a third array of size M to store the starting state $\Psi_{\mathbf{n}}(t_s)$, but using $48 \times M$ bytes in total.

The computational load in terms of operations scales as

$$\lambda M \log M \tag{49}$$

which is slightly faster than the brute force approach that would use (12) directly on a vacuum padded lattice $\underline{M}$ and take $\sim\lambda M \log(\lambda M)$ operations. The speed-up is mostly marginal—an improvement by a factor of $(1 + \log\lambda/\log M)$. The somewhat surprising result that there is any speed up at all compared to the highly-optimized FFT on $\underline{M}$ is due to the fact that so much of the initial system is vacuum and does not contribute. For this reason, there is no advantage to be gained by trying to use the maximum starting lattice size M that will fit in memory (perhaps after padding the nonzero part of the field that comes out of the conversion phase with vacuum), and minimum magnification λ.

The memory needed is of course strongly reduced—by a factor of at least $\lambda/2$, comparing to the most memory-efficient in-place FFT on the huge $\underline{M}$ lattice.

An expansion in d directions requires summing of a number of terms that grows as $(t_{\text{flight}})^d$. This can eventually become fairly time intensive as was seen for the calculation of Fig. 6. Memory use never budges above the baseline no matter how long the flight takes.

The time needed can be alleviated by an extremely basic parallelization. Namely, distributing the evaluation of the $B^{(\mathbf{q})}$ on many processing cores. Up to λ cores could be used to obtain the result in a time $\sim M \log M$. However, a significantly smaller number will be optimal since λ FFTs in parallel will require $\sim\lambda M$ numbers stored in memory again, which is what one is trying to avoid.

7.2. Relationship to fast Fourier transforms

The algorithm presented in (42) bears some rough resemblance to a Cooley–Tukey FFT algorithm [32] with radix-λ_j. The similarity is that the end results of the FFTs on the smaller M lattice are multiplied by twiddle factors in (42b). These involve $e^{i\alpha_j \bar{x}_j}$ that introduces fractional phase shifts compared to those available on the FFT lattice. However, the overall procedure is quite different to Cooley–Tukey and relies heavily on the vacuum padding assumption (26). This allows it to e.g. perform the two sequential FFTs on each $\mathbf{q}$th term in the final sum (42a). This is also what allows the effort to scale as λ instead of the λ^2 that would be expected from a manual summation of smaller M-size FFTs.

Looking at (42e), one can see that when λ and M have common factors, this n'' index only accesses a part of the field $B^{(q)}_{\mathbf{p}}$. One can try to gain some computational advantage from this by using a pruned FFT [33] for the step (42c), to calculate only the required $\mathbf{p}$ values. The advantage of pruned FFTs is not huge though. This would reduce the overall computational effort at most from $2M \log M$ to $2M \log M - M \log \lambda$.

7.3. Loss of phase information

An important feature of the algorithm (42) to be aware of is that while the density in x-space at t_{final} is calculated precisely, the $\overline{\Psi}_{\bar{\mathbf{n}}}(t_{\text{final}})$ is not generally viable for further evolution, and does not store the correct momentum distribution. This is because of the phase aliasing (25) discussed in Section 4.3.

The wavefunction that can be reconstructed from $\overline{\Psi}_{\bar{\mathbf{n}}}(t_{\text{final}})$ is:

$$\widetilde{\overline{\Psi}}_{\widetilde{\bar{\mathbf{m}}}} = \frac{\lambda \Delta V}{(2\pi)^{d/2}} e^{-i\bar{\mathbf{a}}\cdot\bar{\mathbf{k}}_{\widetilde{\bar{\mathbf{m}}}}} \,\text{DFT}\left[\overline{\Psi}_{\mathbf{n}}\right]_{\widetilde{\bar{\mathbf{m}}}}. \tag{50}$$

This sits on a fine k-space lattice $\bar{k}_j(\widetilde{\bar{m}}_j) = \widetilde{\bar{l}}_j\,(2\pi/\bar{L}_j)$ with $\widetilde{\bar{l}}_j = \text{mod}\left[\widetilde{\bar{m}}_j + \frac{1}{2}M_j\,,\, M_j\right] - \frac{1}{2}M_j$. The resulting momentum distribution is shown in Fig. 7, for the same 1D system that was studied in Section 3. This time, the initial wavefunction $\Psi_{\text{ic}}(x)$ was evolved to times t_{flight} using the prescription (42) on the initial $M = 2048$ lattice rather than the standard step-by-step evolution (16) on $\underline{M} = 81\,920$ that was used in Section 3. The green case at $t_{\text{flight}} = 10/\omega$ might still be passable for some purposes, though the high momenta are already lost. The red longer-time case is completely scrambled.

