Perfect pattern formation of neutral atoms in an addressable optical lattice

J. Vala,[1,3] A. V. Thapliyal,[2,3] S. Myrgren,[1] U. Vazirani,[2,3] D. S. Weiss,[4] and K. B. Whaley[1,3]
[1]*Department of Chemistry and Pitzer Center for Theoretical Chemistry, University of California, Berkeley, California 94720, USA*
[2]*Department of Computer Science, University of California, Berkeley, California 94720, USA*
[3]*Mathematical Sciences Research Institute, 1000 Centennial Drive, Berkeley, California 94720-5070, USA*
[4]*Department of Physics, Pennsylvania State University, University Park, Pennsylvania 16802-6300, USA*
(Received 10 July 2003; revised manuscript received 15 June 2004; published 18 March 2005)

We propose a physical scheme for formation of an arbitrary pattern of neutral atoms in an addressable optical lattice. We focus specifically on the generation of a perfect optical lattice of simple orthorhombic structure with unit occupancy, as required for initialization of a neutral atom quantum computer. The scheme employs a compacting process that is accomplished by sequential application of two types of operations: a flip operator that changes the internal state of the atoms, and a shift operator that selectively moves the atoms in one internal state along the lattice principal axis. Realizations of these elementary operations and their physical limitations are analyzed. The complexity of the compacting scheme is analyzed and we show that this scales linearly with the number of lattice sites per row of the lattice.

DOI: 10.1103/PhysRevA.71.032324 PACS number(s): 03.67.Lx, 32.80.Pj, 42.50.Vk

I. INTRODUCTION

Neutral atoms trapped in an optical lattice constitute an attractive system for implementation of scalable quantum computation [1–3], simulation of many-body systems [4], and implementation of topological quantum computing [5–9]. In standard optical lattices, small lattice constants present a serious obstacle to implementing quantum computation, since it is difficult to address individual qubits with an external field. An optical lattice with a large lattice constant is in principle *addressable*, and can allow for the quantum state manipulation of individual atoms by an optical field. High addressability and controllability and low decoherence make addressable lattices promising candidates for large-scale quantum computer implementation.

The present work focuses on preparation of the initialization of an addressable optical lattice for the purposes of quantum computing. The objective is a perfectly filled, regular optical lattice, with each site occupied by a single atom in its motional ground state and in a specific internal state. We consider one-dimensional (1D), orthorhombic two-dimensional (2D), and three-dimensional (3D) lattices. After loading and laser-cooling atoms in the optical lattice, half the sites have a single atom and half are vacant. In order to use this system for scalable quantum computation, a perfect lattice with each site occupied by a single atom is required. We propose here an efficient, feasible scheme for *compacting* the optical lattice, i.e., for removing vacant sites to the edge of the lattice, thus creating a smaller lattice, but one more suitable for quantum computation.

The scheme presented here can as well be used to make arbitrary patterns of neutral atoms in an addressable optical lattice. These include lattices with fractional occupation, a specific translational and rotational lattice symmetry, a broken symmetry, and heteroatomic patterns. Another important property of the scheme is that it can be applied recursively to reach any desired accuracy of the pattern formation. After a large number of elementary operations, the lattice can be cooled and imaged again. The remaining defects can be eliminated by repeating the compacting scheme. This recursion increases the total pattern formation time by the total duration of additional cooling and imaging cycles, but does not result in any increase in the scaling of the pattern formation, i.e., the algorithmic complexity of the scheme is unchanged. The possibility of preparing any homoatomic or heteroatomic pattern of neutral atoms to an arbitrarily high degree of perfection makes this scheme attractive for initialization of quantum simulations of condensed phase systems, in addition to initialization of quantum computation.

Before proceeding, we briefly discuss (i) the possibility of quantum computation using an imperfect pattern of atoms with vacant sites, and (ii) other possible approaches to preparation of an optical lattice with single occupancy at each site. One can imagine starting with a known imperfect lattice pattern, and then instead of simplifying the distribution, devising a quantum algorithm that accounts for the known locations of the vacancies. Our analysis of this procedure suggests that even if the vacancy locations are known, they will cause bottlenecks in quantum information flow. These bottlenecks eventually occur when the computer size, i.e., the number of atomic qubits, or equivalently, the number of occupied lattice sites, is scaled up. In fact, we maintain that the probability of finding a "good" sublattice, where "good" means that each filled site is connected to another filled site, is exponentially small for any constant filling factor f. This is because the probability that there will be "insurmountable" blocks of defects (gaps) in any chosen sublattice increases rapidly with its size. The site percolation threshold has the following values for lattices of various dimensions: 1 (1D), 0.59 (2D), and 0.31 (3D) [10,11]. So an initial filling factor $f \approx 0.5$ does exceed the percolation threshold in a three-dimensional lattice. While this implies a nonzero probability that distant qubits are connected, it does not guarantee that they are connected via independent routes, nor even that they are connected at all. Perhaps more importantly, the mapping of a quantum algorithm onto an imperfectly filled lattice may be a hard classical computational problem. Based on these considerations, an imperfect lattice structure does not appear

practical for realization of a scalable quantum computer.

Three other approaches to the preparation of an optical lattice with single occupancy at each site have been theoretically explored. The quantum phase transition from the superfluid (SF) to Mott insulator (MI) phase [12,13], observed in [14], can prepare a singly occupied optical lattice with more than 90% fidelity when the lattice constant is small ($a \approx 0.5\ \mu\mathrm{m}, U_0 \approx 1\ \mu\mathrm{K}$) [13,14]. The lattice depth must be increased adiabatically to move from the SF to the MI phase. For the larger lattice constants required to make the lattice addressable ($a \approx 5.0\ \mu\mathrm{m}$), the tunneling strength diminishes exponentially, rendering the time scale of this adiabatic transfer unrealistically long. Another proposal for preparing unit occupancy in an optical lattice uses the dipole interaction between atoms in excited Rydberg states [3,16] to create an energy barrier to double site occupations [17]. A third alternative involves adiabatic loading of one optical lattice from another that has been preloaded with one or more atoms at every site by making use of the Mott-insulator quantum phase transition [18]. The adiabatic loading error is estimated to be 10^{-4}. Use of this approach with an *addressable* lattice, or starting from a smaller lattice which is then transformed to an adiabatic lattice, would relax the difficult adiabaticity requirement discussed above.

The alternative schemes mentioned above are designed to prepare an optical lattice where each site is occupied by a single atom. In contrast, the scheme proposed in this paper can also be used to design an arbitrary pattern of atoms and vacancies in an optical lattice, as well as patterns of different types of atoms. Another difference between the scheme proposed here and those based on the SF-MI transition is that, in the scheme proposed here, the result can be checked before a quantum calculation is begun, and any errors corrected.

The paper is structured as follows. After briefly characterizing the optical lattice system (Sec. II), we first explain the principles of the compacting scheme (Sec. III). Description of the compacting elements—site-selective *flip* of internal atomic states and state-selective *shift* operations—follows in Sec. IV, together with analysis of the corresponding heating. The flip operation we propose here can also be used to implement single-qubit operations, and provides an approach for doing this in an addressable optical lattice. We make a detailed analysis of the heating of atoms due to their interaction with an addressing laser beam during the flip operations and provide a systematic analysis of the motional heating of atoms in a general time-dependent optical lattice potential, to assess and minimize the extent of heating during shift operations. Section V provides a detailed analysis of the compacting algorithm and its corresponding complexity, demonstrating a favorable scaling of the compacting resources with the lattice size. Section VI summarizes and concludes. Technical details of both physical and algorithmic aspects are summarized in Appendixes.

II. ADDRESSABLE OPTICAL LATTICE

A 1D optical lattice can be realized by interfering two linearly polarized laser beams [19]. Higher-dimensional lattices with a simple orthorombic (nearly cubic) structure can be implemented using perpendicular arrangements of two (2D) or three (3D) pairs of laser beams with slightly different frequencies. The difference in frequencies for each pair eliminates undesired optical interference between them [20]. Such an optical lattice with $a \approx 5\ \mu\mathrm{m}$ can be made with CO_2-laser beams [19] or with blue-detuned light [21]. The blue-detuned standing waves consist of two beams propagating at a shallow angle θ_b with respect to each other, giving a lattice constant $a = \lambda/[2\sin(\theta_b/2)]$.

