Adaptive Grid Refinement For a Model of Two Confined and Interacting Atoms ⋆

William F. Mitchell [a,*]

[a] *Mathematical and Computational Sciences Division*
100 Bureau Drive Stop 8910
National Institute of Standards and Technology
Gaithersburg, MD 20899-8910

Eite Tiesinga [b]

[b] *Atomic Physics Division*
100 Bureau Drive Stop 8423
National Institute of Standards and Technology
Gaithersburg, MD 20899-8423

Abstract

We have applied adaptive grid refinement to solve a two-dimensional Schrödinger equation in order to study the feasibility of a quantum computer based on extremely-cold neutral alkali-metal atoms. Qubits are implemented as motional states of an atom trapped in a single well of an optical lattice of counter-propagating laser beams. Quantum gates are constructed by bringing two atoms together in a single well leaving the interaction between the atoms to cause entanglement. For special geometries of the optical lattices and thus shape of the wells, quantifying the entanglement reduces to solving for selected eigenfunctions of a Schrödinger equation that contains a two-dimensional Laplacian, a trapping potential that describes the optical well, and a short-ranged interaction potential. The desired eigenfunctions correspond to eigenvalues that are deep in the interior of the spectrum where the trapping potential becomes significant. The spatial range of the interaction potential is three orders of magnitude smaller than the spatial range of the trapping potential, necessitating the use of adaptive grid refinement.

Key words: adaptive grid refinement, elliptic eigenvalue problem, quantum computing, Schrödinger's equation

⋆ Contribution of NIST; not subject to copyright.

* Corresponding author.
Email address: william.mitchell@nist.gov (William F. Mitchell).

1 Introduction

The idea of using the rules of quantum mechanics as a paradigm for computing has started a flurry of research over the last five to ten years [14]. Entanglement of "quantum" bits, the quantum analog of a classical bit, promises to be a significant advantage with the potential of providing a tool for the solution of problems that are intractable on classical computers. The discovery of a quantum algorithm for factorization [17] has renewed interest in research into the foundations of computer science, and has far reaching implications in cryptography, computer security, and e-commerce. Advances in diverse fields of physics such as ultra-cold physics on atoms and ions, nuclear magnetic resonance imaging (NMR), and nanofabrication of quantum dots and Josephson junctions have led to several realizations of a quantum bit and quantum gates, the quantum analog of one- and two-bit operations. So far, a quantum computer of seven quantum bits has been shown to work using NMR techniques [21].

Currently, limitations of the physical realizations pose a serious challenge for increasing the number of quantum bits by an order of magnitude. Decoherence, i.e. unwanted coupling to the environment, changes the fragile entanglement irreversibly and thereby makes quantum computing hard. Nevertheless, upper bounds for an acceptable rate of decoherence have been derived [14] and show that quantum computing is still feasible. Theoretical research is not only directed toward the understanding of the decoherence but also toward the characterization of coherent behavior of a small number of quantum bits during the quantum gate operation.

For this paper we are interested in modeling a quantum gate with quantum bits that are based on ultra-cold atoms [15,18]. Ultra-cold atoms can be confined by counterpropagating laser beams. The light creates a three-dimensional washboard potential or optical lattice. A single atom is held in each potential minimum or lattice site of the washboard. Two energy levels of an atom form a quantum bit. By bringing two atoms from individual sites together and having them collide or interact to cause entanglement, a two-quantum-bit operation, or quantum gate, can be realized.

First-principle modeling of the interactions of two atoms in a lattice site is numerically challenging, as length scales for the lattice and the mutual atom-atom interaction differ by orders of magnitude, creating the need for carefully chosen spatial grids. This paper presents the application of the adaptive grid techniques in [10,11] to the finite element solution of the Schrödinger's equation that models this interaction.

The remainder of the paper is organized as follows. In Sec. 2 we define the

mathematical model of the interaction of atoms in a trap. Section 3 describes the numerical methods used to solve the equations of this model. In Sec. 4 we describe the adaptive processes involved in the solution method. Section 5 presents the results of applying these methods to solve a model of two Cesium atoms in a trap. Finally, we conclude in Sec. 6.

2 Mathematical Model

For appropriately chosen laser intensities and geometries the lattice sites are, to good approximation, harmonic and cylindrically symmetric. This leads to a model for the relative motion of two atoms in a single site by a two-dimensional Schrödinger equation in cylindrical coordinates, i.e., an elliptic eigenvalue problem of the form

$$-\frac{1}{\rho}\frac{\partial}{\partial\rho}\rho\frac{\partial\psi}{\partial\rho}-\frac{\partial^2\psi}{\partial z^2}+V(\rho,z)\psi=\lambda\psi \text{ in } (0,\infty)\times(0,\infty)\,, \tag{1}$$
$$\frac{\partial\psi}{\partial n}=0 \text{ on } z=0 \text{ or } \rho=0\,,$$
$$\psi\to 0 \text{ as } r\to\infty\,.$$

Here, λ is the energy eigenvalue, the atom-atom separation is $r=\sqrt{\rho^2+z^2}$, and $V(\rho,z)=V_t(\rho,z)+V_{aa}(r)+(m/\rho)^2$, where V_t is the trapping potential that describes the optical well, V_{aa} is a short-ranged atom-atom interaction potential, and $(m/\rho)^2$, with quantum number m a nonnegative integer, is a consequence of the cylindrical symmetry (henceforth we use $m=0$).

For the purposes of this paper the atom-atom interaction is modeled as a Lennard-Jones potential

$$V_{aa}(r)=4\epsilon(\frac{\sigma^{12}}{r^{12}}-\frac{\sigma^6}{r^6})\,. \tag{2}$$

As illustrated in Fig. 1, the length σ determines the location of the potential well and the energy ϵ determines its depth. In addition, V_{aa} is zero for $r=\sigma$ and the minimum of the potential is $-\epsilon$ at $r=\sqrt[6]{2}\sigma$. At large r, the atom-atom potential approaches $-C_6/r^6$, where $C_6=4\epsilon\sigma^6$ is the van-der-Waals dispersion coefficient. A realistic atom-atom potential also satisfies $V_{aa}\propto -1/r^6$ for large r, but will have a complex short-range shape. (See for example [20]).