The fact that the phase structure in x-space remains small-scale despite a magnification of the density stymies several superficially promising ideas on how to increase the efficiency of the expansion calculation:

First, one could be tempted to try to reduce the processing load to only $\sim\log\lambda$ FFTs on M-points instead of the present λ FFTs, by implementing several sequential expansions (42) by small factors, say $\lambda_j = 2$. However, at each such step we are left with a discretized wavefunction that has its momentum-space tails truncated. This will soon come to resemble the red line in Fig. 7 and become useless for further evolution.

Overall, the algorithm presented here relies essentially on the possibility to discard the high frequency phase behavior in the final expanded state, because it is not measured. This is in some contrast to the long tradition of adaptive or nonuniform meshing techniques in fields such as ultrashort optical pulses in laser induced fusion and nonlinear optics ([34,35] and references therein), going back also earlier work on adaptive rezoning in aerodynamics and flows with shocks (see e.g. [36–38] and references therein). These approaches saved computational resources by adapting the numerical lattice in a time-dependent way to the state of the system at a given moment. They were able to deal with a changing system size in situations where the potential size of an accurate computational lattice remains broadly constant, but the physical dimensions vary. A time-dependent snapshot of the state that is obtained can be used for later evolution — a matter that is particularly important for nonlinear evolution. For wave dynamics, such a situation is more typical of focusing phenomena, shocks, caustics, etc. than expansion.

What we have here, instead, is a case where the size of a computational lattice that can hope to encode the full state is growing rapidly. We are saved, though, by the option to ignore most of the information in the expanded state because: (i) It will

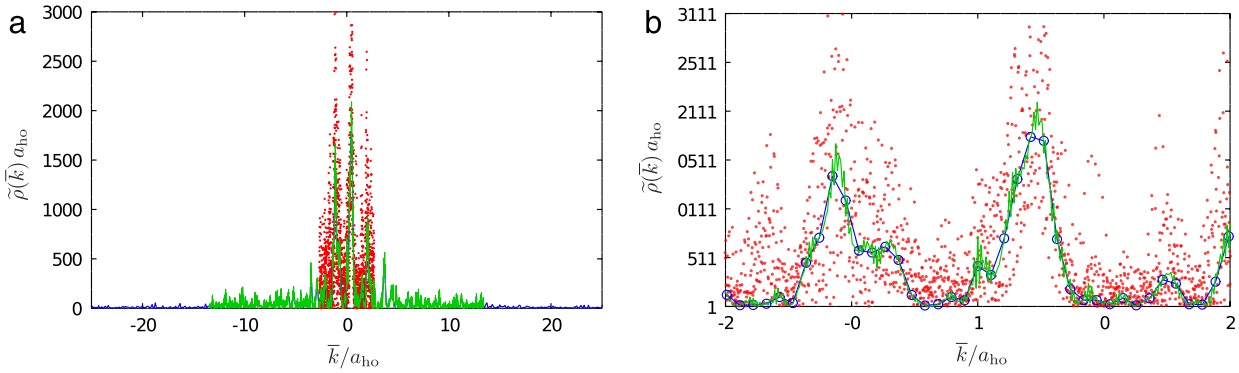

Fig. 7. Loss of momentum information in the final wavefunction generated by (42). The plot shows: In blue: the true k-space density $\widetilde{\rho}(k) = |\widetilde{\Psi}(k)|^2$ in the cloud on the initial (M) lattice. Physically, this is preserved during free evolution; In green and red: The apparent k-space density $\widetilde{\rho}(\overline{k}) = |\widetilde{\overline{\Psi}}(\overline{k})|^2$, inferred by a DFT (50) of the coarse grained final wavefunction $\overline{\Psi}(\overline{x}, t_{\rm final})$ in x-space. Green is for a magnification of $\lambda = 8$ at $t_{\rm flight} = 10/\omega$, while red is for a longer time of $t_{\rm flight} = 40/\omega$ with $\lambda = 40$. Other parameters as in Section 3. Panel (b) is a magnification of part of Panel (a). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

not be measured or needed for subsequent evolution, and: (ii) The intervening evolution is linear so that the influence of each data point at the start can be propagated forward to the end without the need for interactions at intermediate times. That, in turn, is more typical of expansion in field theories, such as quantum wavefunctions, or electromagnetic fields, and this is where the method presented here would rather find utility.