A pair of counterpropagating (along the z axis) linearly polarized laser beams of identical wavelength λ generates a standing wave characterized by the electric field

$$\mathbf{E}(z) = \sqrt{2}E_0[e^{i\theta/2}\cos(kz - \theta/2)\mathbf{e}_- - e^{-i\theta/2}\cos(kz + \theta/2)\mathbf{e}_+]. \tag{1}$$

Here θ is the relative angle between the linear polarization vectors of both beams, $k = 2\pi/\lambda$, and E_0 is the single-beam field amplitude. In the absence of any additional external field, the resulting 1D periodic lattice potential $U(z)$ (z is the propagation axis of the beams) depends on the magnetic hyperfine sublevel [1]. It can be characterized by the following relation:

$$U(z) = \frac{U_0}{2}\cos(\theta)\cos(2kz) + \frac{U_1}{2}\sin(\theta)\sin(2kz), \tag{2}$$

where $U_0 = \frac{2}{3}\tilde{\alpha}E_0^2$ and $U_1 = \frac{1}{3}\tilde{\alpha}E_0^2 m_F/F$ describe the well depths of the potential at $\theta = 0$ and π, respectively, $\tilde{\alpha}$ is the characteristic polarizability of a given transition, F is the total angular momentum of the relevant atomic hyperfine level, and m_F is the magnetic hyperfine sublevel. Note that the potentials for all magnetic hyperfine sublevels coincide for $\theta = 0$. If the linear polarization of one of the lattice beams is rotated, the periodic potential is shifted, to an extent that depends on the internal state of the trapped atom.

We focus here on an optical lattice filled with atoms of ^{133}Cs, although the proposed scheme is also applicable to other alkali-metal atoms. The $6s\ ^2S_{1/2}$ electronic ground state of ^{133}Cs consists of two hyperfine levels of total angular momentum $F = 3$ and 4, with energy splitting $\Delta E = E_{F=4} - E_{F=3} = 9.1926$ GHz. Any pair of the available $2F+1$ magnetic hyperfine sublevels of a single atom can in principle be used to define a qubit.

One-dimensional as well as multidimensional orthorhombic arrangements with approximately 20 sites per dimension and 100 μK depth are readily achievable. The lower-dimensional arrangements can be conveniently realized as sublattices of a 3D lattice, which has the benefit of removing the requirement for additional confining potentials when working with 1D and 2D lattices. After loading, cooling, and imaging the optical lattice (as described in Appendix A), each lattice site is either vacant or occupied with a single atom with equal probability. For 1D, 2D, and 3D lattice potentials and the occupation probability of 0.5, this results in 10, 200, and 4000 atoms, respectively, and hence in the same number of qubits. Imaging of the lattice carried out during the cooling process provides a map of the lattice occupancy in which the presence or absence of an atom at each site is specified. The vacant sites are then removed to the lattice

edge to create a smaller but perfect lattice, as described in the next section.

III. COMPACTING SCHEME

We now describe the scheme for compacting the imperfectly filled optical lattice into a smaller lattice in which each site is occupied by a single atom with 100% fidelity. As noted above, the scheme can also be applied to form any arbitrary desired lattice occupation pattern. Our procedure exploits the ability to move a subset of trapped atoms to fill vacant sites. Mathematically, this compacting of the lattice is equivalent to sequentially removing vacant sites to the lattice edges. We define two elementary physical operations that are sufficient to compact the lattice. These are (1) the flip operator $\hat{F}_I$, which toggles the internal state of an atom at position $I=(i,j,k)$ between two different (hyperfine) levels, and (2) the shift operator $\hat{S}_{IJ}$, which moves an atom from position I to a neighboring position $J=(i',j',k')$ with $\alpha'=\alpha\pm1$ for one of $\alpha=i,j,k$.

Compacting of the imperfectly filled lattice results from translating some of the atoms, which we refer to as *mobile* atoms, to fill vacant sites, while the rest, which we call *stationary* atoms, remain fixed within their lattice structure. Atoms are translated by rotating the linear polarization of one of the lattice beams by an angle θ relative to the counterpropagating beam [see Eq. (2)]. We refer to the state of a mobile atom as the *mobile* state, and that of a stationary atom as the *storage* state. The identity of these atoms can change during the compacting procedure. We do require, however, that all the participating atoms are trapped in the lattice potential at all times. This requirement translates into the condition that only internal atomic states with nonzero magnetic hyperfine quantum number m_F can be used in the compacting procedure, since $m_F=0$ states would become deconfined during the shift operation at $\theta=\pi/2$. Possible choices of suitable compacting states are $|\mathrm{Cs}\ 6S_{1/2},F=3,m_F=-1\rangle$ and $|\mathrm{Cs}\ 6S_{1/2},F=3,m_F=+1\rangle$, or $|\mathrm{Cs}\ 6S_{1/2},F=3,m_F=+1\rangle$ and $|\mathrm{Cs}\ 6S_{1/2},F=4,m_F=+1\rangle$. We demonstrate the scheme here for storage states $m_F=+1$ and mobile states $m_F=-1$.

The compacting scheme is illustrated in Fig. 1 for a short segment of a 1D lattice. Here atoms are compacted to the left. Note that in this simple example, the stationary atom is at the same location in space after the compacting is completed. This is not essential and is of course not implemented in the realistic case when an imperfect lattice is compacted to make a smaller perfect lattice in its center. We start with all atoms in the storage state $m_F=+1$, and use the occupancy map obtained from the lattice imaging to identify which atoms are to be made mobile and which are to remain stationary in the first stage of compacting. We then undertake the following steps.

(1) Change the relative polarization $\theta=0\rightarrow\pi$. All atoms are initially in the storage state $m_F=+1$ and all move by $-a/2$, where a is the lattice constant and the negative sign indicates direction to the left. Formally, this operation keeps the lattice occupation structure invariant and hence acts on this as identity.

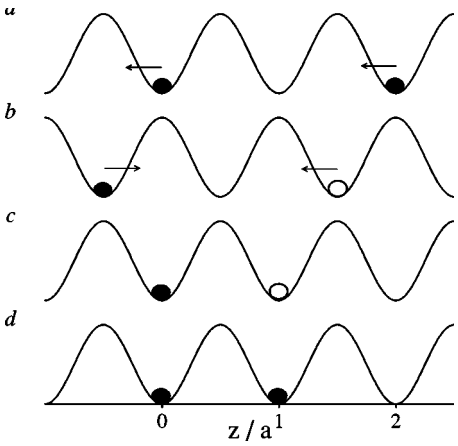

FIG. 1. Compacting scheme. (a) All atoms of the lattice are in a single internal state and are shifted by $-a/2$ as a result of rotation of the relative polarization angle θ from 0 to π. (b) The state of the mobile atoms is then flipped, e.g., from $m_F=+1$ (filled circles) to $m_F=-1$. This is accompanied by a change in shift direction, denoted by the arrows. (c) Rotation of the angle θ back from π to 0 moves the $m_F=+1$ atom back to its original position and the $m_F=-1$ atom forward to the next lattice site, providing the desired compacting of the lattice structure. (d) The mobile atoms are then flipped back to the original state. The net effect is a shift of the right-hand atom by one site to the left, or, equivalently, of the corresponding vacancy by one site to the right.

(2) At $\theta=\pi$, the mobile atoms are selectively flipped into the mobile internal state $m_F=-1$, while the states of stationary atoms remain unchanged.

(3) Change the relative polarization $\theta=\pi\rightarrow0$. The stationary atoms in $m_F=+1$ states are shifted back to their original positions (shift of $+a/2$) while the mobile atoms in $m_F=-1$ states are shifted forward to the next vacant lattice site (shift of $-a/2$). This step changes the lattice occupation map.

(4) The mobile atoms in $m_F=-1$ states are selectively flipped back to the storage state $m_F=+1$.

The procedure is repeated until all vacant sites have been moved to the edge of the lattice. The same methodology can clearly be applied to generate any partial occupation pattern of one or more atom types.

IV. ELEMENTARY COMPACTING OPERATIONS

A. Flip operation and site selectivity

The compacting scheme requires that we can make an atom at a single site undergo a flip transition, while none of its neighbors make the transition. We describe here our approach to accomplishing this in multidimensional optical lattices. It uses an independent "addressing" beam, which is a far-off-resonant, circularly polarized laser beam tightly focused on the atom to be flipped. The circularly polarized beam shifts the stationary and mobile atomic states differently, so that the resonance transition between them is shifted. The addressing beam will have nonzero intensity at nontarget atoms, especially those that lie along the addressing beam axis. However, as long as the Rayleigh range is

reasonably much smaller than the lattice spacing, the resonance frequency shift will be much smaller at nontarget sites. In the presence of the addressing beam, a pulse from a spatially homogeneous source can be used to flip the target atom. The requirement is that the pulse time be long compared to the inverse of the difference between the resonance frequency shifts of the target and the nearest nontarget atoms. In that case, nontarget atoms will be far enough off resonance that they will not be flipped. These flip pulses can be driven either by direct microwave excitation or by copropagating stimulated Raman beams. The parameters required for the addressing beam are easily attained. For instance, a 2 μW laser beam at 877 nm, focused to a waist of 1.2 μm, gives a 1 MHz relative Stark shift, 2.5 times larger than that of the nearest neighbor.

Another way to get site selectivity is to drive an off-resonant stimulated Raman transition using tightly focused perpendicular laser beams, so that only atoms at the target site experience appreciable intensity from both beams. An advantage to this approach is that, since there is no need to frequency-resolve the transitions at different sites, the flip can be accomplished much more quickly.

