

Adaptive quantum design of atomic clustersJason Thalken,¹ Yu Chen,² A. F. J. Levi,^{1,2} and Stephan Haas¹¹*Department of Physics and Astronomy, University of Southern California, Los Angeles, California 90089-0484, USA*²*Department of Electrical Engineering, University of Southern California, Los Angeles, California 90089-2533, USA*

(Received 19 November 2003; published 20 May 2004)

Adaptive quantum design identifies the best broken-symmetry configurations of atoms and molecules that enable a desired target function response. In this work, numerical optimization is used to design atomic clusters with specified quasiparticle densities of states. The dominant self-assembled building blocks of these engineered quantum systems are found to depend on the symmetry of the target function. For example, particle-hole symmetric spectra can be constructed from a dilute configuration of atomic dimers, whereas more complex structures such as trimers and quadrumers are required for asymmetric target functions. The convergence of the optimization algorithms depends on the shape of the target function, the density of atoms, and constraints due to substrates and boundary conditions. Hybrids of steepest-descent methods, simulated annealing, and genetic algorithms are found to be most efficient.

DOI: 10.1103/PhysRevB.69.195410

PACS number(s): 73.22.-f, 73.22.Gk, 61.46.+w, 74.78.Na

I. INTRODUCTION

In the near future, it will likely be possible to control the precise spatial positions of atoms and molecules using the experimental techniques now being developed by nanoscience. To complement these emerging capabilities it is clear that a new set of theoretical tools has to be developed to assist in the exploration of a potentially vast number of atom configurations and a corresponding enormous range of physical properties. In this paper, we outline an approach, which we call “adaptive quantum design,” that sets out to address this challenging task. In contrast to classical systems, atomic scale devices exhibit quantum fluctuations and collective quantum phenomena caused by particle interactions. Besides offering an excellent testing ground for models of correlated electrons, they also force us to reconsider conventional paradigms of condensed matter physics, such as crystal symmetries that are imposed by nature. In some instances, such symmetries need to be explicitly broken in order to enable or optimize a desired system response. Consider, for example, the quasiparticle density of states in tight-binding systems, which is the subject of the present study. In translationally invariant structures, i.e., crystals, it is well known that the spectral response function exhibits Van Hove singularities at positions of low dispersion, such as the band edges in a one-dimensional chain or the band center in a two-dimensional square lattice. These enhancements of the density of states can be very useful in amplifying system responses such as optical conductivity at specific incident energies. It is therefore important to be able to control the positions and shapes of such features by adaptive design techniques applied to models which capture the essential degrees of freedom of interacting atomic clusters.

Traditional *ad hoc* methods for the design of nanoscale devices will likely miss many possible configurations. At the same time, it is unrealistic to expect individuals to manually explore the vast phase space of possibilities for a particular device function. The proposed solution to this difficult design problem is to employ machine-based searches of configuration space that enable user-defined target functions.

Adaptive quantum design solves the inverse problem by numerically identifying the best broken-symmetry spatial configuration of atoms and molecules that produce a desired target function response.

The two major ingredients of adaptive quantum design are the physical model, which in this work evaluates the electronic density of states for a particular spatial arrangement of the atoms, and the search algorithm that finds the global minimum in the parameter space of all possible configurations. This problem is typically highly underdetermined in the sense that there can be several atomic configurations that yield a system response very close to the desired target function. Often, the associated landscape of solutions is shallow and has many nearly degenerate local minima.

Adaptive quantum design therefore relies heavily on efficient algorithms that accurately model the interacting system and find its optimal configuration, i.e., the global minimum in the available parameter space that yields the best match to the desired target function. This target response may be a specific angular transmission of a photonic crystal, a fine-tuned Josephson current along a superconducting junction array, or an energy-dependent density of states profile of an atomic cluster—the case on which we focus here. In this work, various numerical search tools are applied, including guided random walk, simulated and triggered annealing, and genetic algorithms. In particular the last set of techniques is most efficiently implemented on parallel computers. It is observed that in practice hybrids of these methods yield the best results, i.e., fast convergence towards a global minimum.

