Waveplate design report.

David Paredes Barato[1]
[1]Department of Physics, Durham University, Rochester Building, South Road, Durham DH1 3LE, United Kingdom*

Keywords: Fourier propagation

Dated: February 6, 2012

CONTENTS

I. MOTIVATION

Our experiment consist on an EIT scheme to probe a singly blockaded sphere of ^{87}Rb atoms. This would allow for the creation of a single photon source or the implementation of a CNOT quantum gate.

Previous experiments have shown that long range dipole-dipole interactions between atoms have effects on the propagation of light through a cloud of atoms [1]. That experiment was performed with an atom cloud that was bigger than the blockade radius of the atoms being probed.

Our setup tries to remedy this by using a smaller cloud confined in a tightly focussed optical dipole with a volume smaller than $100\,\mu$m^3. As we load directly from a MOT into the dipole trap and the dipole trap volume is small, we suffer from poor loading (less than 100 atoms/s). The probe time, which was of the order of tens of microseconds, was enough to lose all the atoms from the trap, as we have to shut down the dipole beam while probing.

Therefore, we searched for solutions to the loading problem. Ultimately, we arrived at a solution which has shown to be satisfactory so far: probe repeatedly for shorter periods of time and alternate those periods with periods where the dipole trap is ON to avoid losing the atoms.

However, we thought for a while on other solutions to the poor loading. Given the successful application of Bottle Beam (BoB) traps [2, 3], we decided that making one of these traps might help us. A BoB trap (see Figure 1) creates a region in space with zero electric field completely surrounded by regions of higher intensity. The particles are then trapped in the regions where the electric field is practically zero, and thus the light shifts that they experiment are negligible.

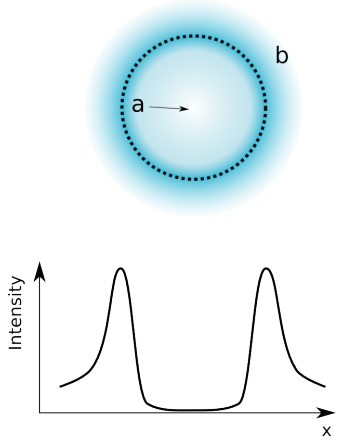

FIG. 1: Schematic view of a Bottle Beam (BoB) trap. It consists on a region of space (a) with no electric field completely surrounded by regions of higher intensity (b). The particles would be trapped in region (a).

It is possible to obtain a BoB trap with a dark spot in the focus of a lens if we apply a π phase difference to half of the power of the a Gaussian beam under adequate conditions. This phase difference makes the beam cancel out in the axis of propagation.

One of the problems that we wanted to target is the difficulty of aligning probe and trap beams simultaneously at such small scales (less than $1\,\mu$m). Therefore, we de-

* david.paredes@durham.ac.uk

cided that the best way to have at least axial alignment is to couple both beams through the same monomode fiber.

How can we create a BoB trap at the trapping frequency while leaving the probe beam unchanged, when both beams have to follow the same optical path? We can design a "magic number" waveplate that applies a π phase shift to part of the trap beam, while the optical path length of the probe beam varies only by integer numbers of the wavelength.

In section II we will briefly describe dipole trapping of neutral atoms and in III the discussion on the wave plate design to perform Bob trapping will be exposed. Section IV will briefly explain the background behind the Fourier propagation method and the implementation in Python. I will present the results in section V, and the discussion will follow in section VI.

II. BOTTLE BEAM TRAP

One of the ways to trap neutral atoms [4] is via dipole traps. Dipole traps rely on the dipolar force created by far detuned electromagnetic fields. The following (brief) discussion on the crucial expressions for the dipole traps is taken from [5]. Please, refer to it for more complete information.

If we assume (using a semi-classical model) that an atom is placed in laser light with electric field $\vec{E}$, it induces an atomic dipole moment $\vec{p}$ that oscilates at the driving frequency ω. The amplitude of both quantities are related by the *complex polarizability*, α, as follows:

$$p = \alpha E. \tag{1}$$

Here, the (complex) polarizability usually depends on the driving frequency, $\alpha = \alpha(\omega)$. The interaction potential induced by the oscillating dipole moment is then given by

$$U_{dip} = -\frac{1}{2}\left\langle \vec{p}\vec{E} \right\rangle = -\frac{1}{2\epsilon_0 c}\Re(\alpha) I, \tag{2}$$