Site-selective flip operations can cause heating in several distinct ways. To our knowledge, these have not been analyzed in detail before. The first is due to the impulse that both target and nontarget atoms can feel during the process of turning on the spatially inhomogeneous beams, including either the addressing beam or the tightly focused stimulated Raman beams. However, when the frequency of these beams is tuned between the first-excited-state fine structure levels ($6P_{1/2}$ and $6P_{3/2}$ in Cs), the ac Stark shifts due to the two levels are opposite. For any individual ground state sublevel, there is a magic frequency where the two Stark shifts cancel. To avoid this heating, one simply needs to use the magic frequency for the storage hyperfine sublevel, for instance, 877 nm for the $F=3$, $m_F=1$ hyperfine sublevel. The difference in frequency between the Raman beams is negligible on the scale needed to avoid this heating effect.

A second heating mechanism only applies when an addressing beam is used, and it results because the addressing beam necessarily Stark-shifts the mobile sublevel. Therefore the trapping potential of an atom in that state is the sum of those due to the optical lattice light and the addressing light. The vibrational frequencies are thus different for the two hyperfine sublevels. Flip transitions between the storage and mobile sublevels can be made in two limits. If the pulse time is long compared to the inverse of the vibrational state splitting, so that the vibrational states are resolved, then the atom can make a transition to the new vibrational ground state, and there will be no heating. If the pulse time is shorter than the inverse of the vibrational state splitting, then the original vibrational wave function will be projected onto a superposition of states in the new basis. These levels will tend to dephase, and the atom will no longer be in the vibrational ground state when it is returned to the storage state. This heating is significantly reduced when the addressing beam is weakened, which requires that the microwave pulse be longer. Calculations of this heating effect have been performed in the context of making single-qubit operations in a site-addressable optical lattice, and the resulting heating amounts to $\sim 10^{-4}$ vibrational quanta per 30 μs flip operation [23].

Another heating mechanism is due to the photon recoil from the radiation used to implement the flip operations. In the case of microwave excitation or copropagating Raman beams, the photon recoil is negligible. But for the orthogonal Raman beams, unless the pulse is slow enough to resolve the atom's vibrational states, the target atom will receive a significant photon recoil kick. In the Lamb-Dicke limit, the probability of vibrational excitation per pulse is not high, but this would still likely be the dominant source of heating in the compacting sequence.

Driving the site-selective flip transition with π pulses can also be carried out if the pulse is shaped, e.g., like a Blackman pulse [24], in order to minimize off-resonant excitation. One could also use a square pulse, but this requires that all nontarget atoms lie close to the first minimum of the resulting sinc function [25]. In either realization, π pulses require very stable and repeatable field intensity. A flip operation can alternatively be implemented using adiabatic two-photon passage. We describe this option in Appendix B.

The site-selective flip operation described here can also be used to perform single-qubit operations on neutral atoms in a 3D optical lattice. In that case, it can be applied to atoms in the magnetic-field-insensitive $m_F=0$ states. The essence of the single-qubit gate is that an addressing laser beam shifts the hyperfine transition frequency in a spatially selective way, while microwave radiation actually drives the transitions. To understand the benefits of this approach, it is useful to consider alternatives. Doppler-free stimulated Raman transitions using tightly focused laser beams have been used to address individual ions [26], and could in principle be applied to neutral atoms. Such an approach, however, is problematic in three-dimensional lattices where other atoms are necessarily present in the path of a laser beam. This problem could be avoided by using two orthogonal Raman beams, but this introduces the possibility of unwanted vibrational excitation due to the Raman beams. With pulsed transitions, the size of the pulse is quite sensitive to the alignment and power of the focused laser beams. The optical-shift–microwave transition combination is far less sensitive to the intensity of the laser light at the atom. Site selectivity can also be obtained by using a magnetic (or electric) field gradient to create a position-dependent Zeeman (or Stark) shift. A drawback of this approach is that the field gradient inevitably introduces undesired phase shifts on other qubits in a register. In contrast, the operation presented here has only a minimal effect on qubits that are not specifically being addressed. In summary, the optical-microwave combination gate allows a single atom to be addressed in the middle of the 3D array, in a relatively alignment-insensitive way, while barely affecting the other atoms in the ensemble.

B. Shift operation and state selectivity

The shift operator moves the lattice trapped atoms over a distance $a/2$, where a is the lattice constant. All the atoms to be moved are initially prepared in a storage state, e.g., $m_F=+1$. It is easily seen from Eq. (2) that the atoms will move

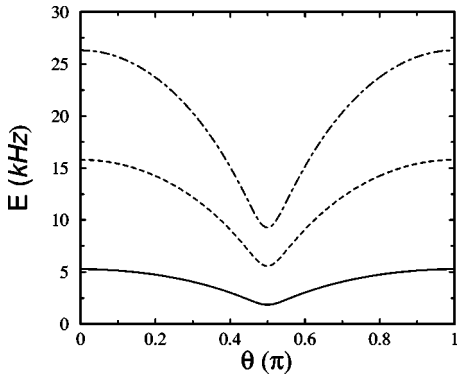

FIG. 2. Change of the lowest vibrational eigenstates of the periodic potential for the atomic state $|\mathrm{Cs}\ 6S_{1/2}, F=4, m_F=-1\rangle$ as a function of the rotation of the relative polarization angle θ. The energy is measured relative to the potential minimum. The vibrational frequency for the transition $0\rightarrow 1$ is 10.5 kHz for $\theta=0$ and 3.66 kHz for $\theta=\pi/2$.

when the linear polarization vector of one of the lattice beams is rotated. Rotating the relative angle of polarization θ by π moves all the atoms by $-a/2$. Switching the internal state of the mobile atoms to the mobile state, and slowly returning the polarization to its original value $\theta=0$, returns the stationary atoms back to their original positions, while each mobile atom is moved forward another $-a/2$, so that it is separated by a from its original location.

As seen from Eq. (2), at $\theta=0$ the atoms are confined only by the first term of the potential. As θ increases, the second term starts to dominate the optical lattice potential. At $\theta=\pi/2$, where the well depth is U_1, only this state-sensitive term confines the atoms. Further rotation to $\theta=\pi$ increases the well depth back to its original value U_0. We assume that the relative angle between the polarization vectors of the lattice beams, θ, is a linear function of time, i.e., $\theta=\beta t$.

The highest vibrational heating occurs for the largest variation of the potential during the shift operation. For ground state Cs atoms with $|m_F|=1$ in the blue-detuned optical lattice, the potential well depth changes by a factor of approximately $U_0/U_1=8$ during one shift operation. The corresponding change in the frequency of the lowest vibrational levels is illustrated in Fig. 2. Vibrational excitation of atoms during shift operations has been addressed experimentally by Mandel *et al.* [27] in a small-wavelength, nonaddressable optical lattice. The analysis was restricted to the simplified case when the well depth is independent of θ and provided an estimate of vibrational excitation to the $v=1$ state when the well depth decreases by a factor of 2 at $\theta=\pi/2$.

We have made a detailed study of vibrational heating deriving from shift operations in an addressable lattice, in order to determine the extent of vibrational heating over a broad range of optical lattice parameters. Initially each lattice atom is in its vibrational ground state, with energy U_0. A shift operation displaces the lattice potential and vibrationally excites atoms in the moving potential reference frame. This heating can be quantified by the total energy of a moving particle relative to the zero-point energy $\langle E\rangle - U_0$. We have employed both analytical and numerical techniques to get insight into the dynamics of atoms in such a time-varying lattice potential (characterized by $U_0=100\ \mu$K and $a=5.3\ \mu$m) and to establish upper bounds on the resulting vibrational heating of the trapped atom. We first studied a *simplified model* both analytically and numerically (see Appendix C), in order to demonstrate the analytical validity of our numerical tools. The numerical approach is based on a Fourier grid representation of quantum states and operators, and the Chebyshev polynomial method for quantum propagation. The methodology is described in detail in [28–30]. The explicit time dependence of the potential was approximated in discrete steps, with a time step $\Delta t \approx 1/U_0$ chosen to ensure the accuracy of the final results to within $\sim 10^{-5}$.

In the simplied situation when the lattice potential minimum is assumed to be independent of the polarization vector difference, the heating can be described as a transformation of the initial stationary state into a coherent state with maximal displacement determined by the energy transferred from the lattice to the atom (see Appendix C). The extent of heating can be measured in terms of the decrease in fidelity associated with the shift, $\tilde{F}=|\langle\Psi_i|\Psi_f\rangle|^2=|\langle v=0|\Psi_f\rangle|^2$, which corresponds to the projection of the final coherent state onto the vibrational ground state of the translated potential. We refer to the vibrational heating resulting from this transformation as "inertial." It is seen to give rise to a sigmoidal increase of $\tilde{F}$ as the pulse duration τ increases and can be suppressed by suitable choice of τ (see Appendix C). The *realistic case*, when the potential well depth is not constant with θ, has to be studied numerically. For the parameters $U_0=100\ \mu$K, $a=5.3\ \mu$m, the potential minimum decreases from U_0 to $U_0/8$ and increases again to U_0 as the angle between polarizations is rotated through π. The vibrational frequency is also modulated. Physically, the trapped atoms are first accelerated as θ approaches $\pi/2$ and are then decelerated when $\theta>\pi/2$. This process of acceleration and deceleration of the motion, absent in the simplified case, gives rise to a second contribution to vibrational heating that dramatically modifies the total amount of vibrational heating found for atoms trapped in the realistic potential.