The trapping potential is given by

$$V_t(\rho,z)=\rho^2+\left(\frac{\omega_z}{\omega_\rho}\right)^2 z^2\,, \tag{3}$$

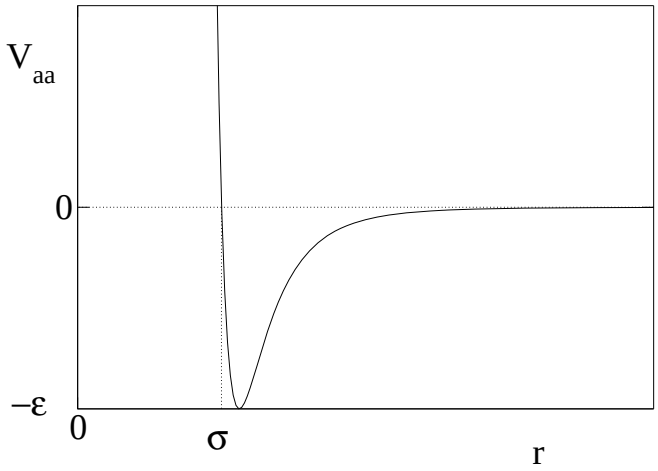

Fig. 1. Schematic of the atom-atom potential $V_{aa}(r)$.

where ω_ρ and ω_z are the trapping frequencies along the ρ and z direction, respectively. The ratio ω_z/ω_ρ determines the shape of the trap. If this ratio is 1.0, the trap is spherical (when rotated about the z axis). If it is greater than 1.0, then the trap is elongated in the ρ direction and "pancake" shaped. If it is less than 1.0, then the trap is elongated in the z direction and "cigar" shaped.

Equation 1 has been expressed in terms of dimensionless length and energy scales given by

$$l = \sqrt{\frac{\hbar}{\mu\omega_\rho}} \tag{4}$$

for the unit of length and

$$U = \frac{1}{2}\hbar\omega_\rho = \frac{\hbar^2}{2\mu l^2}\,, \tag{5}$$

for the unit of energy. Here, $\hbar$ is Planck's constant and μ is the reduced mass of the two interacting atoms. These units are the "natural" length and energy scales for a trapping potential with frequency ω_ρ. For example, in the absence of V_{aa}, Eq. 1 can be solved analytically and the smallest eigenvalue is $2 + (\omega_z/\omega_\rho)$, which is between 2 and 100 for typical values of ω_z/ω_ρ.

The semi-infinite domain is truncated to a rectangle $(0, R_z) \times (0, R_\rho)$ with sufficiently large R_z and R_ρ, and homogeneous Dirichlet boundary conditions are applied on $z = R_z$ and $\rho = R_\rho$.

For eigenpairs of interest λ is positive and small compared to the $-\epsilon$ depth of the potential. For these eigenenergies the scattering length, a, and the van-

der-Waals length

$$x_0 = \frac{1}{2}\sqrt[4]{C_6} = \frac{1}{2}\sqrt[4]{4\epsilon\sigma^6}\,, \tag{6}$$

summarize the interaction potential. For realistic values of σ and ϵ, $\sigma \ll x_0 \ll 1$.

For given σ and ϵ, the scattering length can be computed by solving Eq. 1 in the absence of V_t. The equation reduces to an ordinary differential equation (ODE) in the atom-atom separation r,

$$\begin{aligned} \frac{d^2\phi}{dr^2} &= (V_{aa}(r) - \lambda)\phi \quad r > 0\,, \\ \phi(0) &= 0\,, \\ \frac{d\phi}{dr}(0) &= 1\,. \end{aligned} \tag{7}$$

In physics and chemistry this corresponds to a scattering problem. At large r, ϕ exhibits a periodic behavior $B\sin(\sqrt{\lambda}[r-a])$, where a is the scattering length and B is an (irrelevant) normalization constant. The wave function ϕ reaches its periodic behavior for r larger than the van-der-Waals length. Figure 2 illustrates the solution for a wave function with $\lambda = 1$, $a = 0.1$, and $x_0 = 0.87$. The wave function oscillates very rapidly for $r < x_0$ and much slower for $r > x_0$ indicating the need for adaptive grids.

The scattering length can be interpreted as the effective size of the colliding atoms and can be used to construct approximate delta-function-like atom-atom potentials that have successfully been used in analytical solutions of two-atom Schrödinger equations as well as a starting point for many-atom theories. (See [4] for a recent use of the approximate potential.)

In addition, one can determine the number of nodes, N, of the solution of Eq. 7 for $\lambda = 0$. In practice, for $\lambda > 0$ it is simpler to count the number of extrema that occur before the large-r periodic behavior begins.

In our software the ratio ω_z/ω_ρ, ϵ and σ can be entered. Typical values for ω_z/ω_ρ range between 0.01 and 100. Realistic values for σ and ϵ are $\sim 10^{-3}$ and $\sim 10^9$, respectively. In other words, the extent of the atom-atom interaction is orders of magnitude smaller than the natural length of the trapping potential and the depth of the potential is orders of magnitude larger than natural energy of the trap. Both relationships indicate the need for a carefully chosen grid.

We find it convenient to define the atom-atom potential in alternative ways. In all cases, the ratio ω_z/ω_ρ must be given. Equations 6 and 7 show that ϵ and

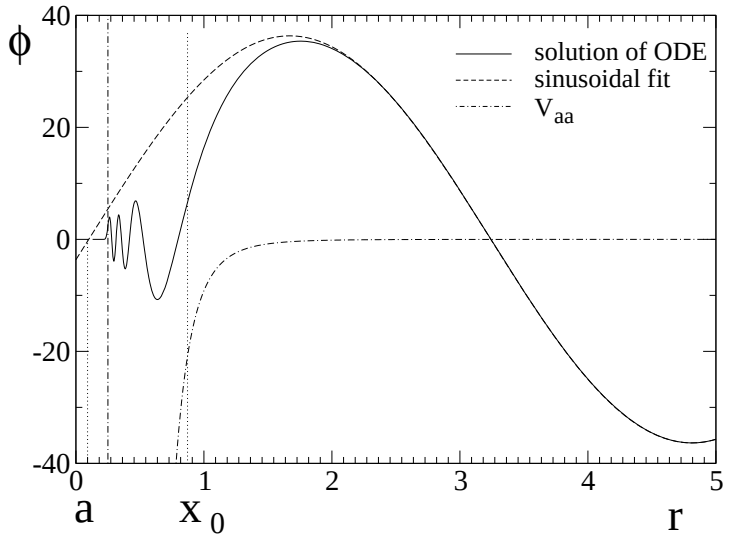

Fig. 2. The solution of the ODE, using $\lambda = 1$, is shown with a solid line. The interaction potential is shown with a dash-dot line. The vertical dotted lines indicate the scattering length a and the van-der-Waals length, x_0, respectively. The dashed wave function, $B\sin(\sqrt{\lambda}[r-a])$, is a fit to the ODE solution for $r \gg x_0$. Here $a = 0.1$, $x_0 = 0.87$, and $N = 6$.