The particular example of a long-range tight-binding model is chosen for this study because it captures essential features of correlated quantum mechanical systems, and yet permits fast numerical diagonalization of relatively large clusters with broken translational symmetry. In a typical optimization run, these “function calls” occur 100–10 000 times. This model is therefore suitable for developing and testing adaptive design algorithms, and should be viewed as an initial step towards the design of atomic clusters. The goal of this work is to test the applicability and limits of

adaptive design techniques on a simple but nontrivial quantum system.

II. LONG-RANGE TIGHT-BINDING MODEL

The tight-binding approach is an effective tool to describe the band structure of electronic systems. It is commonly used to model the relevant bands close to the Fermi level, obtained from complex density-functional theory calculations. Since in this work symmetry-breaking, nonperiodic configurations are considered, a long-range variant of the tight-binding model with overlap integrals depending on the variable interatomic distance has to be used. Its Hamiltonian is given by

$$H = - \sum_{i,j} t_{i,j} (c_i^\dagger c_j + c_i c_j^\dagger), \quad (1)$$

where c_i^\dagger and c_i are creation and annihilation operators at a site \vec{r}_i , and the sum over pairs of atoms is restricted to avoid double counting. Here the spatial decay of the overlap integral $t_{i,j}$ is parametrized by a power law,

$$t_{i,j} = \frac{t}{|\vec{r}_i - \vec{r}_j|^\alpha}, \quad (2)$$

where the exponent α is taken to be 3.0 throughout the paper unless mentioned otherwise.¹ This parametrization reflects an algebraic variation of the overlap integral with interatomic separation. The choice of sign in the Hamiltonian follows the convention for s orbitals. However, this simple implementation of the tight-binding model does not account for the orbital directionality of realistic Au, Ag, or Pd atomic wave functions. More sophisticated and numerically expensive techniques, such as the local density approximation, would be needed to make quantitative predictions for these systems.

The Hamiltonian matrix of the long-range tight-binding model in the basis of single-particle states is nonsparse, and only its diagonal matrix elements vanish. In order to obtain the spectrum, the matrix is diagonalized numerically for finite clusters. In Fig. 1, the resulting densities of states of translationally invariant chains and square lattices are shown for the nearest-neighbor tight-binding model with $t_{i,j} = t \delta_{i,j}$ in Figs. 1(a) and 1(c), and for the case of long-range overlap integrals in Figs. 1(b) and 1(d). Characteristic Van Hove singularities are observed, in one dimension at the band edges, and in two dimensions at the band center. For the long-range model, the particle-hole symmetry is broken because of frustration introduced by the competing overlaps, leading to asymmetries in the density of states. While the system sizes in Fig. 1 are chosen to be rather large in order to make contact with the familiar thermodynamic limit, there are still some visible finite-size remnants, i.e., a faint pole structure due to the discreteness of the system. These features become much more pronounced for the few atom clusters that are studied in the following sections of this paper.

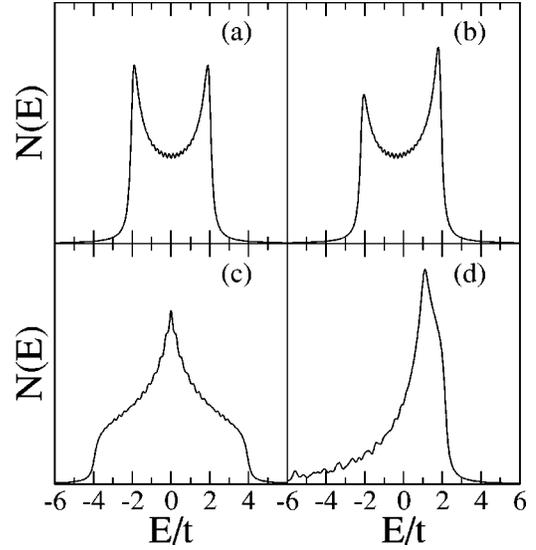


FIG. 1. Density of states in spatially invariant tight-binding systems with periodic boundary conditions. (a) Nearest-neighbor chain (no long-range overlap integral) with 30 atoms. (b) Same as (a), but with long-range overlaps according to Eq. (2). (c) Nearest-neighbor square lattice with 400 atoms. (d) Same as (c), but with long-range overlaps.