This analysis shows that in the realistic case, the vibrational heating may be decomposed into two contributions, the inertial contribution and a second contribution deriving from the acceleration of the lattice potential reference frame. The inertial contribution can be essentially eliminated by suitable timing of the shift operation just as in the case of the simplified potential (Appendix C), although the precise time duration must now be determined by simulations. The combined effect of both heating mechanisms will depend on the ratio U_0/U_1, which in turn is determined by the lattice detuning and by the choice of atomic internal states. Comparison of Fig. 3 below with the corresponding behavior for the simplified potential (Appendix C, Fig. 11) shows that for small values of $U_0/U_1<3$ the inertial contribution to heating is dominant. The experimental estimates by Mandel *et al.* [27], obtained for $U_0/U_1\approx 2$, appear to fall into this category, although we note that these are made in a different parameter range. For higher ratios U_0/U_1, the potential acceleration dominates the heating mechanism, giving rise to a slower rise in $\tilde{F}$ as a function of τ.

Optimal conditions can be achieved by varying time duration τ, well depth U_0 (see Fig. 4) and the ratio U_0/U_1 (Fig.

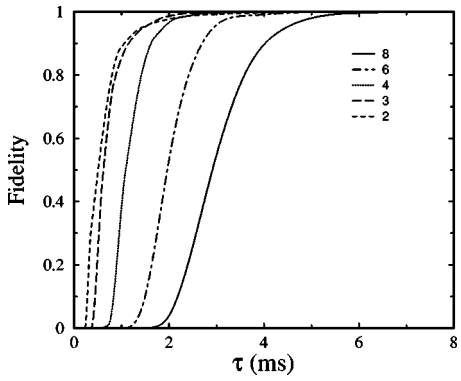

FIG. 3. Minimal fidelity $\tilde{F}$ as a function of the shift operation time τ, calculated numerically with the full potential at five different values of U_0/U_1. The minimal fidelity provides a measure of the vibrational heating, with small $\tilde{F}$ corresponding to greater heating. At small values U_0/U_1, the vibrational heating is dominated by the inertial contribution and the rise of $\tilde{F}$ with τ is very similar to that obtained with the simplified potential (see Fig. 11). As U_0/U_1 increases, the additional vibrational heating contribution from acceleration of the potential now increases, making the rise of $\tilde{F}$ slower. The inertial contribution has the same origin as that obtained with the simplified potential (Appendix C) and is similarly fully controllable by suitable timing of the shift operation. The plot shows that for given duration τ, a higher fidelity can be obtained with a larger value of the potential ratio U_0/U_1.

3) by suitable choice of internal atomic states. Figure 3 shows that for $U_0/U_1=8$, the atoms can be shifted on a time scale of 5–6 ms without appreciable heating (see also Fig. 4). It is evident that even for fairly fast shift operations, e.g., with $\tau<2$ ms, the atoms remain deeply trapped in the lattice, although recooling of the lattice may be required after many such operations.

V. EFFICIENCY OF COMPACTING

To systematically analyze the efficiency of the compacting scheme, we consider here finite 1D, 2D, and 3D lattices with 50% initial site occupation p_{occ} (i.e., the initial filling).

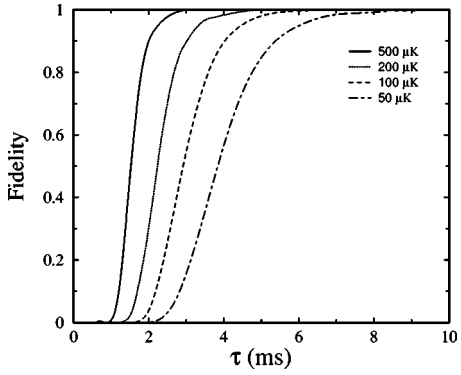

FIG. 4. Minimal fidelity $\tilde{F}$ during a shift operation for time duration τ, shown as a function of the well depth U_0. Calculations were carried out using the full (realistic) potential with $U_0/U_1=8$.

From imaging during laser cooling, we have a map of the lattice occupation. Our goal is to move the atoms so that they form a single contiguous block with no vacancies, i.e., a perfect finite lattice. The primitive operation in this analysis is the shift of an atom or group of atoms by one lattice site. We define the cost (time) for this operation to be one unit. The physical implementation of this operation was discussed in Sec. IV B, where we showed that moving a group of n atoms through one lattice site requires two movements of the trapping potential (two elementary shifts) and $2n$ flip operations (see Fig. 1). Note that there are only two elementary shifts because all atoms move in a shift. For fault-tolerant computation [31] we need to have $O(n)$ parallel operations [32], so for evaluating the scalability of the scheme we assume that the spin flips can be achieved in parallel. This is a good assumption, since the time required for a spin flip is much smaller (by a factor of 100–1000) than the time to move the trapping potential. In fact, we can ignore the spin flip cost altogether. We may also ignore the cost of classical computations required to plan and implement the compacting procedure since, with proper optimization, this cost scales at most linearly with the total number of lattice sites. Therefore, the total cost is essentially determined by the shift operations.

We first study the 1D lattice to show the solution in a simple setting. Then we consider the 2D lattice, which presents additional challenges. Finally, we generalize the techniques of the 2D lattice to the 3D lattice.

A. One-dimensional lattice

The one-dimensional lattice L_n has n sites. The site occupation probability is p_{occ}. Our goal is to move the atoms so that they form a line of atoms with no gaps in between.

Here, we use the simplest possible algorithm, which we call COMPACT, to remove the vacancies. The algorithm moves the atoms to the left so that all the vacancies move to the right. Thus after running the algorithm we end up with a line of atoms at the left of the lattice. The one-dimensional algorithm COMPACT consists of the following two steps.

(1) Find the leftmost vacancy v. If there are no atoms to the right of this vacancy, the compacting is complete and the algorithm terminates.

(2) Let G be the set of atoms that are to the right of v. Shift all the atoms in G left by one step. Return to the step 1 and repeat.

We say that a vacancy is at the right side if all the atoms are to its left. Since there are at most n vacancies and each operation can take one vacancy to the right side, the number of operations needed is n. On average there are $(1-p_{occ})n$ vacancies, so the expected number of shifts required is $(1-p_{occ})n$. Also, if there are exactly v vacancies that are not on the right side, then the number of operations needed will be v.

Clearly, the cost of this algorithm in the average case can be lowered by finding the center of mass of the set of atoms and compacting them around this. However, this modification does not improve the worst-case cost.

B. Two-dimensional lattice

We consider a finite two-dimensional square lattice $L_{n,n}$ with $N=n^2$ sites. We denote a sublattice $S_{i,j}$ $(1\leqslant i\leqslant j\leqslant n)$ as

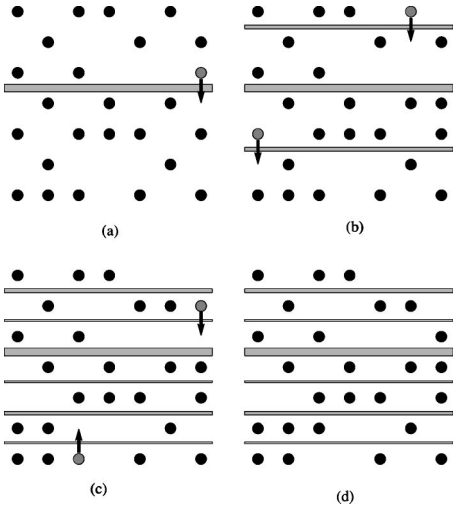

FIG. 5. Movement of atoms during the balancing process for a 7×7 lattice. (a) shows the first halving and balancing, corresponding to the zeroth level of the binary tree from Fig. 6. (b) shows the balancing corresponding to the second level of the tree. (c) shows the balancing corresponding to the third level of the tree. (d) is the final result where all the rows are balanced. The horizontal shaded rectangles show the partitioning of the lattice into halves. The gray circles represent atoms that are to be moved in the corresponding step. The black circles represent atoms that do not move in that step.

all rows between and including rows i and j. Let $N(S_{ij})$ denote the total number of atoms in the sublattice $S_{i,j}$ and n_k the number of atoms in the row k.

The two-dimensional compacting procedure has two parts.

(1) ^{2}BALANCE($S_{1,n}$). Here the atoms are first moved so that they are equally distributed among all the rows. The general idea is to divide the lattice into two halves, e.g., the top half and the bottom half (Fig. 5). Then, atoms are moved from the top half to the bottom half or vice versa, so that each half contains the same number of atoms (a difference of one atom is allowed when the total number of atoms is odd). After this, each half of the lattice is given as input to this procedure again, thus recursively balancing all the rows. Figure 5 illustrates this procedure for a 7×7 lattice. A more detailed description of the algorithm can be found in Appendix D. The recursive procedure can be summarized graphically by a binary tree structure (Fig. 6), described below, in order to find its cost.