where $\Re(\alpha)$ is the real part of the susceptibility and I is the light power. Therefore, the force on the dipole will depend on the gradient of the light power, $F = -\nabla U_{dip} \sim \nabla I$. Another important quantity, related to the imaginary part of the polarizability, $\Im(\alpha)$, is the power absorbed by the oscillator from this driving field,

$$P_{abs} = \left\langle \dot{\vec{p}}\vec{E} \right\rangle = \frac{\omega}{\epsilon_0 c}\Im(\alpha) I, \tag{3}$$

where we can consider the power absorbed to be re-emitted as dipole radiation. If this is the case, we can obtain the *photon scattering* rate as

$$\Gamma_{sc} = \frac{P_{abs}}{\hbar\omega} \tag{4}$$

The polarizability for a two level system can be calculated, classically or semiclassically, as an oscillating field driving an electron, and the only difference is whether we consider saturation effects or not. The classical formula reads

$$\alpha = 6\pi\epsilon_0 c^3 \frac{\Gamma/\omega_0^2}{\omega_0^2 - \omega^2 - i\,(\omega^3/\omega_0^2)\,\Gamma} \tag{5}$$

In both cases, the polarizability takes into account the damping rate of the electron: in the classical picture, Larmor's formula applies; however, in the semiclassical approach one needs to take into account the dipole matrix element between the ground and excited state. This provides an on-resonance damping rate

$$\Gamma = \frac{\omega_0^3}{3\pi\epsilon_0\hbar c^3}\left|\langle e\,|\mu|\,g\rangle\right|^2. \tag{6}$$

For dipole trapping, where we are interested in the far-detuned regime where we find low saturation and low scattering rates ($\Gamma_{sc} \ll \Gamma$), equation (5) holds.

If we define the detuning as $\Delta \equiv \omega - \omega_0$, the condition $|\Delta| \ll \omega_0$ allows us to apply the *rotating wave approximation*, which neglects fast oscillating terms in the Hamiltonian. Additionally, setting the condition of negligible saturation, we can obtain exressions for both dipole potential and scattering rate:

$$U_{dip}(\vec{r}) = \frac{3\pi c^2}{2\omega_0^3}\frac{\Gamma}{\Delta} I(\vec{r}), \tag{7}$$

$$\Gamma_{sc}(\vec{r}) = \frac{3\pi c^2}{2\hbar\omega_0^3}\left(\frac{\Gamma}{\Delta}\right)^2 I(\vec{r}). \tag{8}$$

Two points can be emphasized on these equations:

- The sign of the detuning determines the repulsive or attractive nature of the dipole force. For red detuned ($\Delta < 0$) dipole traps, the interaction attracts atoms to the maxima of the light field, whereas a blue detuned ($\Delta > 0$) potential repels the atoms out of the maxima.
- The scaling of potential and scattering rate with light intensity and detuning makes convenient working in a situation with large detunings and high optical intensities.

In order to implement a far off-resonance, red detuned trap (FORT), one could use a tightly focused Gaussian which will provide a unique intensity maximum where the atoms will be attracted to.

This technique, however, has some disadvantages. Quantum-mechanically speaking, if we place an atom in a region where there are electromagnetic fields, the energy levels will be *shifted* by the interaction between atom and field. The expression for these *light shifts* or *ac Stark shifts* can be obtained as a second-order perturbation in the field

(they are first order in the light intensity), and in the case of a two-level atom it reduces to

$$\Delta E = \pm \frac{|\langle e|\mu|g\rangle|^2}{\Delta} |E|^2 = \pm \frac{3\pi c^2}{2\omega_0^3} \frac{\Gamma}{\Delta} I \qquad (9)$$

for the ground and excited state (plus and minus sign, respectively).

The light shifts associated with this intensity maxima might be undesirable. It might, then, be possible to trap atoms in a region of space surrounded by intensity maxima using blue detuned light.

In our case, we are interested in creating a trap that provides zero intensity in some volume and then steep "walls" of light. This will constitute a "dark spot" in a beam of light.

Several implementations produce "dark spots" to trap atoms. All of them rely on the cancellation of the power produced at some point along the axis of propagation of a beam. Amongst them, there are many ways to obtain a dark spot using a single beam of blue-detuned light:

- Interferometer [2, 6]: Some combine two Gaussian beams with equal power at the focus with a π phase shift using a Mach-Zehnder interferometer with a telescope in one of the arms, but it cannot be used easily in our case as it requires additional alignment.