7.4. Generalizations

The algorithm is readily adapted to cases where several complex-valued fields are present. One such case that may be aided with the algorithm presented here are positive-P simulations of supersonic BEC collisions [39,40,6,41]. Here, two independent complex-valued fields $\psi(\mathbf{x})$ and $\psi^+(\mathbf{x})$ that correspond to the $\widehat{\Psi}(\mathbf{x})$ and $\widehat{\Psi}^\dagger(\mathbf{x})$ Bose fields are used, and allow for the exact treatment of quantum fluctuations. The comparison of calculated and experimental pair velocity correlation widths has been problematic in these systems, because of the narrowness of the correlation peak in velocity [6]. The detected peak is distorted in comparison with its k-space prediction due to not yet being in the far-field regime. There is no hope of a direct calculation of the free flight because the quantity $\varepsilon t_{\rm flight}/\hbar$ of (18) is very high (up to $\sim 10^4$) in BEC collision experiments.

The approach can also be trivially adapted to cases of other spectra than the free particle one. This simply requires a modification of (9) to $\widetilde{\Psi}(\mathbf{k}, t_{\rm final}) = \widetilde{\Psi}(\mathbf{k}, t_s) \exp\left[-it_{\rm flight}\omega_{\mathbf{k}}\right]$, with appropriate tweaks in (42b) and (42c). The crucial element is the presence of the vacuum assumption (26).

7.5. Conclusions

To conclude, an algorithm (42) has been presented that allows the exact calculation of the density of a wavefunction freely expanding into vacuum for practically arbitrary flight times without filling up the computer memory. The memory requirements do not depend on flight time and are the same size as the initial input state. Computation time is slightly faster than using an FFT on a large vacuum padded lattice. It is implemented using standard FFT libraries and some summing of terms. The approach relies crucially on two physical inputs: (1) That the initially compact wavefunction expands into vacuum, and (2) that the density length scale of the expanded cloud grows approximately linearly with time. The approach makes no assumptions about symmetries of the system or about the input wavefunction, so that it is a black box tool that can be immediately applied to general cases. This makes it well suited to the study of wavefunctions containing defects or samples of a thermal ensemble, a topic of many recent experiments [1,20–25]. The flight times over which nontrivial defect evolution occurs during free flight are estimated in Section 2.4.

Acknowledgments

I would like to thank Ray Ng and Mariusz Gajda for helpful discussions on this matter, and Julius de Hond whose feedback identified an omission in the manuscript. The work was supported by the National Science Centre (Poland) Grant No. 2012/07/E/ST2/01389.