(2) ^{2}ROW-COMPACT. The atoms in each row are then compacted as in the 1D case, except that here all rows can be compacted in parallel since mobile atoms can be moved together as a group. Figure 7 illustrates this procedure for a 7×7 lattice that has already been balanced according to Fig. 5.

A detailed list of the logical steps of ^{2}BALANCE($S_{1,n}$) is given in Appendix D. We now analyze the cost of the entire procedure. First we analyze the cost of ROW-COMPACT. It is the same as doing the 1D compacting on each row, but we can do all the rows in parallel, so the worst-case cost is n [33] (there being only n sites per row, thus requiring at most

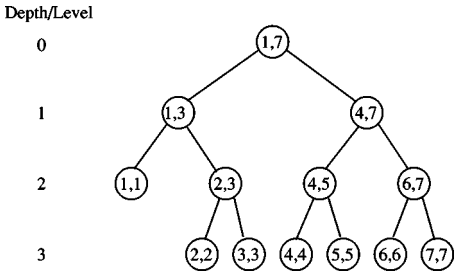

FIG. 6. Binary tree describing the recursion of BALANCE when it is given $S_{1,7}$, i.e., a 7×7 lattice, viz., the one considered in Fig. 5. The nodes of the tree show the starting and ending rows of the sublattices balanced in each step. The root of the tree has depth or level 0. The maximum depth is $\lceil \log_2 n \rceil$ [34] which is 3 in this case, since $n=7$. Note that at the last level, i.e., at the leaves of the tree, all the lattices consist of one row only. Hence no balancing needs to be done at this point. Thus the total number of balancing steps is $\lceil \log_2 n \rceil$.

$n-1$ moves to get an atom to the left). The average-case cost is again $(1-p_{occ})n$, since there are on average $(1-p_{occ})n$ vacancies per row. Additionally, if v is the maximum number of internal vacancies among all rows, then the cost is v, since we have to move all the vacancies beyond the last atom on the right.

The cost of the balancing procedure is somewhat more difficult to evaluate. We can represent the recursion by a binary tree of depth $\lceil \log_2 n \rceil$ [34] as shown in Fig. 6. The nodes of the tree are labeled by the first and last rows of the sublattice $S_{i,j}$ passed as input to BALANCE. The set of nodes at the same depth d are said to be at level d. The root is defined to be at depth 0. During each recursive step, the number of rows in the sublattice is halved, so the depth of the tree is $\lceil \log_2 n \rceil$. The total number of shifts required is the sum of the number of shifts required at each level of the tree.

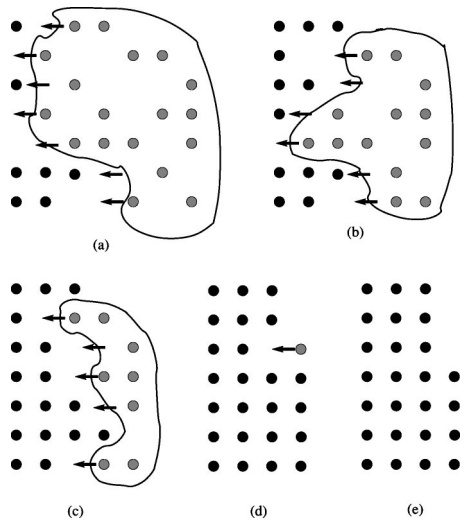

FIG. 7. Movement of atoms during the row-compacting process for the 7×7 balanced lattice obtained at the end of the balancing shown in Fig. 5. The gray circles represent atoms that are to be moved in the corresponding step. The black circles represent atoms that do not move in that step.

Note that the number of atoms to be moved during a balancing step does not cause a problem, since they can move in parallel as a group. Thus we basically need to count the number of shifts that the atoms have to make at each level. For a sublattice of k rows the atoms only need to make $\lceil k/2 \rceil$ shifts at most, since that is the largest number of rows of the two halves that are being balanced. At depth d the number of rows in a sublattice is at most $n/2^{d+1}$ [35]. There are 2^{d+1} such sublattices to be balanced at depth d. Now, the balancing of all these sublattices can be done in parallel. First move all the atoms in all these sublattices that need to be moved up. This takes at most $n/2^{d+1}$ shifts. Then shift all the atoms in all the sublattices that need to be moved down. This also takes at most $n/2^{d+1}$ shifts. Thus the total number of shifts required is $n/2+n/2+n/4+\cdots+(\lceil \log_2 n \rceil \text{terms}) \leq n/2 + (n/2)\Sigma_{n=0}^{\infty}(1/2)^n = 3n/2$, where the first term in the summation is only $n/2$ (not doubled) because there is only one direction in which to move the atoms the first time the lattice is halved. Also note that one would expect to be able to balance each level of the binary tree in one step on average. This would imply that the average complexity of balancing is $\approx \log_2 n$ shift operations. Numerical simulations suggest that this is indeed the case.

Putting together the two components of the cost analysis leads to the conclusion that in the two-dimensional lattice with $N=n^2$ sites, the compacting procedure takes at most $n+3/2n=5/2n=5/2\sqrt{N}$ steps [neglecting terms $O(\log N)$].

C. Three-dimensional lattice

We consider a three-dimensional cubic lattice $L_{n,n,n}$ with $N=n^3$ sites. The compacting procedure for this three-dimensional case is a generalization of that for the two-dimensional case. One key difference is that the notion of rows now needs to be defined. For example, the lattice is in the positive octant of the coordinate system, and its three sides coincide with the axes. Taking the lattice spacing to be the unit along the axes, the lattice sites are given by the coordinates (i,j,k) with $1 \leq i,j,k \leq n$, where the first coordinate is the x coordinate, the second coordinate is the y coordinate, and the third coordinate is the z coordinate. We define row $R_{i,j}$ to be the ordered set of lattice sites $\{(l,i,j) | 1 \leq l \leq n\}$. The lth position in the row $R_{i,j}$ corresponds to the site (l,i,j).

Define a sublattice $S_{i,j;k,l}$ with $(1 \leq i \leq j \leq n, 1 \leq k \leq l \leq n)$ to consist of the set of lattice sites $\{(a,b,c) | 1 \leq a \leq n, i \leq b \leq j; k \leq c \leq l\}$. Let $N(S_{i,j;k,l})$ denote the total number of atoms in the sublattice $S_{i,j;k,l}$.

The three-dimensional lattice compacting procedure consists also of two parts.

(1) ^{3}BALANCE$(S_{1,n;1,n}, 0)$. The atoms are moved so that they are equally distributed among the n^2 rows. The second parameter in the function call corresponds to the current recursion depth. The general idea is the same as in the case of the two-dimensional lattice. The difference is now that we alternately balance the two halves made by planes parallel to the xy plane and the xz plane, using the recursion depth d to implement this alternation of balancing.

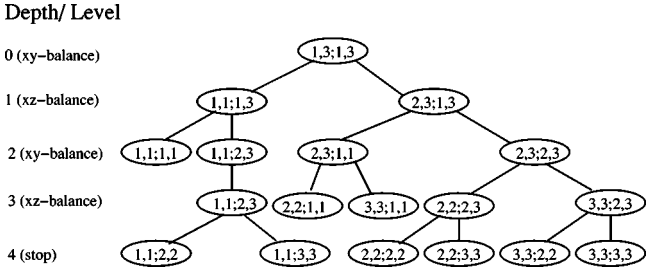

FIG. 8. Binary tree describing the recursion of BALANCE when it is given $S_{1,3;1,3}$ as input, i.e., $n=3$. The nodes of the tree show the indices of the sublattices balanced in each step. The root of the tree has depth or level 0. The maximum depth is $2\lceil \log_2 n \rceil$ [34] which is 4 in this case since $n=3$ to the even levels correspond to balancing with respect to the first set of indices (the halving plane is parallel to the xy plane), while the odd levels correspond to balancing with respect to the second set of indices (the halving plane is parallel to the xz plane). Note that at the last level, i.e., at the leaves of the tree, all the lattices consist of one row only and hence no balancing needs to be done. Thus the number of nontrivial balancing steps is $2\lceil \log_2 n \rceil$.

(2) ^{3}ROW-COMPACT. The atoms in each row are then compacted in a fashion similar to the 1D case, except now the operations for all the rows can be done in parallel, since we can move a three-dimensional group of atoms together in a single shift.

A detailed list of the logical steps for ^{3}BALANCE$(S_{1,n;1,n}, 0)$ is given in Appendix E. Now we analyze the cost of the entire procedure. First we analyze ROW-COMPACT. The analysis is identical to that for the two-dimensional case, namely, a worst-case cost of n and average-case cost of $(1-p_{occ})n$. Additionally, if v is the largest number of internal vacancies in any row, then the cost is v.