- Near-field features [7]: It is possible to obtain field minima in the near field produced by scattering from a circular apperture to trap atoms. However, these traps are usually very close to the aperture which produces them, and they can affect highly excited Rydberg states.

- Holographic plates [8]: They are used to tailor the shape of the beam at the focal spot of a lens. In principle, this is the most flexible method, but we are interested in obtaining a trap that is aligned "by default" in the longitudinal direction by having both probe and trap beams pass through the same fiber. These holographic traps would provide a dark spot in both probe and trap beam, which is not suitable for our purposes.

- Segmented wave plate [9]: They exploit the radial symmetry of a Gaussian beam to obtain one of such traps. They apply a π phase difference to an inner circular region of a Gaussian beam with respect to an outer annulus using a segmented quarter wave plate whose central region is rotated by 90°. The wave plate they use is quite big (several cm), and the losses produced by the cut of the wave plate are negligible. However, these sizes are not practical for our experiment, thus avoiding the cuts would be important.

Furthermore, the constraints in our setup are

- Both trap and probe beams should follow the same optical path. This implies, amongst other things, that both need to travel through the same fiber. In this way, transversal alignment is a default and longitudinal alignment will depend on the dispersive features of the optical elements before the trap.

- The trapping region should be smaller than the dipole blockade radius. In our experiment, this blockade is of the order 5-10 μm.

If we follow the wave plate technique, it is not unfeasible to design one that provides a dark spot in the trapping beam while leaving the probing beam undisturbed. The main problem is to tailor the length such that the probe beam spans an integer number of wavelengths inside the medium, whereas the trapping light spans a half-integer number of them.

In the next section we discuss the design constraints and the solutions we arrived at.

III. WAVE PLATE DESIGN

To create a wave plate that produces a BoB trap at one wavelength and leaves other wavelength almost undisturbed (speaking about phase relationships) it is required that this plate provides different phase shifts in different regions for the trapping light but not for the probe light.

In the case of the segmented waveplate, the outer and inner regions are shifted by $\pm\pi/2$, where the different signs account for the different orientations of the segments in the waveplate. These two, when combined in the trap beam, would provide a π phase shift that cancels the contributions from the outer and inner regions when the beam is focused. However, given the small sizes of our beams to begin with, a cut in a waveplate would produce non-negligible losses which we cannot afford. Therefore, a solution would be to create a waveplate with a stepped profile [see figure] that provides a π phase shift to the trapping light in the region of the feature, be it a bump or a well. The rest of the waveplate provides an homogeneous phase shift across its section, thus not disturbing either of the beams.

Given a beam with a certain wavelength λ, the phase change ϕ inside a medium of wavelength-dependent refractive index $n(\lambda)$ and thickness z is

$$\phi = \frac{2\pi n(\lambda) z}{\lambda} \qquad (10)$$

If that medium is embedded in another one (let's say, air) with refractive index n_0, then the phase difference between

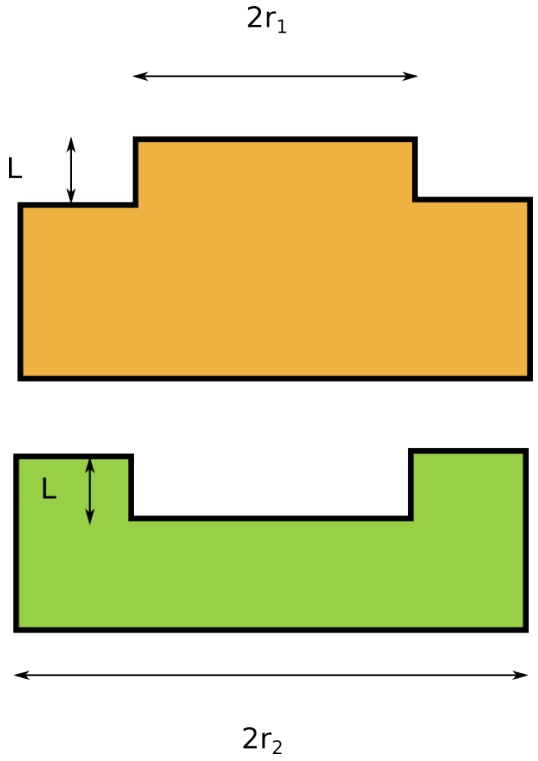

FIG. 2: Two different waveplate profiles. Both give a π phase shift between the inner and outer regions of the beam for an appropriately chosen feature length L.