III. TARGET FUNCTIONS AND CONVERGENCE CRITERION

While “nature” gives us densities of states that are constrained by the dimensionality and symmetry of the underlying lattice, our objective is to engineer spectral responses with specific shapes that are useful in designing nanoscale devices. For example, we may wish to produce a quasi-two-dimensional spectrum in a one-dimensional system or to concentrate spectral weight in particular energy windows. These goals are achieved by placing the atomic constituents into optimized symmetry-breaking configurations which are determined by numerical searches.

The specific target functions that are studied in this work are shown in Fig. 2. They are (a) a flat top-hat density of states centered at $E=0$,

$$N(E)_{target} = \theta(E - E_c) \theta(E + E_c), \quad (3)$$

where $\theta(x)$ is the Heaviside function, and E_c is an energy cut off, (b) a symmetric two-peak function, centered at $E=0$,

$$N(E)_{target} = \theta(E - E_{c2}) \theta(E + E_{c1}) + \theta(E - E_{c1}) \theta(E + E_{c2}), \quad (4)$$

and (c) a particle-hole symmetry-breaking function with two unequal peaks, i.e., more spectral weight on the quasihole side ($E < 0$) than on the quasidelectron side ($E > 0$) of the spectrum,

$$N(E)_{target} = \theta(E - E_{c2}) \theta(E + E_{c1}) + \theta(E - E_{c4}) \theta(E + E_{c3}). \quad (5)$$

In systems with finite numbers of tight-binding atoms, these continuous shapes are approximated by equally spaced

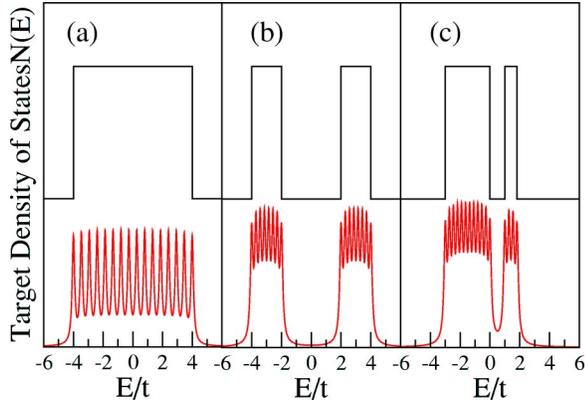


FIG. 2. Target densities of states used in this work: (a) top-hat function, centered at $E=0$, (b) particle-hole symmetric two-peak function, and (c) asymmetric two-peak function. For systems with finite numbers of particles these shapes are approximated by quasiparticle peaks.

poles within the energy windows where $N(E)_{target}$ is a non-vanishing constant. Here, the δ functions are given a finite width of $0.02t$. As more atoms are added to the system, these peaks merge together, approaching the bulk result. For other nonflat target functions the quasiparticle peak spacing can be varied, e.g., following a Gaussian or Lorentzian shape. Naturally, not all targets can be achieved equally well. Factors that influence the achievable match to a target include the number of available atoms and, as will be shown, a continuous or discrete number of accessible spatial positions. The dimensionality of the system poses additional constraints. In particular, there are less available configurations in lower dimensions. This study focuses on three prototype spectral responses which are targeted by numerically optimizing configurations of clusters with up to 48 atoms in two spatial dimensions.

The optimization algorithms seek to minimize the deviation from a given target density of states, defined by the error function