The analysis of balancing for the three-dimensional lattice is similar to that of the two-dimensional lattice. During each recursion step the lattice is halved along the xy or xz plane, so the depth of the recursion is $2\lceil \log_2 n \rceil$ [34]. We can represent the recursion by a binary tree of depth $2\lceil \log_2 n \rceil$, where even levels correspond to halving the sublattice parallel to the xy plane and odd levels correspond to halving the sublattice parallel to the xz plane (see Fig. 8). The sum of the number of shifts required for each level of the tree gives us the total number of shifts required during the execution of the algorithm. We will do the counting separately for the odd and even levels. Since mobile atoms can be moved in parallel, we only need to count the number of shifts that the atoms have to make during each balancing step. For a sublattice of k rows along the direction in which we halve the sublattice atoms, we only need to make $\lceil k/2 \rceil$ shifts. At depth $2d$ and $2d+1$ the relevant number of rows in the sublattice is $n/2^{d+1}$ (taking the root to be at depth 0) [36]. As in the 2D case, we parallelize the required shifts for all the sublattices at a given depth k, by doing all the shifts in the same direction in parallel. Since there are two directions, the number of shifts required at each depth is twice that required for a single sublattice at that depth. Then the total number of shifts at the even levels is $n/2+n/2+n/4+n/8\cdots(\lceil \log_2 n \rceil \text{terms}) \leq n/2$

$+(n/2)\Sigma_{n=0}^{\infty}(1/2)^n = 3/2n$. The first term in the sum is only $n/2$ (not doubled) because there is only one direction in which to move the atoms the first time the lattice is halved. For the odd levels a similar argument holds, except that the first term is n because there are two sublattices at that stage and balancing them may require moving atoms in opposite directions. Then the total number of moves required at the odd levels is $n+n/2+n/4+n/8\cdots(\lceil\log_2 n\rceil \text{terms}) \leqslant n\Sigma_{n=0}^{\infty}(1/2)^n = 2n$. Adding together the total moves at the even and odd levels gives us the total number of moves required: $(3/2)n+2n=(7/2)n$. Thus, in the three-dimensional lattice with $N=n^3$ sites, the lattice can be compacted in at most $(7/2)n+n=(9/2)n=(9/2)\sqrt[3]{N}$ steps [again neglecting terms $O(\log N)$]. As in the case of the two-dimensional lattice, one would expect that the balancing at each level of the binary tree can be typically done in one step. This would imply that the average cost of balancing is $2\lceil\log_2\rceil n$ steps. However, we do not have a rigorous proof of this estimate.

VI. SUMMARY AND CONCLUSIONS

We have presented an experimentally viable scheme for initializing a quantum computer based on neutral atoms trapped in an optical lattice of large lattice constant. The proposed compacting scheme allows preparation of a uniformly filled optical lattice of orthorhombic structure, where each lattice site is occupied with a single atom in a specific internal level and motional ground state. We have proposed physical mechanisms for the realization of two elementary compacting operations, namely, the flip of the internal state of atoms and their shift within the lattice structure. We have analyzed the efficiency of their implementation in some detail, with particular attention devoted to the study of motional heating during the elementary operations. Mechanisms to control this heating were proposed, based on analytic and numerical solutions of the atomic dynamics in the time-dependent lattice potential. Our analysis of the complexity of the compacting process demonstrates its scalability with the size of the quantum computer in one, two, and three dimensions.

The scheme presented here is feasible with current technology. Our detailed study of physical feasibility together with algorithmic scaling shows that we can achieve a single cycle of the entire compacting scheme in less than one second, starting from a half-filled lattice of approximately 8000 sites. If the elementary steps are implemented perfectly, then this guarantees a perfect pattern formation in a single cycle. If this is not the case, the scheme may be applied recursively to achieve any desired accuracy. After a complete compacting cycle, the lattice can be optically cooled and imaged to identify any remaining errors. Thus even if implementation of the elementary steps is imperfect or if decoherence introduces errors during compacting, these accumulated errors are reset to zero. The most serious limitation derives from atom losses due to scattering of lattice photons or to inelastic collisions with stray particles. Fortunately, the rate of these processes is expected to lie below a threshold value of 10^{-4} in terms of the typical elementary operation time scale. This is consistent with fault-tolerance requirements for quantum computation [15]. For example, assuming the error rate of atom loss equal to 10^{-4} per shift operation time scale, compacting a half-occupied 3D lattice with 20 sites per dimension can reach a fidelity above 0.99 in less than one second within the first cycle. A second cycle improves the result to fidelity $\sim$0.998 in less time, since fewer operations are required to implement this cycle. Recursive pattern formation can thus be seen as a process that is correcting these (classical) errors by fixing defects at a speed higher than they occur, thereby removing any propagation of these errors, as required for fault tolerance.

In summary, the scheme presented here allows the preparation of an arbitrary homoatomic or heteroatomic pattern of neutral atoms. State preparation can be accomplished in a fault-tolerant manner, and the prepared state can be checked before it is used. We expect that this approach will find immediate application in the preparation of quantum states in site-addressable large-wavelength optical lattices.

ACKNOWLEDGMENTS

We thank Ronald de Wolf and Kenneth Brown for helpful discussions. The effort of the authors is sponsored by the Defense Advanced Research Projects Agency (DARPA) and the Air Force Laboratory, Air Force Material Command, USAF, under Agreement No. F30602-01-2-0524. D.S.W. also acknowledges support from the National Science Foundation.

APPENDIX A: LOADING AND COOLING OF OPTICAL LATTICE

A magneto-optical trap [37] can be used to load well-spaced far-off-resonant optical lattice sites with many atoms ($N=10-100$) [20]. Atoms are lost in pairs during laser cooling, as a result of photon-assisted collisions [38,39]. Eventually, the sites initially occupied by an odd number of atoms become singly occupied, while sites initially occupied by an even number of atoms are left empty. The filling factor that results from this combined loading and cooling process approaches a maximum of 0.5. Such loading of a far-off-resonant optical lattice with Cs atoms has been demonstrated experimentally for a lattice of a small lattice constant ($\sim$0.5 μm), achieving a filling factor of $\sim$0.44 [38].

Trapped atoms can then be brought to the vibrational ground state of their lattice sites using Raman sideband cooling [40–44]. This procedure has been experimentally demonstrated in 1D [42,44], 2D [41], and 3D [40,43] optical lattices. In [43], the 3D ground state was populated by up to 55% of atoms, corresponding to over 80% in each dimension. It was shown that the cooling was limited by the rescattering of cooling photons. This mechanism is dramatically reduced at the low densities associated with large lattice constants. Even with a 5 μm lattice constant, a 150 μK depth lattice leaves the atoms well within the Lamb-Dicke limit required for efficient sideband cooling. It is reasonable to expect nearly 100% vibrational ground state population after Raman sideband cooling.

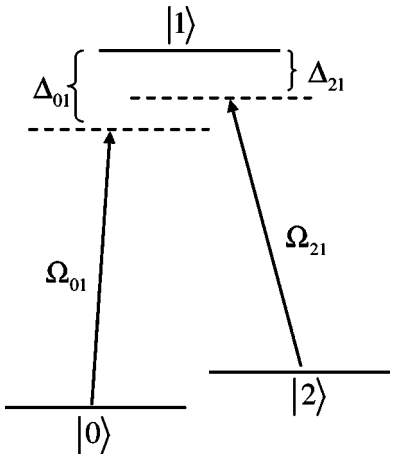

FIG. 9. The three-level quantum system coupled by two optical fields in a Λ configuration. Ω_{ij} is the Rabi frequency for coupling of the states i and j; Δ_{ij} is the detuning from a resonance $\omega_{ij}=(E_i-E_j)/\hbar$.

With a 150 μK deep, 5 μm spaced lattice, the vibrational level spacing is well below the expected minimum polarization gradient cooling temperature of 2 μK, so the polarization gradient cooling limit should be obtained [20]. With a temperature so small compared to the lattice depth, there will be negligible site hopping during cooling. Therefore the atoms can be imaged during the cooling for as long as needed. 1D CO_2 lattices have recently been imaged [19]. In a similar way, it should be possible to use $\sim$0.6 numerical aperture optics to image successive 2D planes, and thereby construct the site occupancy map of a 3D lattice. Since the localized atoms in the 3D lattice will scatter light coherently, there will be significant interference between light scattered from atoms in the image plane and light from nonimaged planes. Information can be obtained from such signals, particularly if the ratio of the lattice spacing to the imaging wavelength is well chosen. A detailed study of the imaging in this situation shows that with $15\times15\times15$ lattices, one can determine the site occupancy of each site in an arbitrarily half-filled lattice with greater than 99% reliability [22]. Once the lattice occupation map is known, the compaction procedure can begin (Secs. III and IV). After the first compaction, one would expect the sites to be 99% filled. From this more regular arrangement of atoms, the site occupancy can be determined with negligible error. When compaction is repeated, imperfect knowledge of the vacancy locations ceases to be a source of error.

APPENDIX B: ADIABATIC FLIP OPERATION

The flip operation can also be implemented using adiabatic two-photon passage. Consider a three-level optical system in a Λ configuration (see Fig. 9) where two ground states $|0\rangle$ and $|2\rangle$ are coupled indirectly via an intermediate state $|1\rangle$, using a pair of laser fields. Adiabatic passage occurs when the initial and target levels coupled off-resonantly through the intermediate state sweep through their dressed-state resonance as a result of interaction with a linearly frequency-chirped field (for a description of the chirped field and its interaction with matter see for instance [45,46]).