σ uniquely define x_0, a and N. It is also the case that x_0, a and N uniquely define ϵ and σ. By solving Eq. 6 for σ and substituting into Eq. 2, we can write V_{aa} in Eq. 7 as a function of x_0 and ϵ. Then, for the given x_0 we use a bisection root finding algorithm on Eq. 7 to determine the ϵ that gives the desired a and N. σ is then computed from Eq. 6.

There are additional combinations of parameters that uniquely define the atom-atom potential, including 1) σ and x_0, 2) σ, a and N, and 3) C_6, ϵ, ω_ρ and μ.

3 Solution Method

We solve the two dimensional Eq. 1 using finite elements with piecewise linear functions over triangles. The adaptive grid approach described in Sec. 4 automatically finds a spatial representation that is refined where the atoms are close together and the atom-atom interaction potential is very strong and negative, and is coarse at larger separation where the trapping potential is largest in value or energy. The standard finite element discretization [19] results in a generalized eigenproblem

$$Ax = \lambda Mx \tag{8}$$

where A is the stiffness matrix with entries that are energy inner products of the basis functions, M is the mass matrix with entries that are L^2 inner products of the basis functions, λ is an eigenvalue (or eigenenergy), and x is the corresponding eigenvector, the discrete form of an eigenfunction (or wave function), ψ. This eigenproblem is real and symmetric, but highly indefinite, i.e., it has many negative eigenvalues as well as positive eigenvalues.

We are interested in obtaining a small number of wave functions whose eigenvalues are closest to zero. These are called the trap states because they are eigenstates in which the eigenfunction extends from short to large atomic separations and the trapping potential plays an important role. To compute the eigenvalues, and corresponding eigenvectors, closest to some given number λ_0, we use the standard "shift-and-invert" transformation [5] which converts Eq. 8 to the eigenproblem

$$(A - \lambda_0 M)^{-1} M x \;=\; \nu x \tag{9}$$

where $\nu = 1/(\lambda - \lambda_0)$. The largest (in magnitude) eigenvalues of Eq. 9 correspond to the eigenvalues of Eq. 8 that are closest to λ_0, and the eigenvectors remain unchanged.

To compute a small number of the largest eigenvalues of Eq. 9 we must use an iterative eigensolver to avoid explicitly computing $(A - \lambda_0 M)^{-1}$. One wants a method that can compute a given number of eigenvalues at the ends of the spectrum, and only requires multiplication of a vector by the matrix. One such method is Arnoldi's method [5], which finds eigenvectors in a low-dimensional Krylov space, although other possibilities exist.

To multiply a vector y by $(A - \lambda_0 M)^{-1} M$, one first multiplies by M, and then solves the linear system

$$(A - \lambda_0 M)x = My \tag{10}$$

for x. When λ_0 is deep inside the spectrum, $A - \lambda_0 M$ is symmetric but highly indefinite. Currently, we solve this system with a direct method. Ongoing research is seeking a good preconditioner to use with GMRES [16] or another iterative method for indefinite systems.

4 Adaptive Methods

In this section we describe the adaptive methods used in the solution of Eq. 1. This is divided into three parts: adaptive grid refinement, sequence of pseudopotentials, and selection of next solution step.

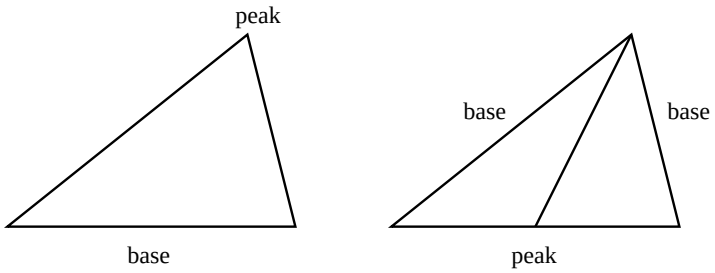

Fig. 3. Newest node bisection of a triangle.

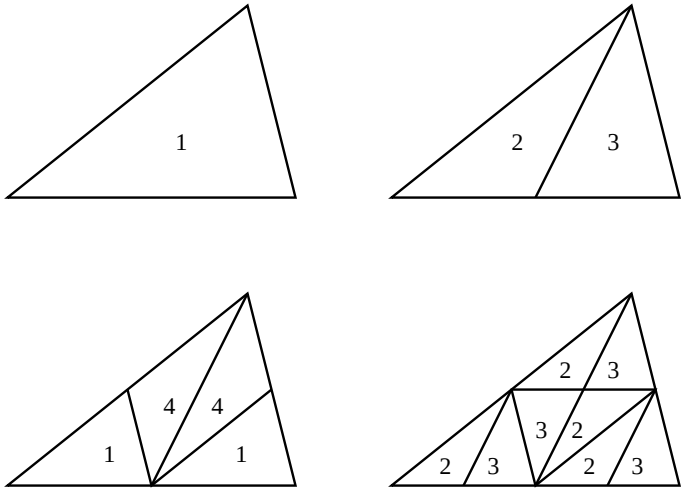

Fig. 4. The four shapes of triangles that arise during repeated application of newest node bisection.

4.1 Adaptive Grid Refinement

The triangulation is adaptively refined by the newest node bisection method. Here we briefly review the method. A detailed description of the method can be found in [10], and its parallelization in [11].

A triangle is divided to form two new triangles by connecting one of the vertices, called the peak, to the midpoint of the opposite side, called the base, as in Fig. 3. The new vertex created at the midpoint of the base is assigned to be the peak of the new triangles. As illustrated in Fig. 4, repeated application of this refinement will result in only four triangle shapes, and therefore the angles are bounded away from 0 and π. It can be shown with simple geometry that triangles labeled with the same number have the same shape, and further refinements continue with these shapes.

It is important that the final triangulation be compatible, i.e., given any two triangles their intersection is either empty, a common vertex, or a common side. Compatibility can be maintained during the refinement process by dividing pairs of triangles rather than individual triangles. A triangle is said to

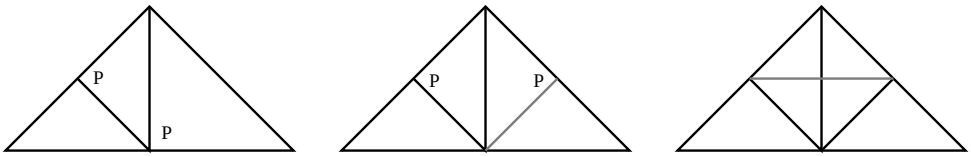

Fig. 5. Maintaining compatibility during refinement. Peaks are labeled P.

be compatibly divisible if its base is either the base of the triangle that shares that side or part of the boundary of the domain. If a triangle is compatibly divisible, then divide the triangle and the neighbor that shares the base simultaneously as a pair. If the triangle is not compatibly divisible, then after a single bisection of the neighbor it will be. So in this case, first divide the neighbor by the same process, and then divide the triangle and the neighbor simultaneously, as illustrated in Fig. 5.