a beam propagating outside as compared to the beam propagating inside is

$$\Delta\phi = \frac{2\pi\,(n(\lambda) - n_0)\,z}{\lambda} \tag{11}$$

We require a $0 \mod (2\pi)$ phase shift for a probe beam with wavelength λ_p, and a $\pi \mod (2\pi)$ shift for the trapping wavelength, λ_t. Therefore, if m_p, m_t are integer numbers, the equations for both wavelengths read

$$m_p 2\pi = \frac{2\pi\,(n(\lambda_p) - n_0)\,z}{\lambda_p} \tag{12}$$

$$(2m_t + 1)\,\pi = \frac{2\pi\,(n(\lambda_t) - n_0)\,z}{\lambda_t} \tag{13}$$

If we divide both equations, we obtain:

$$\frac{m_p}{m_t + 0.5} = \frac{(n(\lambda_p) - n_0)\,\lambda_t}{(n(\lambda_t) - n_0)\,\lambda_p} \tag{14}$$

Hence, if we decide the wavelength of our probe beam to be the center of mass of the D2 line of ^{87}Rb, $\lambda = 780.241\,\text{nm}$, knowing the wavelength dependence of the refractive index of the material will provide us with all the possible trapping wavelengths. The number of them is, of course, infinite, and the trapping wavelength will depend on m_p and m_t, which are the number of full-wavelengths of probe and trap light that fit inside the medium. We choose Zeonex as the material for the wave plate and model its refractive index as a polynomial with known coefficients [10]. Air will be the background material, with $n_0 = 1$

Fixing the probe wavelength, the right hand side of equation (14) is a continuous function of the trap wavelength. In order to be far from resonance, we will use a wavelength between 760 and 770 nm, a range of frequencies which is suitable for diode lasers. The value of this RHS is of the order of 0.97-0.98 in the region of interest, so we need to obtain values of the LHS that are close to these by changing m_t and m_p.

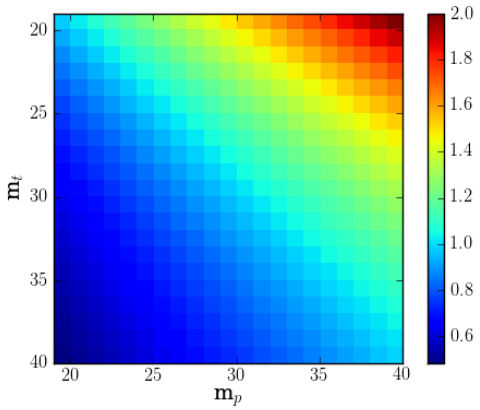

FIG. 3: Value of the LHS of equation (14) with varying m_t and m_p.

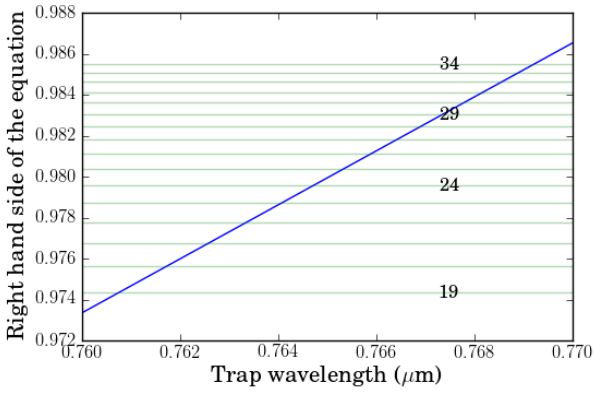

FIG. 4: Plots of the RHS of equation (14) and the LHS with equal values $m = m_p = m_t$. The numbers indicate the value of m for the line they are crossed by.

Given that the radius of the waveplate is comparable to its length in the propagation direction (as scaled to improve

numerical accuracy in the simulations), to minimize boundary effects we should minimize the length of the feature in the waveplate. We obtain this length by reorganizing (12, 13), which gives

$$z = \frac{m_p \lambda_p}{(n(\lambda_p) - n_0)} = \frac{(m_t + 0.5)\,\lambda_t}{(n(\lambda_t) - n_0)} \tag{15}$$

The smaller m_p or m_t, the smaller the length of the feature and the smaller will be the boundary effects.

If we choose m_p annd m_t to be the same, we can check that the values of the LHS are within the range of the RHS in the region of interest.

Given the probe wavelength, we chose a wavelength of 768.2373 nm. This corresponds to the 32nd order of both lights, and gives us a length for the feature of 49.9354 μm.