References

[1] M. Gring, M. Kuhnert, T. Langen, T. Kitagawa, B. Rauer, M. Schreitl, I. Mazets, D.A. Smith, E. Demler, J. Schmiedmayer, Science 337 (6100) (2012) 1318–1322.
[2] D.A. Smith, M. Gring, T. Langen, M. Kuhnert, B. Rauer, R. Geiger, T. Kitagawa, I. Mazets, E. Demler, J. Schmiedmayer, New J. Phys. 15 (7) (2013) 075011.
[3] A. Perrin, R. Bucker, S. Manz, T. Betz, C. Koller, T. Plisson, T. Schumm, J. Schmiedmayer, Nat. Phys. 8 (2012) 195.
[4] K. Gawryluk, M. Gajda, M. Brewczyk, Phys. Rev. A 92 (2015) 043607.
[5] J.-C. Jaskula, M. Bonneau, G.B. Partridge, V. Krachmalnicoff, P. Deuar, K.V. Kheruntsyan, A. Aspect, D. Boiron, C.I. Westbrook, Phys. Rev. Lett. 105 (2010) 190402.
[6] K.V. Kheruntsyan, J.-C. Jaskula, P. Deuar, M. Bonneau, G.B. Partridge, J. Ruaudel, R. Lopes, D. Boiron, C.I. Westbrook, Phys. Rev. Lett. 108 (2012) 260401.
[7] Y. Kagan, B.V. Svistunov, Phys. Rev. Lett. 79 (1997) 3331–3334.
[8] M.J. Davis, S.A. Morgan, K. Burnett, Phys. Rev. Lett. 87 (2001) 160402.
[9] K. Goral, M. Gajda, K. Rzazewski, Opt. Express 8 (2) (2001) 92–98.
[10] M. Brewczyk, M. Gajda, K. Rza zewski, J. Phys. B: At. Mol. Opt. Phys. 40 (2) (2007) R1.
[11] N.P. Proukakis, B. Jackson, J. Phys. B: At. Mol. Opt. Phys. 41 (20) (2008) 203002.
[12] P. Blakie, A. Bradley, M. Davis, R. Ballagh, C. Gardiner, Adv. Phys. 57 (5) (2008) 363–455.
[13] N. Proukakis, S. Gardiner, M. Davis, M. Szymanska (Eds.), Quantum Gases: Finite Temperature and Non-Equilibrium Dynamics, World Scientific, 2013.
[14] M.J. Steel, M.K. Olsen, L.I. Plimak, P.D. Drummond, S.M. Tan, M.J. Collett, D.F. Walls, R. Graham, Phys. Rev. A 58 (1998) 4824–4835.
[15] A. Sinatra, C. Lobo, Y. Castin, J. Phys. B: At. Mol. Opt. Phys. 35 (17) (2002) 3599.
[16] A. Polkovnikov, Ann. Physics 325 (8) (2010) 1790–1852.
[17] A.D. Martin, J. Ruostekoski, New J. Phys. 12 (5) (2010) 055018.
[18] M. Frigo, S.G. Johnson, Proc. IEEE 93 (2) (2005) 216–231. Special issue on "Program Generation, Optimization, and Platform Adaptation.
[19] M. Frigo, S.G. Johnson, Proc. 1998 IEEE Intl. Conf. Acoustics Speech and Signal Processing, vol. 3, IEEE, 1998, pp. 1381–1384.
[20] L. Chomaz, L. Corman, T. Bienaime, R. Desbuquois, C. Weitenberg, S. Nascimbene, J. Beugnon, J. Dalibard, Nature Commun. 6 (2015) 6162.
[21] S. Serafini, M. Barbiero, M. Debortoli, S. Donadello, F. Larcher, F. Dalfovo, G. Lamporesi, G. Ferrari, Phys. Rev. Lett. 115 (2015) 170402.
[22] G. Lamporesi, S. Donadello, S. Serafini, F. Dalfovo, G. Ferrari, Nat. Phys. 9 (2013) 656.
[23] S. Donadello, S. Serafini, M. Tylutki, L.P. Pitaevskii, F. Dalfovo, G. Lamporesi, G. Ferrari, Phys. Rev. Lett. 113 (2014) 065302.
[24] L.E. Sadler, J.M. Higbie, S.R. Leslie, M. Vengalattore, D.M. Stamper-Kurn, Nature 443 (2006) 312.
[25] C. Weiler, T.W. Neely, D.R. Scherer, A.S. Bradley, M.J. Davis, B.P. Anderson, Nature 455 (2008) 948.

[26] C. Mora, Y. Castin, Phys. Rev. A 67 (2003) 053615.
[27] Y. Castin, J. Phys. IV France 116 (2004) 89.
[28] T. Karpiuk, P. Deuar, P. Bienias, E. Witkowska, K. Pawłowski, M. Gajda, K. Rzążewski, M. Brewczyk, Phys. Rev. Lett. 109 (2012) 205302.
[29] J. Pietraszewicz, P. Deuar, Phys. Rev. A 92 (2015) 063620.
[30] Y. Castin, R. Dum, Phys. Rev. Lett. 77 (1996) 5315–5319.
[31] P. Drummond, I. Mortimer, J. Comput. Phys. 93 (1) (1991) 144–170.
[32] J.W. Cooley, J.W. Tukey, Math. Comp. 19 (1965) 297–301.
[33] H. Sorensen, C. Burrus, IEEE Trans. Signal Process. 41 (3) (1993) 1184–1200. http://dx.doi.org/10.1109/78.205723.
[34] M.C. Mattar, F.P. Newstein, Comput. Phys. Comm. 20 (1980) 139.
[35] V.L. Derbov, L.A. Melnikov, A.D. Novikov, S.K. Potapov, J. Opt. Soc. Amer. B 7 (1990) 1079.
[36] G. Moretti, Comput. & Fluids 7 (1979) 191.
[37] G. Moretti, Annu. Rev. Fluid Mech. 19 (1987) 313.
[38] G. Moretti, Comput. & Fluids 15 (1987) 59.
[39] P. Deuar, P.D. Drummond, Phys. Rev. Lett. 98 (2007) 120402.
[40] P. Deuar, J. Chwedeńczuk, M. Trippenbach, P. Ziń, Phys. Rev. A 83 (2011) 063625.
[41] R.J. Lewis-Swan, K.V. Kheruntsyan, Phys. Rev. A 91 (2015) 052114.