To determine which triangles should be refined, one computes an *a posteriori* local error indicator, and refines those elements that have the largest indicators. It is not necessary that an error indicator be a good estimate of the local error, only that it indicate which triangles would most benefit from refinement. However, a local error estimator can be used as an error indicator. Many error indicators and estimators have been proposed and used [22].

In this paper we use a very simple hierarchic error indicator, which is just the difference between the solution at the peak of the triangle and the interpolation of the solution at the older two neighbors of the peak. It corresponds to the coefficient of a hierarchical basis centered at the peak [10]. This error indicator has the advantage of being very cheap to compute, and the numerical results in Sec. 5 demonstrate that it is effective for the type of problems considered in this paper.

Since we are interested in computing a small number of wave functions, not just one, we need a slight modification to the usual use of error indicators. It is important that the grid be fine enough in the right places for an accurate computation of *all* the solutions. The usual error indicator is a function of the solution, so we simply compute the error indicator of a triangle for each of the solutions, and take the maximum of those values as the error indicator for that triangle. Then the triangle will be refined if any of the solutions is sufficiently inaccurate in the vicinity of that triangle.

In addition to guiding adaptive refinement, an error estimate can also be used as a termination criterion for the refinement process. For this purpose we compute a simple extrapolation error estimate to approximate the L^∞ norm of the error. It turns out that this error estimate is just the maximum over all error indicators. The following argument is approximate, and this error estimate is not particularly accurate, but we usually find it to be correct within a factor of 2 to 4 for elliptic boundary value problems where the solution is

known, and not unreasonable to use as a termination criterion.

Consider the error indicator, $\tilde{e}_t$, for triangle t with longest side length h. $\tilde{e}_t = \psi_{t,H} - \psi_{t,h}$ where $\psi_{t,h}$ is the computed solution at the peak of t, and $\psi_{t,H}$ is the interpolation of the older two neighbors of the peak, in other words the solution we would have at that point with a local grid size of $H = \sqrt{2}h$. Convergence theory for linear finite elements with a sufficiently smooth solution tells us that the error in the L^∞ norm is second order, i.e., $e_{t,h} = \mathrm{O}(h^2)$ where $e_{t,h} = \psi_{t,h} - \psi_t$ and ψ_t is the exact solution at the peak of t. Thus, we can assume that $e_{t,h} \approx ch^2$ and $e_{t,H} \approx cH^2$ for some constant c. Then

$$\tilde{e}_t = \psi_{t,H} - \psi_{t,h} = \psi_{t,H} - \psi_t + \psi_t - \psi_{t,h} = e_{t,H} - e_{t,h} \approx c(H^2 - h^2)$$

Thus

$$c \approx \frac{e_{t,H} - e_{t,h}}{H^2 - h^2} = \frac{\tilde{e}_t}{(\sqrt{2}h)^2 - h^2} = \frac{\tilde{e}_t}{h^2}$$

which gives

$$\tilde{e}_t \approx ch^2 \approx e_h$$

and therefore the error indicator is an estimate of the local pointwise error, and the maximum of the absolute values of the error indicators is an estimate of the L^∞ norm of the error.

Also, it was shown in [9] that an estimate of the energy norm of the error can be computed from the error indicators by

$$||e_h||_E^2 \approx \Sigma \tilde{e}_t^2$$

Thus we could use an estimate of either the L^∞ norm or the energy norm of the error as a termination criterion.

4.2 Pseudopotentials

As described in the previous section, the adaptive refinement process begins with a very coarse uniform grid, and automatically adapts the refinement of the grid to the solution. For the interacting atoms in a trap problem, we must choose the truncated domain large enough to contain the large scale waves due to the trap, which means $R_\rho = O(1)$ and $R_z = O(1)$. But the width of the well from the Lennard-Jones potential is $O(\sigma)$ which is $O(10^{-3})$ for the

problems we are interested in. This is orders of magnitude smaller than the size of the triangles in the initial grid, and the solution-based error indicator will not detect this important feature until, by luck more than anything else, the elements containing the well are refined many times.

To overcome this problem of getting the refinement started in the right area we use a sequence of pseudopotentials. These are potentials of the same form as Eqs. 2 and 3, but with milder parameters so that the features of the potential can be detected with coarser grids. For the first pseudopotential we use a very wide and shallow well, i.e. relatively large σ and small ϵ. The wider well will create waves in the solution at a scale that can be detected on the coarse initial grid, while the shallowness causes the number of nodes to be small so that this problem is not too difficult, and does not require large resources to solve. Then we go through some number of increasingly difficult problems, with smaller σ and larger ϵ until we reach the correct potential. This sequence draws the refinement down toward the well of the final potential.

Naively choosing a sequence of values for σ and ϵ will result in erratic behavior of the large scale waves in the solution, making it more difficult for the adaptive grid to settle down on the scale of the trapping potential. So one must carefully choose a strategy to define a trajectory through the space of pseudopotentials. For this paper we use a sequence that smoothly reduces x_0 and simultaneously increases the number of nodes, N. There are other possibilities, and it is not clear which one will be an effective trajectory that minimizes the extra resources used to guide the adaptive refinement in this way. In future work we plan to study this problem.

The sequence of pseudopotentials is defined with the help of the Wentzel-Kramer-Brillouin (WKB) approximation of Eq. 7 [8]. The WKB model leads to an analytic relationship between the (approximate) number of nodes (bound states) of the $\lambda = 0$ solution of Eq. 7 and σ and ϵ. In fact, we have

$$N_{\mathrm{WKB}} = \frac{1}{\pi} \int_{\sigma}^{\infty} \sqrt{-V_{aa}(x)}\, \mathrm{d}x = \frac{C}{\pi} \sigma \sqrt{4\epsilon}\,, \tag{11}$$

where $C \approx 0.42$ is determined by numerical integration. From numerical experiments we find that the absolute difference between the real-valued N_{WKB} and the actual integer number of nodes N is less than 1. It then follows that for given x_0 and N_{WKB} we can invert Eqs. 6 and 11 and find

$$\begin{aligned} \sigma &= 2x_0 \sqrt{\frac{C}{\pi N_{\mathrm{WKB}}}} \\ \epsilon &= \frac{\pi^3 N^3}{16 x_0^2 C^3}\,. \end{aligned} \tag{12}$$

In practice, the user defines the final values σ_f and ϵ_f by specifying the atom-atom potential as in Sec. 2, as well as the number of pseudopotentials n. The final number of nodes N and x_0 are therefore known. We then construct an initial pseudopotential with $N_{\mathrm{WKB}} = 1$ and a van-der-Waals length of $N \times x_0$. This uniquely defines the initial atom-atom potential. Using the factor $f = N^{1/n}$, each subsequent pseudopotential is defined by $N_{\mathrm{WKB}} \leftarrow f N_{\mathrm{WKB}}$ and $x_0 \leftarrow x_0/f$.