Furthermore, there is another requirement: the powers on the outside and the inside of the inner feature should be equal to cancel out. We can compute that number using simple integration in the 1D case: If we integrate symmetrically a Gaussian from the mean, what is the integrating domain that will give us half of the full integral? According to our definition of Gaussian beam, the waist-waveplate radius ratio should be $\gamma = 0.4769$. This will prove to be slightly different in the simulations, probably due to the effects of the feature boundaries in the waveplate.

IV. SPATIAL FREQUENCY DECOMPOSITION

In order to simulate the propagation of a Gaussian beam through such a wave plate, we use the spatial frequency decomposition method [11].

The method yields results that are similar to the Rayleigh-Sommerfeld formula for a scalar field [12], but it is more convenient as the propagation of the field E in the half-space $z > 0$, is reduced to a simple multiplication in the spatial-frequency domain. The half space is bounded by the plane of our initial condition, a complex function $E(0) \equiv E(x, y, 0)$.

We can obtain the angular spectrum of the initial condition performing a two-dimensional Fourier transform:

$$A(0) \equiv A(u, v, 0) = \mathcal{F}[E](u, v) \tag{16}$$

$$= \int\int E(x, y, 0) e^{-\mathrm{i}2\pi(ux+vy)} \mathrm{d}x\mathrm{d}y \tag{17}$$

This is the plane-wave decomposition in angular frequencies of the original field.

Assuming E to be stationary, Maxwell's equations give rise to the Helmholtz equation for the field,

$$\nabla^2 E + k^2 E = 0, \tag{18}$$

which describes the propagation of the field in space.

If we translate that equation for the angular spectrum, we obtain:

$$\frac{d^2 A}{dz^2} + \left(k^2 + k_x^2 + k_y^2\right) A = 0, \tag{19}$$

which has the solution

$$A(z) \equiv A(x, y, z) = \mathcal{H}(z) A(x, y, 0), \tag{20}$$

where the propagator

$$\mathcal{H}(z) = e^{i\sqrt{k^2 - k_x^2 - k_y^2}\, z} \tag{21}$$

appears as a multiplicative factor.

We can assume the paraxial approximation, as the waist of our beam (several μm) when compared to the wavelength ($\lambda = 780\,\mathrm{nm}$) gives a paraxiality estimator [13] $\mathcal{P} \cong 0.9$, which is close to unity. Using this assumption, the propagator reduces to

$$\mathcal{H}(z) \approx e^{ikz} e^{-i\left(+k_x^2 + k_y^2\right) z/2k}. \tag{22}$$

As we are interested in the field, we then perform the inverse Fourier transform, giving

$$E(z) = \mathcal{F}^{-1}[A(z)] = \mathcal{F}^{-1}[\mathcal{H}(z) A(0)] \tag{23}$$

$$= \mathcal{F}^{-1}[\mathcal{H}(z) \mathcal{F}[E(0)]] \tag{24}$$

.

An infinitely thin, diffractive mask with *transmittance function* $t(x, y)$, can be defined as

$$t(x, y) = \frac{E_{\text{after}}(x, y, z)}{E_{\text{before}}(x, y, z)}, \tag{25}$$

where E_{after} and E_{before} are the field after- and before the mask, respectively. We consider a medium with refractive index $n(x, y, z)$ as a set of flat layers with transmittance functions that depend on the refractive index in that layer,

$$t_i = t_i[n(x, y, z_i)], \tag{26}$$

where z_i is the position of the i-th layer. Then, we model the propagation as a sucession of free-space propagation steps followed by the application of transmittance functions.

We can take advantage of the Fast Fourier Transform (FFT) routines available in most computing languages to perform a simulation of this system.

In algorithmic language, this takes the form:

```
for each layer (i) in layers:
    Multiply E0 by transmittance: E=t(i) E0
    Fourier transform: A=FFT (E)
    Multiply by propagator: A1=A H(dz)
    Inverse Fourier transform: Ef=iFFT (A1)
    Set the initial value: E0=Ef
```

V. RESULTS

I developed a program in Python [14] that uses the spatial frequency decomposition method to test the propagation of a Gaussian beam through a custom wave plate.