The atom-field Hamiltonian in the interaction representation is

$$\hat{\mathbf{H}}(t)=\hat{\mathbf{H}}_0+\hat{\mathbf{H}}_I(t)=\sum_{i=0}^{2}\hbar\omega_i|i\rangle\langle i|+[W_{01}(t)\sigma_{01}^{+}+\text{H.c.}]+[W_{12}(t)\sigma_{12}^{+}+\text{H.c.}] \tag{B1}$$

where $\sigma_{ij}^{+}=|j\rangle\langle i|$ and $W_{ij}(t)$ is the time-dependent coupling strength given by the external optical field. Canonical transformation of this system with coupling fields detuned from resonance with the intermediate state results in an effective two-level system [47,48]:

$$\hat{\mathbf{H}}_{\text{eff}}(t)=\hat{\mathbf{H}}_0+\hat{\mathbf{H}}_{\text{Stark}}+\hat{\mathbf{H}}_{I,\text{eff}}(t)=\sum_{i=0}^{2}\hbar\omega_i|i\rangle\langle i|+\frac{|W_{01}(t)|^2}{\Delta_{01}}\sigma_{01}^{z}+\frac{|W_{12}(t)|^2}{\Delta_{12}}\sigma_{12}^{z}+W_{02}\sigma_{02}^{+}+W_{02}^{*}\sigma_{02}^{-} \tag{B2}$$

where $W_{02}=(W_{01}W_{12}/2)(1/\Delta_{12}-1/\Delta_{01})$ is the effective coupling strength between the states $|0\rangle$ and $|2\rangle$, and Δ_{ij} is the detuning of the field from the atomic transition between levels i and j. Application of a linearly chirped field for one of the transitions $0\rightarrow1$ or $2\rightarrow1$ results in a robust and complete flip operation.

APPENDIX C: VIBRATIONAL HEATING IN SHIFT OPERATION

This appendix complements Sec. IV B where the vibrational excitation (heating) of atoms during the shift operation is investigated in the most general situation characterized by a strong variation of the lattice vibrational spectrum with time. This investigation is carried out using numerical simulation of the dynamics of atoms in a time-varying potential, and requires benchmark calculations that demonstrate the accuracy and reliability of the simulation results. These are provided here, where we focus on the simplified case when the lattice potential well depth is constant with variation of θ. In the vicinity of the minimum, we can approximate the periodic potential of the lattice with a second-order Taylor expansion, yielding the harmonic vibrational frequency $\omega=k\sqrt{2U_0/m}$, where $k=2\pi/\lambda$. Comparison with the results of direct diagonalization of the periodic potential shows that the fractional difference from the true frequency is $\sim10^{-3}$, and that the anharmonicity of the periodic potential [manifested in the deviation of $\omega_{k,k+1}=(E_{v=k+1}-E_{v=k})/\hbar$ from the harmonic frequency] is linear in the vibrational quantum number k and is only 3% for $k=20$. These facts justify the use of the harmonic approximation to get analytical insights into the process of vibrational heating.

The motion of the simplified potential, induced by rotation of the polarization vector of one of the lattice beams, is linear in time and induces a transfer of energy $\Delta E=mv^2=m\Delta z^2/\tau^2$ into the system. Here m is the atom mass, and Δz is the total distance of the potential translation over the du-

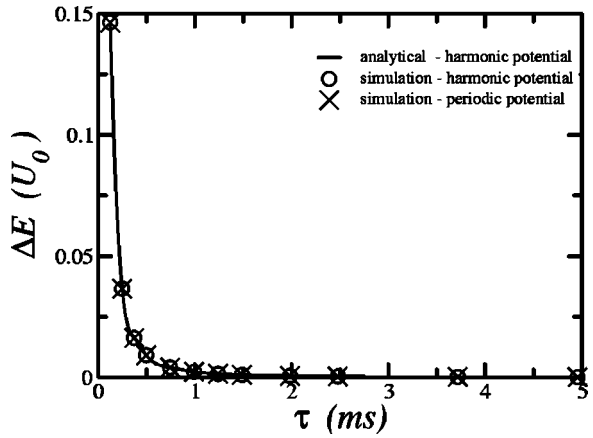

FIG. 10. Maximal vibrational heating of the initial ground state population ΔE as a function of the total time scale τ of the shift operation. The analytical and numerical simulation results in the harmonic approximation are compared with each other and with the numerical simulation using the simplified periodic potential (see the text for explanation). The analytic and numerical results are indistinguishable on this scale.

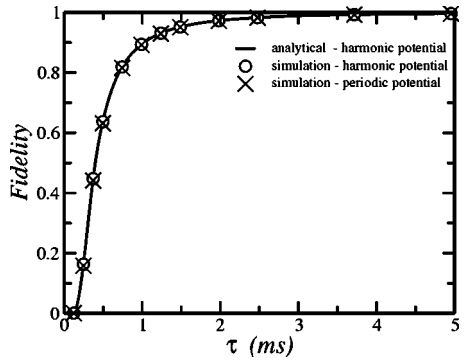

FIG. 11. Minimal fidelity $\tilde{F}$ as a function of the shift operation time scale for the simplified case (see the text for explanation). The excellent agreement between analytic and numerical results shows that the classical picture of the process based on the coherent state representation correctly reproduces the atomic dynamics induced by the time-varying potential. For short time scales, the anharmonicity of the periodic potential slightly reduces the fidelity compared to the results obtained within the harmonic approximation.

ration τ. ΔE is the maximum energy that can be transferred into vibrational motion during the beginning or end of a translation. Onset of the potential translation at time $t=0$ causes displacement of the initially stationary vibrational state with respect to the moving potential reference frame. The displacement transforms the stationary initial state into a vibrating coherent state with a maximal displacement of x_{max}. This displacement is related to the transferred energy via $\Delta E=\frac{1}{2}m\omega^2x_{max}^2$, whence $x_{max}=(\sqrt{2}/\omega)(\Delta z/\tau)$. The maximal energy in the system at $\theta=\pi$ is $2E$. The analytical and numerical results for the harmonic and periodic potential show perfect agreement, demonstrated in Fig. 10, and hence justify the applied harmonic approximation and the coherent state representation.

To further characterize the quality of the shift operation, we define a fidelity measure $\tilde{F}=|\langle\Psi_i|\Psi_f\rangle|^2=|\langle v=0|\Psi_f\rangle|^2$, which corresponds to the projection of the final coherent state onto the vibrational ground state of the translated potential. In this simplified case, the fidelity can be evaluated analytically, yielding the expression $\tilde{F}=\exp[-(m\omega/4\hbar)x_{max}^2]=\exp[-(m/2\omega\hbar)\Delta z^2/\tau^2]$. As can be seen in Fig. 11, the fidelity remains low for shift operation times less than 0.5 ms, whereafter it rises sharply. For shift times longer than 2 ms the fidelity is very close to unity, and thus provides a good estimate of the time scale necessary to preserve adiabaticity.

Exploration of this simplified model provides insight into the control and minimization of vibrational heating. The onset of the potential motion generates a coherent state from the initial vibrational ground state. In the moving reference frame, the potential motion first displaces the initial state to the negative momentum and negative coordinate region of the phase space, while setting its motion along the classical phase space trajectory corresponding to the coherent state. When $\theta=\beta\tau=\pi$ the actual motional phase of the coherent state accumulated during the evolution ($e^{-i\omega\tau}$) determines whether halting the potential at this time will reduce or increase the total vibrational energy of a trapped atom. Consequently, choosing the duration of the shift operation τ to correspond to an integer multiple of the vibrational period of the coherent state $T=2\pi/\omega$ can result in complete elimination of the vibrational heating within the simplified model.

APPENDIX D: 2D BALANCE

^{2}BALANCE($S_{i,m}$).; see Sec. V for definitions.

Input. $S_{i,j}$.

Goal. Balance the rows in $S_{i,j}$ so that $\forall k,l$ we have $|n_k-n_l|\leq 1$, where $i\leq k,l\leq j$. Equivalently, each row has at least $n_{min}=\lfloor N(S_{i,j})/(j-i+1)\rfloor$ atoms.

Algorithm

(1) Let $l=(j-i+1)$ denote the number of rows in $S_{i,j}$. If $l=1$ RETURN.

(2) Let $m=i+\lfloor l/2\rfloor$ be a middle row of $S_{i,j}$. For balancing, we need to have $n_{req}=(m-i+1)n_{min}$ atoms in the sublattice $S_{i,m}$.

(3) If $[N(S_{i,m})>n_{req}]$, shift $[N(S_{i,m})-n_{req}]$ atoms down from $S_{i,m}$ to $S_{m+1,j}$.

(4) If $[N(S_{i,m})<n_{req}]$, shift $[n_{req}-N(S_{i,m})]$ atoms up from $S_{m+1,j}$ to $S_{i,m}$.

(5) ^{2}BALANCE($S_{m+1,j}$).