4.3 Selecting the Next Solution Step

In addition to adapting the grid to the solution, there are several other parameters that need to be determined adaptively. These are the number of iterations of an iterative linear system solver for Eq. 10, the size of the Krylov space for the eigensolver, and the maximum number of iterations for the eigensolver. We assume that as the problem gets harder (i.e., σ gets smaller and ϵ gets larger), these parameters will increase, so we only adapt these numbers upward and never decrease them. As illustrated in the flowchart of Fig. 6, the outer level of the algorithm is a loop in which Eq. 1 is approximately solved, and a determination is made of what to change the next time through the loop. The loop is exited when a sufficiently accurate solution is found for the final potential.

There are three tolerances that help guide this process. τ_{ls} is a bound on how large the residual of the linear system can be, τ_e is a bound on how large the residual of the eigenproblem can be, and τ_d is a bound on how large an estimate of the discretization error can be. As pointed out in Sec. 4.1, τ_d can be related to either the L^∞ or energy norm of the discretization error. One could start with relatively large values for these tolerances and adapt them downward as you progress through the pseudopotential trajectory to reduce the computation in the early pseudopotentials, but for this paper we will keep the tolerances constant.

To determine the next solution step, several tests are performed to determine if a parameter should be changed, the grid should be refined, or the next pseudopotential should be used.

The first test is for accuracy of the linear system solver. Here we compute the ℓ^2 norm of the residual of Eq. 10, $||(A-\lambda_0 M)x - My||_2/||My||_2$, for each such system that is solved in the course of the eigensolver, and compare the average of them to τ_{ls}. If the linear system is not solved with sufficient accuracy, the eigensolver will fail, so this is checked first. Of course, if a direct solver is used then this test is skipped.

Second, we check if the eigensolver converged. Usually the eigensolver software

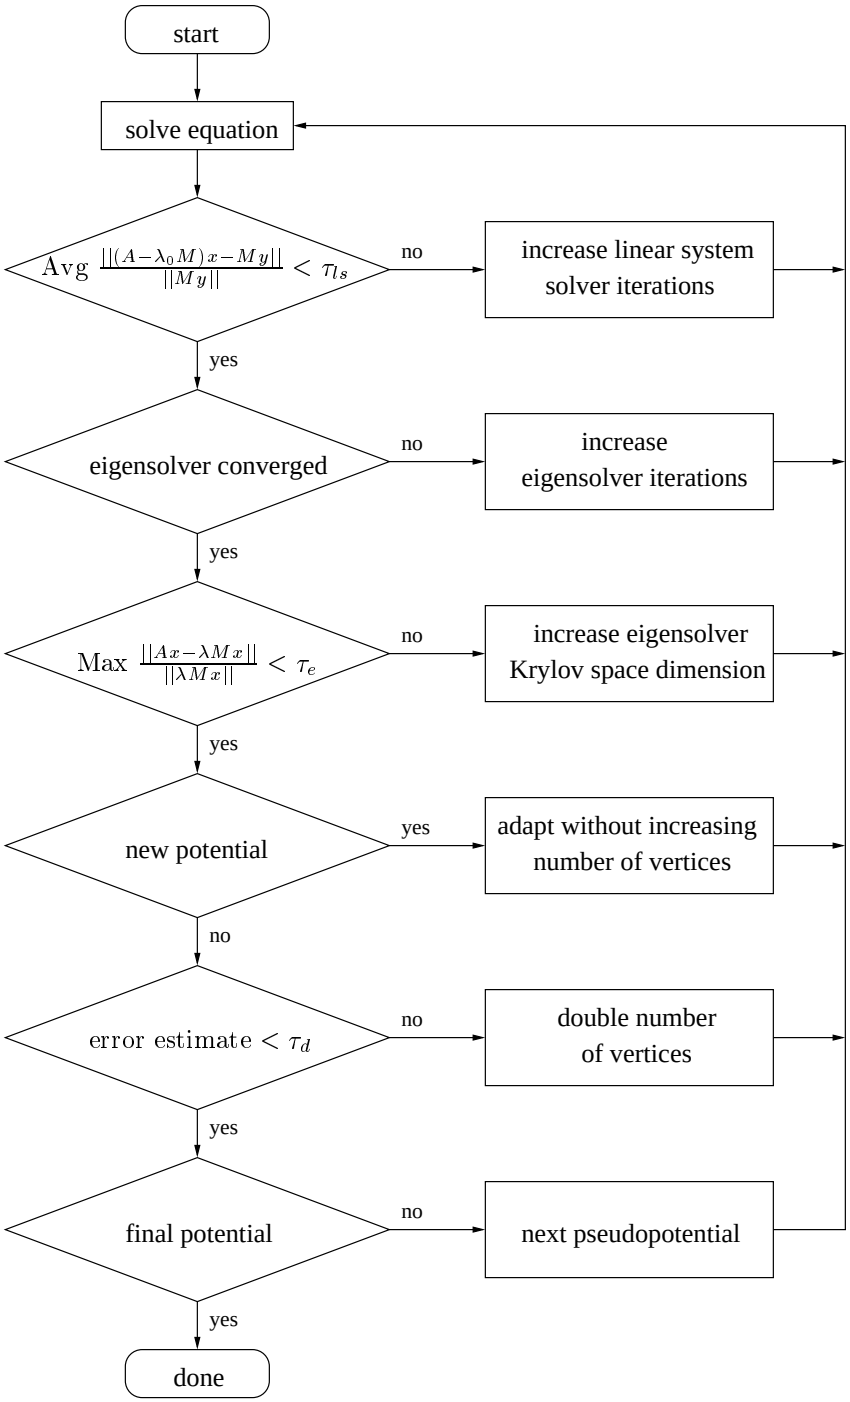

Fig. 6. Flowchart of the process for determining the changes for the next solution.

will report whether or not convergence was achieved.

If it did converge, then the third test is whether the eigensolver gave sufficient accuracy. For this we compute the norm of the eigensystem residual $||Ax - \lambda Mx||_2/||\lambda Mx||_2$ for each eigenvalue computed, and compare the largest with τ_e. If this test fails, then the size of the Krylov space for the eigensolver is

increased.