The program gives qualitative results, given that the simulations were carried out in a 1D regime. The difference between 1D and 2D simulations can be easily observed if one consider the diffraction pattern due to a square slit and a circular one: in the first one one obtains a sinc dependence, whereas in the latter the cylindrical symmetry gives rise to Bessel functions. It is, in principle, possible to obtain quantitative results on the propagation of a cylindrically symmetric beam by reducing the 2D Fourier transform to a Hankel transform of order zero [15].

The field and its domain are represented as one dimensional arrays, perpendicular to the propagation direction. The domain in this radial direction is characterized by points that lie uniformly between the minimum and maximum limits, ρ_{min} and ρ_{max}, and the number of points n_ρ determines the coarseness in this dimension,

$$\Delta\rho = \rho_{max} - \rho_{min}/n_{rho}. \tag{27}$$

One needs to be very careful with spatial aliasing in the signal: Nyquist-Shannon theorem requires that the spatial sampling must be finer than half of the minimum wavelength appearing in the spectral decomposition of the signal. As we are dealing with a complex field with a wavelength of 780nm, our points should be separated, at most, by 390nm. It gives us an estimate of the minimum amount of points that we need to take into account given the limits of the domain. One then sees that simulations bigger than a few mm can become intractable very fast.

We check the propagation through the waveplate and observe that the central part of the trap beam undergoes a π phase shift. As we can see in Figure 5, the edges of the waveplate create disturbances in the beam as it propagates. If we choose to propagate a beam with the wavelength of the probe, the central part of the beam does not undergo a phase shift; however, the *ringing* near the edges of the feature persist, as shown in Figure 6.

In the case that the waveplate is close to a thin lens, the Fourier transform of this function is proportional [11] to the value of the field in the focus, f, of the lens. The spatial scaling will be

$$x' = k_\rho \lambda f, \tag{28}$$

where λ is the wavelength of the field and k_ρ is the spatial frequency, measured in m^{-1}. In Figure 7, where we have performed the FFT on the field right after the waveplate, we can see that the trapping potential (for a suitable set of parameters) resembles a quartic polynomial near the axis

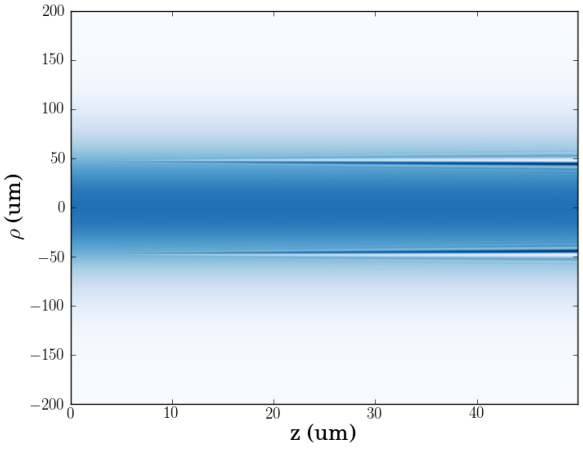

FIG. 5: The edges of the waveplate create disturbances in the propagation of the otherwise Gaussian beam.

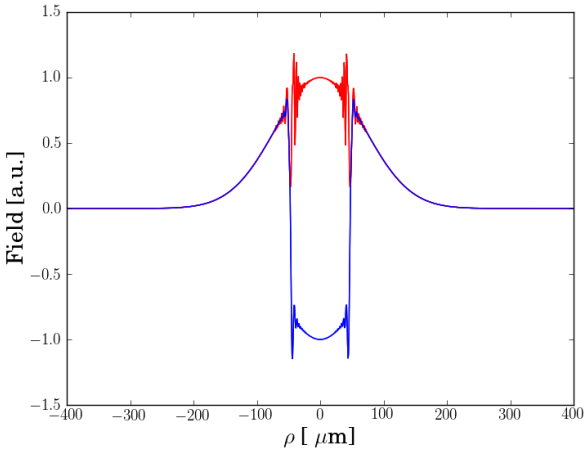

FIG. 6: Profiles of the electric field of the beams at the end of the waveplate. In blue, the trap beam undergoes a π phase shift in the central region and shows perturbances due to the edges of the waveplate. In red, the probe beam shows no such phase shift, but the effect of the edges is visible.

at the focal spot of a lens, whereas the probe beam fits neatly between the maxima of the trapping potential.