APPENDIX E: 3D BALANCE

^{3}BALANCE($S_{i,m;k,l},d+1$); see Sec. V for definitions.

Input. Sublattice $S_{i,j;k,l}$ and recursion depth d.

Goal. Balance the rows in $S_{i,j;k,l}$ so that $\forall_{a,b}$ such that $i\leq a,b\leq j$ and $\forall_{c,d}$ such that $k\leq c,d\leq l$ we have $|n_{a,c}-n_{b,d}|\leq 1$, where $n_{a,c}$ represents the number of atoms in the row

$R_{a,c}$. Equivalently, each row should have at least $n_{min} = \lfloor N(S_{i,j;k,l})/(j-i+1)(l-k+1) \rfloor$ atoms.

Algorithm

(1) If $i=j$ and $k=l$ then RETURN. Only one row so no balancing needed.

(2) If d is odd, go to 10; else if $i=j$ ^{3}BALANCE$(S_{i,j;k,l}, d+1)$. RETURN.

(3) If sublattice $S_{i,j;k,l}$ contains only one xy plane (i.e., $i=j$), then go to 10.

(4) Let $l=j-i+1$ be the number of planes parallel to the xy plane. Then $m=i+\lfloor l/2 \rfloor$ defines a middle plane parallel to the xy plane.

(5) Let $n_{req}=n_{min}(m-i+1)(l-k+1)$ be the minimum number of atoms that are required to be in $S_{i,m;k,l}$.

(6) If $[N(S_{i,m;k,l}) > n_{req}]$, shift $[N(S_{i,m;k,l})-n_{req}]$ atoms from $S_{i,m;k,l}$ to $S_{m+1,j;k,l}$.

(7) If $[N(S_{i,m;k,l}) < n_{req}]$, shift $n_{req}-N(S_{i,m;k,l})$ atoms from $S_{m+1,j;k,l}$ to $S_{i,m;k,l}$.

(8) ^{3}BALANCE$(S_{m+1,j;k,l}, d+1)$.

(9) RETURN.

(10) If $k=l$ ^{3}BALANCE$(S_{i,j;k,l}, d+1)$. RETURN.

(11) Do steps 4 to 8, replacing indices (i,j) with indices (k,l).

(12) RETURN.

[1] I. H. Deutsch and P. J. Jessen, Phys. Rev. A **57**, 1972 (1998).

[2] D. Jaksch, H.-J. Briegel, J. I. Cirac, C. W. Gardiner, and P. Zoller, Phys. Rev. Lett. **82**, 1975 (1999).

[3] D. Jaksch, J. I. Cirac, P. Zoller, S. L. Rolston, R. Côté, and M. D. Lukin, Phys. Rev. Lett. **85**, 2208 (2000).

[4] R. P. Feynman, Int. J. Theor. Phys. **21**, 467 (1982).

[5] M. Freedman, C. Nayak, K. Shtengel, K. Walker, and Z. Wang, e-print quant-ph/0307511.

[6] M. Freedman, C. Nayak, and K. Shtengel, Phys. Rev. Lett. **94**, 066401 (2005).

[7] A. Yu. Kitaev, Ann. Phys. (N.Y.) **303**, 2 (2003).

[8] E. Dennis, A. Kitaev, A. Landahl, and J. Preskill, J. Math. Phys. **43**, 4452 (2002).

[9] L.-M. Duan, E. Demler, and M. D. Lukin, Phys. Rev. Lett. **91**, 090402 (2003).

[10] B. D. Hughes, *Random Walks and Random Environments* (Oxford University Press, New York, 1995).

[11] G. Grimmet, *Percolation* (Springer, Berlin, 1999).

[12] M. P. A. Fisher, P. B. Weichman, G. Grinstein, and D. S. Fisher, Phys. Rev. B **40**, 546 (1989).

[13] D. van Oosten, P. van der Straten, and H. T. C. Stoof, Phys. Rev. A **63**, 053601 (2001).

[14] M. Greiner, O. Mandel, T. Esslinger, T. W. Hänsch, and I. Bloch, Nature (London) **415**, 39 (2002).

[15] J. Preskill, Proc. R. Soc. London, Ser. A **454**, 385 (1998).

[16] M. D. Lukin, M. Fleschhauer, R. Côté, L. M. Duan, D. Jaksch, J. I. Cirac, and P. Zoller, Phys. Rev. Lett. **87**, 037901 (2001).

[17] M. Saffman and T. G. Walker, Phys. Rev. A **66**, 065403 (2002).

[18] P. Rabl, A. Daley, P. Fedichev, J. Cirac, and P. Zoller, Phys. Rev. Lett. **91**, 110403 (2003).

[19] R. Scheunemann, F. S. Cataliotti, T. W. Hansch, and M. Weitz, Phys. Rev. A **62**, 051801 (2000).

[20] S. L. Winoto, M. T. DePue, N. E. Bramall, and D. Weiss, Phys. Rev. A **59**, R19 (1999).

[21] D. Weiss, J. Vala, A. Thapliyal, S. Myrgren, U. Vazirani, and K. Whaley, Phys. Rev. A **70**, 040302 (2004).

[22] B. Vaishnav and D. Weiss (unpublished).

[23] B. Vaishnav (unpublished).

[24] S. A. Rice and M. Zhao, *Optical Control of Molecular Dynamics* (Wiley Interscience, New York, 2000).

[25] M. Kasevich and S. Chu, Phys. Rev. Lett. **69**, 1741 (1992).

[26] D. J. Wineland, C. Monroe, W. M. Itano, D. Leibfried, B. E. King, and D. M. Meekhof, J. Res. Natl. Inst. Stand. Technol. **103**, 259 (1998).

[27] O. Mandel, M. Greiner, A. Widera, T. Rom, T. Hänsch, and I. Bloch, Phys. Rev. Lett. **91**, 010407 (2003).

[28] R. Kosloff, in *Numerical Grid Methods and Their Application to Schrödinger's Equation*, edited by C. Cerjan (Kluwer Academic, Dordrecht, 1993), pp. 175–194.

[29] R. Kosloff, Annu. Rev. Phys. Chem. **45**, 145 (1994).

[30] R. Kosloff, in *Dynamics of Molecules and Chemical Reactions*, edited by R. E. Wyatt and J. Z. H. Zhang (Marcel Dekker, New York, 1996), pp. 185–230.

[31] M. Nielsen and I. Chuang, *Quantum Computation and Quantum Information* (Cambridge University Press, Cambridge, U.K., 2000).

[32] We employ standard computer science notation that $O(n)$ means scaling at most linear with n for large n.

[33] The precise value is $n-1$ but we ignore the constant for clarity and since we are more concerned with the scaling behavior.

[34] The function $\lceil x \rceil$ rounds up to the smallest integer greater than or equal to x, and the function $\lfloor x \rfloor$ rounds down to the largest integer less than or equal to x.

[35] Actually this should be $\lceil n/2^{d+1} \rceil$, but for clarity and simplicity we ignore this. It only contributes an extra term of $\lceil \log_2 n \rceil$ to the worst-case cost, which is negligible compared to the other terms.

[36] Actually this should be $\lceil n/2^{d+1} \rceil$, but for clarity and simplicity we ignore this. It only contributes an extra term of $2\lceil \log_2 n \rceil$ to the worst-case cost, which is negligible compared to the other terms.

[37] E. Raab, M. Prentiss, A. Cable, S. Chu, and D. Pritchard, Phys. Rev. Lett. **59**, 2631 (1987).

[38] M. T. DePue, C. McCormick, S. L. Winoto, S. Oliver, and D. S. Weiss, Phys. Rev. Lett. **82**, 2262 (1999).

[39] N. Schlosser, G. Reymond, I. Protsenko, and P. Grangier, Nature (London) **411**, 1024 (2001).

[40] A. J. Kerman, V. Vuletić, C. Chin, and S. Chu, Phys. Rev. Lett. **84**, 439 (2000).

[41] S. E. Hamann, D. L. Haycock, G. Klose, P. H. Pax, I. H. Deutsch, and P. S. Jessen, Phys. Rev. Lett. **80**, 4149 (1998).

[42] V. Vuletić, C. Chin, A. J. Kerman, and S. Chu, Phys. Rev. Lett. **81**, 5768 (1998).

[43] D.-J. Han, S. Wolf, S. Oliver, C. McCormick, M. T. DePue, and D. Weiss, Phys. Rev. Lett. **85**, 724 (2000).
[44] H. Perrin, A. Kuhn, I. Bouchoule, and C. Salomon, Europhys. Lett. **42**, 395 (1998).
[45] J. Vala, O. Dulieu, F. Masnou-Seeuws, P. Pillet, and R. Kosloff, Phys. Rev. A **63**, 013412 (2001).
[46] J. Vala and R. Kosloff, Opt. Express **8**, 238 (2001).
[47] O. Madelung, *Introduction to Solid-State Theory* (Springer, Berlin, 1978).
[48] A. Imamoğlu, D. D. Awschalom, G. Burkard, D. P. DiVincenzo, D. Loss, M. Sherwin, and A. Small, Phys. Rev. Lett. **83**, 4204 (1999).