$$\Delta = \int_{-\infty}^{\infty} dE [N(E) - N(E)_{target}]^2, \quad (6)$$

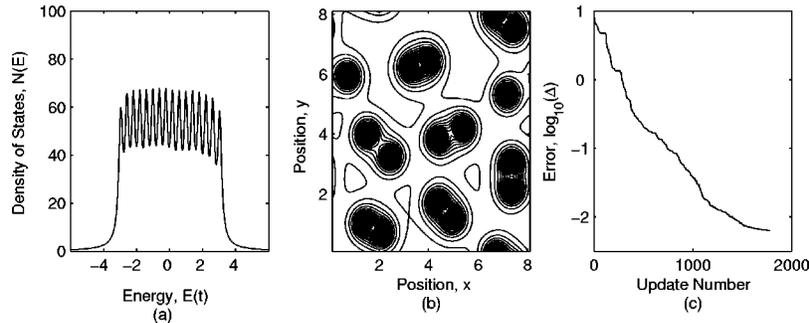


FIG. 3. Adaptive quantum design applied to 16 atoms in a two-dimensional box with periodic boundary conditions. The target density of states is a symmetric top-hat function with bandwidth $6t$. The atomic configurations are optimized by applying a guided random walk algorithm to the long-range tight-binding model. (a) The best matching solution. (b) Contour plot of the potential of the resulting spatial configuration. (c) Convergence to the target function, $\log_{10}(\Delta)$, with the number of updates.

which is the least-square difference between the system response for a given configuration and the target response function. We have explored a number of numerical techniques, including the Newton-Raphson steepest descent method, simple downhill random walk, simulated and triggered annealing, and genetic algorithms. The advantages and disadvantages of these techniques are briefly discussed in the Appendix. In general, it is found that hybrids of these methods tend to work best. In the following section, we focus on adaptive design of atomic clusters in continuous configuration space without any restrictions to underlying discrete lattices. In this case, the search space is infinite which generally allows better convergence to a given target response than in finite configuration space. However, some experimental realizations of such structures require deposition of atoms on substrates with discrete lattice structures.²⁻⁴ Therefore, this case is addressed separately in the subsequent section.

IV. ATOMS UP DESIGN OF TIGHT-BINDING CLUSTERS IN CONTINUOUS CONFIGURATION SPACE

The first target density of states we would like to study in detail is the particle-hole symmetric top-hat function, centered at energy $E=0$, shown in Fig. 2(a). This function represents a constant density of states for a bulk solid between the energy cutoffs $\pm E_c$, giving a bandwidth $W=2E_c$. Here, we choose $E_c=3t$. However, it should be noted that with the adaptive quantum design approach we are not restricted to this choice, and target densities of states with quite different bandwidths can be matched, although often to a lesser degree of accuracy.

In Fig. 3, the solution for a system with 16 atoms confined to a box with periodic boundary conditions is shown. The guided random walk method is applied to optimize the configuration of atoms by iterative local updates of their positions in order to match the top-hat density of states. As observed in Figs. 3(a) and 3(c), good convergence to the target function can be achieved for this case after less than 2000 updates. A contour plot of the potential $\sum_{ij} -t_{ij}/|\vec{r}_i - \vec{r}_j|^\alpha$ for the resulting spatial configuration is shown in Fig. 3(b). For this target function one discovers the formation of dimers with a wide range of interdimer and intradimer spacings.

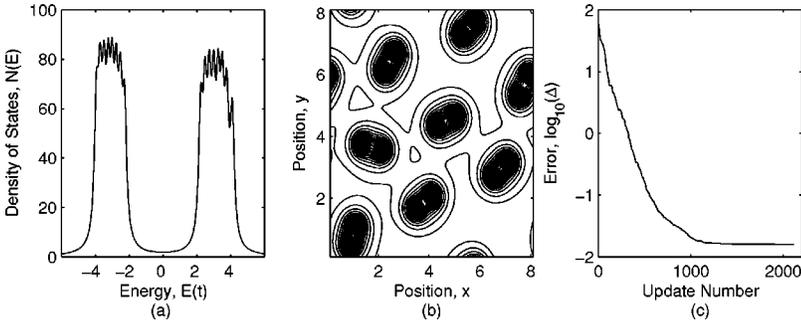


FIG. 4. Same as Fig. 3, but for a symmetric two-peak target function with peaks of bandwidth $2t$ centered at $-3t$ and $3t$.

These self-assembled building blocks have variable directional orientation and are closely packed.