The first thing to test is the size of the waveplate that provides the maximum extinction in the center. We define the extinction as the ratio between the power at the focus and the power at the neighbouring maxima. If we assume that the outer radius of the waveplate is comparable to the waist, we can change both the inner (r_1) and the outer (r_2) radii of the waveplate and obtain the conditions for minima at the focus. In Figure 8 we can see that, when we change r_2 to be much greater than the other sizes, the

waist-inner radius ratio γ that gives the maximum extinction is $\gamma = 0.472$, very close to the prediction using the cumulative distribution function of a Gaussian. I assume that the discrepancy between these two results, close to 1%, comes from the effects of the central feature boundaries on the propagation.

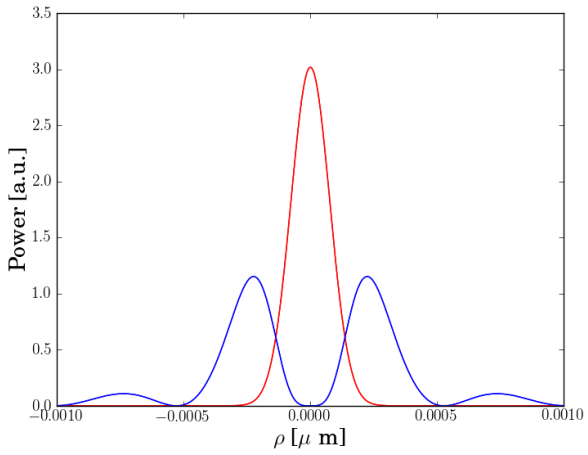

FIG. 7: These are the profiles of the powers at the focal point of the lens. The profile of the trapping beam (blue) resembles a quartic at the origin. The probe beam (red) is confined within the peaks of the trapping potential.

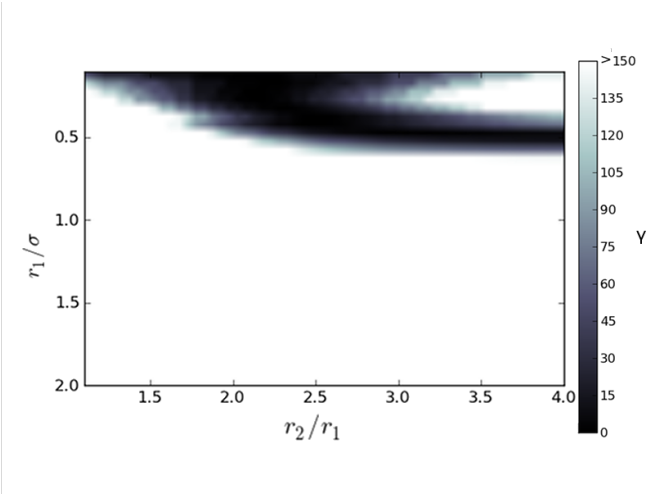

FIG. 8: Shows different values of the waist to radius ratio (γ) for different waveplate inner (r_1) and outer (r_2) radii with fixed beam waist. The minima reaches an asymptotic value, assuming $r_2 \gg r_1$, of $\gamma = 0.472$.

Using a refractive index $n = 1.5$ for both wavelengths and the waveplate described above, a beam waist of $\sigma = 10^{-4}$m and $\gamma = 0.472$, considering the transformation produced by a thin lens of focal $f = 10\,\text{cm}$, we obtain an extinction of $\eta < 5 \cdot 10^{-4}$.

It is possible to perform the propagation of the field obtained at the focus, as we are interested in knowing the longitudinal shape of the trap. But given that the results of the 1D simulation do not follow the laws of propagation of radially-symmetric, 2D beams, the results that we get are meaningless in terms of scaling. Nonetheless, performing such simulation we can obtain qualitative results about the shape of the trap. Using an unrealistic 0.1mm lens, we see in Figure beam evolves into a shape that creates a potential along all directions. The smallest barrier corresponds to about 25% of the maximum height, and occurs along the diagonals.

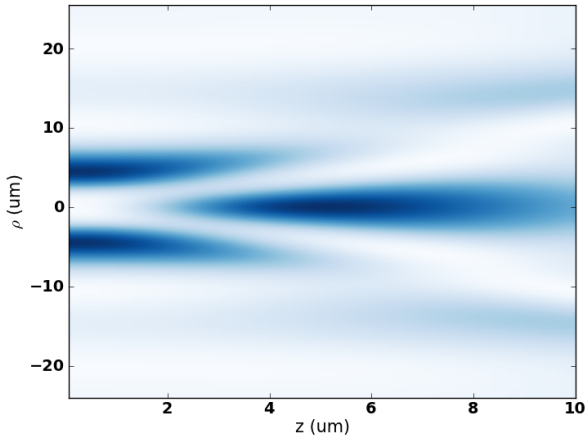

FIG. 9: We can obtain qualitative results about the shape of the trap propagating the result of the FFT. It shows that potential effectively traps the particles in the minimum, being the smallest trapping height that of potential along the diagonals.