If the first three tests are passed, then we have confidence that the eigenproblem of Eq. 9 has been solved for the current pseudopotential and grid.

The fourth test prevents an unnecessary increase in the number of vertices. The first time Eq. 1 is solved with a new pseudopotential, it is solved on a grid that was adapted for the previous pseudopotential. Since the location of the well in the interaction potential (which is roughly where the grid should be finest) has moved, the grid is fine in the wrong place. The first solution is likely to be inaccurate enough that the error estimate will be larger than τ_d, but hopefully good enough to indicate where the grid should be finer. So on the second solution with a given pseudopotential, the grid is adapted while keeping the same number of vertices. This coarsens the grid around the well of the old pseudopotential and refines it around the well of the the new pseudopotential.

The fifth test checks the error estimate of Sec. 4.1 to see if the grid should be refined. Finally, if all these tests are passed then we move on to the next pseudopotential.

5 Numerical Results

We solved Eq. 1 using PHAML 0.9.11 [12,13], a parallel implementation of the methods described in Sections 3 and 4. The eigenproblems were solved with parallel ARPACK [5–7], which implements Arnoldi's method. To solve the indefinite linear system at each iteration of the eigensolver we used MUMPS 4.3 [1–3], a parallel sparse direct solver.

Computations were performed on 8 processors of a PC cluster operating under the Red Hat 7.3 distribution of Linux [1]. Each node of the cluster is a dual 1.66 GHz AMD Athlon MP 2000 processor with 4 Gbyte RAM. The programs were compiled with PGI Fortran 90 version 4.0. The message passing library was the LAM 6.5.6 implementation of MPI.

We solved Eq. 1 with a potential relevant to Cesium. The parameters are given in Table 1. The top section of the table gives the parameters that were used to define the potential, the middle section contains the length and energy scales,

[1] The mention of specific products, trademarks, or brand names is for purposes of identification only. Such mention is not to be interpreted in any way as an endorsement or certification of such products or brands by the National Institute of Standards and Technology. All trademarks mentioned herein belong to their respective owners.

Table 1
Parameters for model equation of Cesium. (1 amu is 1/12 of the mass of a carbon atom.)

C_6/k_B	43.91 K nm^6
ϵ/k_B	100 K
2μ	132.905429 amu
$\omega_\rho/2\pi$	1000 Hz
ω_z/ω_ρ	10
l	$3.90002358 \times 10^{-7}$ m
U	$3.31303438 \times 10^{-31}$ J
ϵ	4.16732862×10^{9}
σ	$1.77425759 \times 10^{-3}$
a	$-2.42519392 \times 10^{-2}$
x_0	$1.34268596 \times 10^{-2}$
N	30

and the bottom section contains the derived dimensionless quantities. k_B is Boltzmann's constant. The semi-infinite domain was truncated with $R_\rho = 5.0$ and $R_z = R_\rho/\sqrt{\omega_z/\omega_\rho} \approx 1.58$. The initial grid is a uniform 24×8 grid. We computed the five eigenvectors with eigenvalues closest to 0.0. The tolerance for the eigenproblem residual was $\tau_e = 0.01$ and the tolerance for the L^∞ discretization error estimate was $\tau_d = 0.05$. The linear system solver tolerance τ_{ls} is not relevant because we used a direct method. We defined a sequence of 8 pseudopotentials as described in Sec. 4.2. Table 2 shows the interleaving of the sequence of pseudopotentials and refinement of the grid. For this problem, less than 100 iterations with a Krylov subspace of dimension 15 was always sufficient for ARPACK to converge.

To obtain the desired accuracy, a grid with approximately 1.5 million vertices was required. The grid over the region where the atom-atom interaction potential dominates is shown in Fig. 7. Note the ring structure in the refined grid corresponding to the peaks and valleys of the waves. The computation took approximately 6 hours. The smallest grid elements were the result of 50 refinements of an initial grid element. Since each refinement doubles the number of elements, and there were 384 elements in the initial grid, a uniform grid with elements of that size would contain $384 \times 2^{50} \approx 4 \times 10^{17}$ elements. Clearly adaptive refinement is a requirement for a problem of this nature.

Figure 8 shows the computed wave function for the first trap state. The first part of the figure shows the entire domain, where the elongated large wave

Table 2
The sequence of problems solved.

σ	ϵ	x_0	N	WKB estimate of N	vertices	next action
0.301401	1.53e+02	0.41184	1	1.00	225	refine
0.301401	1.53e+02	0.41184	1	1.00	512	refine
0.301401	1.53e+02	0.41184	1	1.00	1046	next pseudopotential
0.144729	1.77e+03	0.25254	1	1.63	1042	adapt w/o increase
0.144729	1.77e+03	0.25254	1	1.63	1067	refine
0.144729	1.77e+03	0.25254	1	1.63	2173	next pseudopotential
0.069497	2.04e+04	0.15486	3	2.66	2182	adapt w/o increase
0.069497	2.04e+04	0.15486	3	2.66	2260	refine
0.069497	2.04e+04	0.15486	3	2.66	4673	next pseudopotential
0.033372	2.36e+05	0.09496	4	4.34	4687	adapt w/o increase
0.033372	2.36e+05	0.09496	4	4.34	4956	refine
0.033372	2.36e+05	0.09496	4	4.34	10107	next pseudopotential
0.016025	2.72e+06	0.05823	7	7.07	10154	adapt w/o increase
0.016025	2.72e+06	0.05823	7	7.07	10467	refine
0.016025	2.72e+06	0.05823	7	7.07	21232	refine
0.016025	2.72e+06	0.05823	7	7.07	43178	refine
0.016025	2.72e+06	0.05823	7	7.07	87175	next pseudopotential
0.007695	3.13e+07	0.03571	11	11.53	87241	adapt w/o increase
0.007695	3.13e+07	0.03571	11	11.53	89255	next pseudopotential
0.003695	3.61e+08	0.02190	19	18.81	91098	adapt w/o increase
0.003695	3.61e+08	0.02190	19	18.81	92152	refine
0.003695	3.61e+08	0.02190	19	18.81	185527	next pseudopotential
0.001774	4.17e+09	0.01343	30	30.67	186042	adapt w/o increase
0.001774	4.17e+09	0.01343	30	30.67	189293	refine
0.001774	4.17e+09	0.01343	30	30.67	383068	refine
0.001774	4.17e+09	0.01343	30	30.67	767587	refine
0.001774	4.17e+09	0.01343	30	30.67	1539217	done