To explore the generic aspects of this result, let us now turn to the two-peak target function shown in Fig. 2(b). This density of states has a gap in the center of the spectrum, as may be desired for the construction of two-state systems such as q bits. As shown in Figs. 4(a) and 4(c), the convergence to this target function is not quite as good as for the top-hat function, but saturation occurs already at half the number of iterations compared to the previous example. Interestingly, the optimized spatial configuration that is found in Fig. 4(b) also displays a preference for dimer formation.

In the long-range interacting systems we are studying it is, at least at first sight, not obvious why these target functions prefer dimer building blocks. Let us address this question by examining the individual spectra of the molecular building blocks shown in Fig. 5. Each dimer molecule contributes to the density of states a positive and negative pole with energies $E = \pm t_{12}$, where t_{12} is the hopping integral between the two participating atoms. The zero-energy quasiparticle peaks of two isolated atoms are split into bonding and antibonding combinations once they form dimer molecules. Therefore, isolated dimers are ideal building blocks for particle-hole symmetric densities of states, such as the top-hat and the two-peak target function. The intradimer spacing determines the positions of the $E = \pm t_{12}$ poles via Eq. (2). With an appropriate distribution of these distances the full-target spec-

trum of the top-hat function can be covered. The poles close to the band center ($E=0$) are provided by the less tightly bound dimers. For the two-peak target function, the dimer building blocks required to realize the target spectrum are more tightly bound, and the intradimer spacings need to vary less to achieve this target.

The idea that dimers can be used as building blocks for the particle-hole symmetric target functions strictly applies only to isolated dimers, i.e., the dilute limit, or when potential gradients across dimers pairs from the presence of adjacent atoms do not break particle-hole symmetry. The absence of such gradients, even for relatively high atom densities in a long-range interacting system, accounts for the success of dimers in satisfying the target function.

Lower symmetry building blocks, such as trimers and quadrumers shown in Fig. 5, can achieve more complex target functions, in particular, those with broken particle-hole symmetry. Due to frustration, trimer molecules intrinsically have asymmetric densities of states with unequal spectral weights on the electron and the hole side of their spectra. Quadrumers have a symmetric spectral response in the absence of longer-range frustrating interactions. As an example of how these building blocks enable more complex target functions, let us consider the asymmetric two-peak density of states given by Eq. (5) and shown in Fig. 2(c), which has a narrow upper peak and a wider lower peak, separated by a gap. In Fig. 6 it is observed that an approximate match can

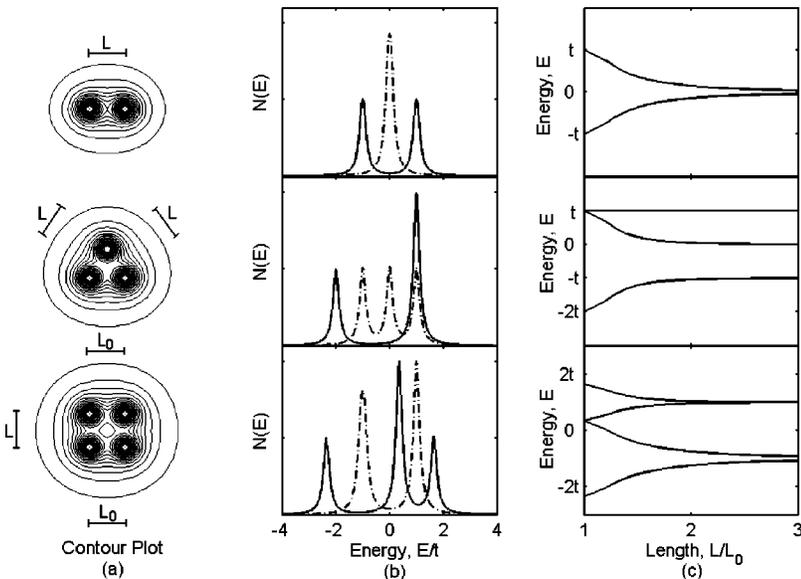


FIG. 5. Spectral response of tight-binding dimers, trimers, and quadrumers. (a) Potential contour plot of dimer, trimer, and quadramer of size L_0 . (b) In the dimer, the single-atom level (dashed line) is split symmetrically (solid line). Competing interactions (frustration) in the trimer and quadramer lead to asymmetric densities of states. The strength of the frustration depends on the separation L . (c) Spectral peak positions as a function of normalized distance L/L_0 for dimer, trimer, and quadramer. The spectra for $L/L_0=1$ are shown as solid lines and $L/L_0=3$ as dashed lines in (b).