VI. DISCUSSION

When discussing the results outlined above one needs to bear in mind that the simulations we have performed hold only for the 1D case.

On the computational side, the amount of points that define our potential after we have performed the FFT depend on the ratio between the size of the window and the size of the initial Gaussian: the FFT will transform big features (with respect to the window size) into small features and vice versa. This causes troubles when we need to consider milimetre-sized simulations: we need to have a resolution high enough to satisfy Nyquist-Shannon theorem and, at the same time, the window needs to be so big that the resolution in the FFT plane satisfies those conditions as well. The number of points to work with lies on

the range $N \approx 10^4 - 10^6$. A single FFT of an array with $2^{23} \sim 10^7$ elements takes about a minute to run on a Pentium IV at 3 GHz. That makes 2D simulations unfeasible at those scales using a single processor, where we need to consider matrices with N^2 complex elements. A possible solution might be to use the Hankel transform, as stated in the previous section, to obtain the correct spread of the profile of the beam as it propagates in space, but this has not been implemented yet.

In our experiment, we need the atoms to absorb the probe to block it completely; that is, that the optical depth of our sample must be as high as possible. If we take a look at Figure 6, we see that the probe beam fits neatly between the two maxima of the trapping beam. This situation is not entirely satisfactory: the trapping region will actually be smaller than the region bounded by the two maxima - about one third of the distance, as a rule of thumb. In this case, the atoms will not be able to block completely the power from the probe beam, thereby reducing the maximum absorption that we can obtain from our system. The ideal case would be where the waist of the probe beam is about 3 or 4 times smaller than the distance between the peaks. Therefore, this system is, at this stage, not suitable for our experiment, as any spurious transmission can mask the effect we are trying to measure.

Finally, even though the system described in this report is not suitable for our purposes, it might be interesting to try it in a system where the signal to noise ratio is not so critical. The compact solution provided by this waveplate could reduce the size of an apparatus that requires trapping and probing simultaneously. For example, one can think of portable aerosol trapping and probing.

[1] J. D. Pritchard, D. Maxwell, A. Gauguet, K. J. Weatherill, M. P. A. Jones, and C. S. Adams, Phys. Rev. Lett. **105**, 193603 (2010).
[2] L. Isenhower, W. Williams, A. Dally, and M. Saffman, Opt. Lett. **34**, 1159 (2009).
[3] P. Xu, X. He, J. Wang, and M. Zhan, Opt. Lett. **35**, 2164 (2010).
[4] W. Phillips, Reviews of Modern Physics **70**, 721 (1998).
[5] R. Grimm, M. Weidemuller, and Y. B. Ovchinnikov (Academic Press, 2000) pp. 95 – 170.
[6] S. Zhang, F. Robicheaux, and M. Saffman, Phys. Rev. A **84**, 043408 (2011).
[7] T. N. Bandi, V. G. Minogin, and S. N. Chormaic, Phys. Rev. A **78**, 013410 (2008).
[8] G. D. Bruce, S. L. Bromley, G. Smirne, L. Torralbo-Campo, and D. Cassettari, Phys. Rev. A **84**, 053410 (2011).
[9] J. L. Chaloupka and D. D. Meyerhofer, J. Opt. Soc. Am. B **17**, 713 (2000).
[10] "Refractive index database," `http://refractiveindex.info/`.
[11] J. W. Goodman, *Introduction to Fourier Optics*, 3rd ed. (Roberts & Company Publishers, 2004).
[12] G. C. SHERMAN, J. Opt. Soc. Am. **57**, 546 (1967).
[13] P. Vaveliuk, B. Ruiz, and A. Lencina, Opt. Lett. **32**, 927 (2007).
[14] There is an important point to bear in mind when performing FFTs in any language: test the scaling of the FFT with the size of the arrays. Make sure that it is corrected. In the case of Python, the correction consists, simply, in multiplying by the size of the divisions of the array.
[15] M. Birkinshaw, in *Astronomical Data Analysis Software and Systems III*, Astronomical Society of the Pacific Conference Series, Vol. 61, edited by D. R. Crabtree, R. J. Hanisch, & J. Barnes (1994) p. 249.