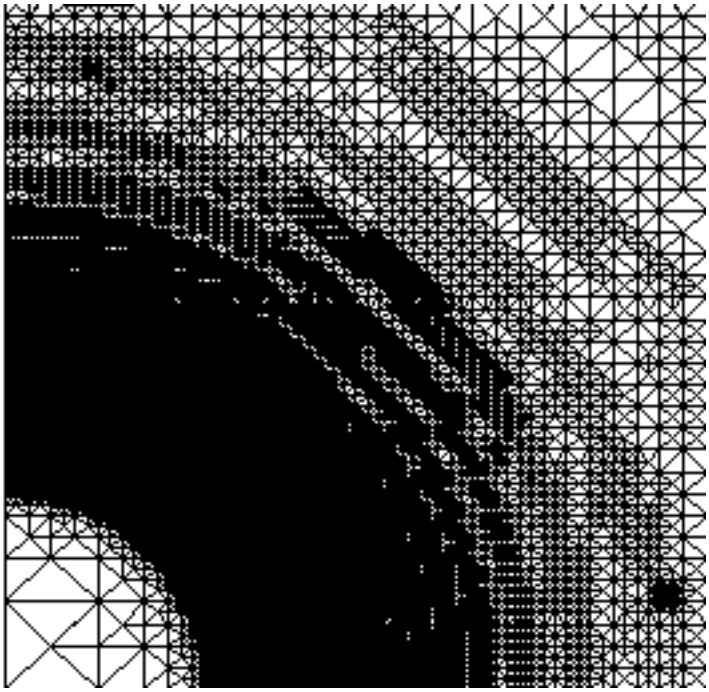

Fig. 7. The adaptively refined grid over the area where the atom-atom interaction potential dominates.

Table 3
Computed values for the first five positive eigenvalues.

vertices	1	2	3	4	5
383068	2.465	14.124	18.602	23.008	27.441
767587	12.583	16.718	20.943	25.217	29.633
1539217	12.226	16.246	20.268	24.295	28.375

due to the nonsymmetric trapping potential is clearly shown. By zooming into a region that covers the well of the interaction potential, we begin to see some features of the fine scale waves. A second zoom level is required to see the details of the 30 node wave created by the interaction potential. Figure 9 shows the computed wave function for the fifth trap state. Here we can see the 5 large scale waves. The small scale waves are not shown, but are nearly identical to those of Figure 8.

Table 3 shows the computed values for the eigenvalues for the last three grid refinements, all of which used the final potential. The eigenvalues exhibit a nearly constant separation, as expected from knowledge of the physical system. Given the difference between the eigenvalues on the grid with 767587 vertices and the grid with 1539217 vertices, the eigenvalues are probably accurate to at least 2 digits. This is reasonable given $\tau_d = 0.05$. We remark that, although $\tau_e = 0.01$, the actual eigenproblem residuals were on the order of 10^{-12} to 10^{-10}. Thus the accuracy of the eigenvalues is limited by the grid size, not the

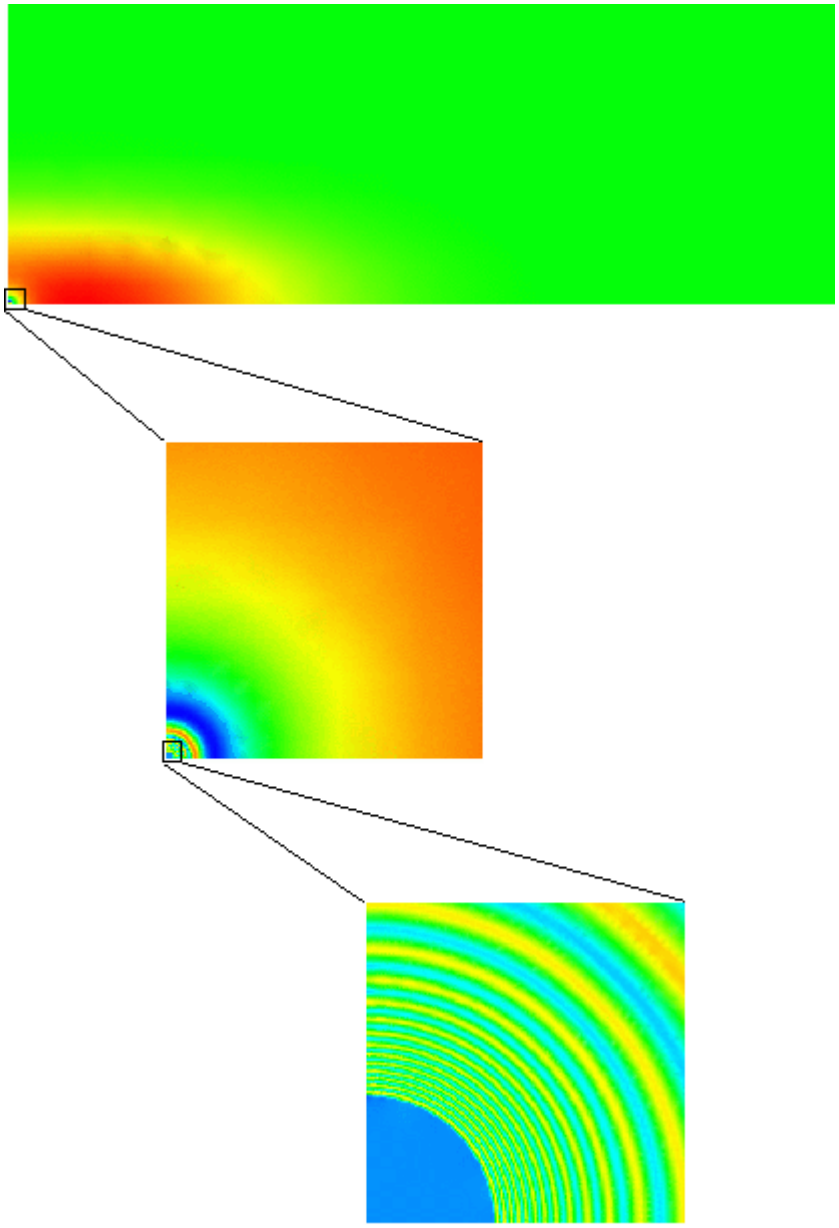

Fig. 8. The computed wave function for the first trap state, shown at three zoom levels.

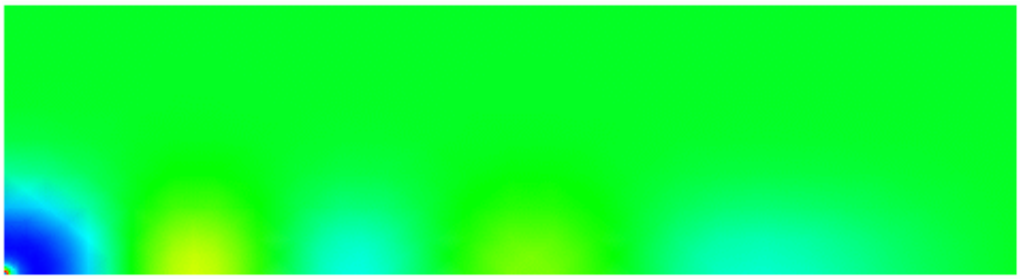

Fig. 9. The computed wave function for the fifth trap state.

eigensolver.

6 Conclusion

We have applied finite elements with adaptive grid refinement techniques to the solution of a model of two interacting Cesium atoms in an elongated optical trap. The model reduces to the solution of a two-dimensional Schrödinger equation with a potential that consists of a short range atom-atom interaction potential plus a long range trapping potential. Because the scales of these two parts of the potential differ by orders of magnitude, adaptive grid refinement is necessary to obtain a grid with small elements where the atom-atom interaction potential dominates and large elements where the trapping potential dominates.

An adaptive refinement strategy based on newest node bisection of triangles, combined with a carefully chosen sequence of pseudopotentials, was found to be very effective for this problem. The solution techniques easily found the areas where the grid should be finest. They accurately computed the wave functions containing two very different scales: 30 nodes with a wavelength 3 orders of magnitude smaller than the diameter of the initial triangles, and 1 to 5 nodes with a wavelength larger than the diameter of the initial triangles.

The work presented in this paper represents a first step toward an efficient solver for this class of problems, but there is still room for improvement in some aspects of the method. A very simple error indicator is used to guide the adaptive refinement. Although the resulting 1.5 million vertex grid is very effective for the model problem, it may be further reduced in size by using a more sophisticated error indicator. Additionally, a better error estimate may be useful as a stopping criterion. The definition of a good sequence of pseudopotentials has not yet been investigated, and may provide significant speedup by reducing the number of extraneous equations that are solved to draw the refinement into the correct region. Finally, the promise for the largest gain is in the linear system solver. Currently a direct method is used because the linear systems are highly indefinite and preconditioned Krylov space iterative methods converge much too slowly with the standard preconditioners. Future research will be directed toward the development of an effective preconditioner for this problem.

References

[1] P.R. Amestoy, I.S. Duff and J.-Y. L'Excellent, Multifrontal parallel distributed symmetric and unsymmetric solvers, Comput. Methods in Appl. Mech. Eng. **184** (2000) 501–520.

[2] P.R. Amestoy, I.S. Duff, J. Koster and J.-Y. L'Excellent, A fully asynchronous multifrontal solver using distributed dynamic scheduling, SIAM Journal of Matrix Analysis and Applications **23** (2001) 15–41.

[3] P.R. Amestoy, I.S. Duff, J. Koster and L'Excellent, J.-Y., MUMPS Homepage, http://www.enseeiht.fr/lima/apo/MUMPS/

[4] E.L. Bolda, E. Tiesinga and P.S. Julienne, Pseudopotential model of ultracold atomic collisions in quasi-one- and two-dimensional traps, Physical Review A **68** 032702 (2003).

[5] R. Lehoucq, D. Sorensen and C. Yang, *ARPACK Users' Guide*, SIAM, Philadelphia, 1998.

[6] R. Lehoucq, K. Maschhoff, D. Sorensen and C. Yang, ARPACK Homepage, http://www.caam.rice.edu/software/ARPACK/

[7] K.J. Maschhoff and D. Sorensen, P_ARPACK: An Efficient Portable Large Scale Eigenvalue Package for Distributed Memory Parallel Architectures, in: Jerzy Wasniewski, Jack Dongarra, Kaj Madsen and Dorte Olesen, (Eds.), *Applied Parallel Computing in Industrial Problems and Optimization*, Proceedings of The Third International Workshop on Applied Parallel Computing, PARA '96, Springer Lecture Notes in Computer Science, Vol. 1184, Springer-Verlag, Berlin, 1996, pp. 478–486.

[8] A. Messiah, *Quantum Mechanics*, John Wiley & Sons, New York, 1958.

[9] W.F. Mitchell, Unified Multilevel Adaptive Finite Element Methods for Elliptic Problems, Ph.D. thesis, Technical Report UIUCDCS-R-88-1436, Department of Computer Science, University of Illinois, Urbana, IL, 1988.

[10] W.F. Mitchell, Adaptive Refinement for Arbitrary Finite-element Spaces with Hierarchical Bases, J. Comp. Appl. Math. **36** (1991) 65–78.

[11] W.F. Mitchell, The Full Domain Partition Approach to Parallel Adaptive Refinement, in: *Grid Generation and Adaptive Algorithms*, IMA Volumes in Mathematics and its Applications Vol. 113, Springer-Verlag, 1998, pp. 151–162.

[12] W.F. Mitchell, The Design of a Parallel Adaptive Multi-Level Code in Fortran 90, Proceedings of the 2002 International Conference on Computational Science, 2002.

[13] W.F. Mitchell, PHAML Homepage, http://math.nist.gov/phaml

[14] M.A. Nielsen and I.L. Chuang, *Quantum Computation and Quantum Information*, Cambridge University Press, 2000.

[15] NIST Physics Laboratory, Quantum Information, http://qubit.nist.gov/index.html

[16] Y. Saad and M.H. Schultz, GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems, SIAM J. Sci. Comput. **7** (1986) 856–869.

[17] P.W. Shor, Algorithms for Quantum Computation: Discrete Logarithms and Factoring, in: S. Goldwasser, (Ed.), Proc. 35th Annual Symp. on Found. of Comput. Science, IEEE Comput. Soc. Press, 1994, pp. 124–134.

[18] K. Southwell, ed., Nature Insight on Ultracold Matter, Nature **416** (2002) 205–238.

[19] G. Strang and G.J. Fix, *An Analysis of the Finite Element Method*, Prentice-Hall, Englewood Cliffs, 1973.

[20] E. Tiesinga, S. Kotochigova and P.S. Julienne, Scattering length of the ground state Mg+Mg collision, Physical Review A **65** 042722 (2002).

[21] L.M.K. Vandersypen, M. Steffen, G. Breyta, C.S. Yannoni, M.H. Sherwood and I.L. Chuang, Experimental Realization of Shor's Quantum Factoring Algorithm Using Nuclear Magnetic Resonance, Nature **414** (2001) 883–886.

[22] R. Verfürth, *A Review of a posteriori Error Estimation and Adaptive Mesh-Refinement Techniques*, Wiley-Teubner, New York, 1996.