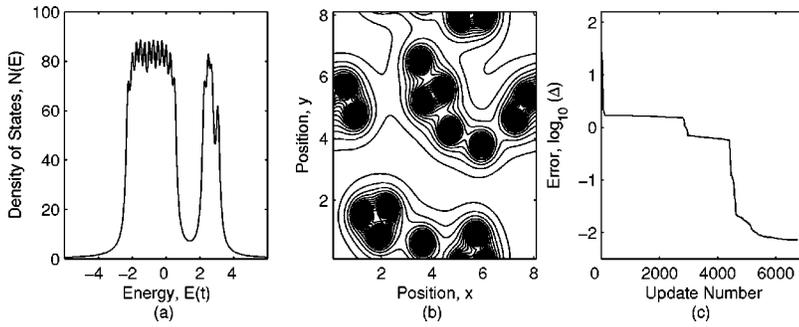


FIG. 6. Asymmetric two-peak quasiparticle density of states. The solution contains dimer, trimer, and quadrumer building blocks.

be achieved by the adaptive method. As expected, the building blocks for this particle-hole asymmetric target function are combinations of dimers, trimers, and quadrumers, which partially combine into larger clusters. Obviously, in order to achieve asymmetric target functions more complex building blocks are required. Especially for systems which contain only a small number of atoms it is important whether the required building blocks are available, and whether there are nonparticipating unbound atoms that may deteriorate convergence and match to the target function.

Let us finish this section by addressing the convergence properties of the three cases that were discussed. So far, only the simplest guided random walk optimization method was considered in which every “downhill step” is accepted. These downhill steps are local updates of individual atomic position that lead to a better match of the spectrum to the target density of states. Especially for the most symmetric target function, this algorithm converges efficiently, with a small remaining error. For the symmetric two-peak function, convergence occurs even faster because the required dimer building blocks are more uniform than for the first case. However, the remaining error is slightly larger, indicating that this may be a metastable solution which could be improved by global updates in which whole subclusters are simultaneously updated. Finally, the convergence plot for the asymmetric two-peak target function [Fig. 6(c)] shows several plateaus, indicating metastable configurations that exhibit a high resistance against local updates. For this case, a much larger number of local updates is required to achieve an acceptable match. Hence, more complex target functions clearly call for more sophisticated numerical search tools, including annealing steps, parallelization, and global updating schemes when available.

V. ADAPTIVE DESIGN IN DISCRETE CONFIGURATION SPACE

In the preceding section local updates were considered which allow atomic positions to change continuously within a given radius. However, for the case of atoms deposited on a substrate with a given lattice structure, the set of available positions is usually discrete, although it may be very large. This has significant consequences for the adaptive design approach. The search space of solutions is finite in this case, which makes it feasible to study more atoms with similar computational effort compared to the continuous case. At the same time, the discreteness of the lattice can prohibit favor-

able configurations that are available in the continuous case, thus deteriorating convergence properties of the optimization procedure. In practice, we find that the feasibility of computations for larger numbers of atoms in the discrete case helps to achieve better matches, as long as the lattice spacings remain sufficiently small. In this section, we explore the effects of an underlying grid on the adaptive design procedure. Also, more advanced techniques are implemented for the numerical optimization, including hybrids of the genetic algorithm, simulated annealing, and the guided random walk method.

In Fig. 7, adaptive design results are shown for systems with 48 atoms on a square lattice with spacing 0.01. The length scale is set by the linear size of the two-dimensional box (48×48) to which the particles are confined. These systems are in the dilute limit, and hence the convergence to the target functions, chosen to be the same as in the preceding section, is good. Also, because of the larger number of particles, there are less finite size effects. For the top-hat target function one obtains a small matching error of $\Delta = 0.000643$, for the symmetric two-peak function one finds $\Delta = 0.000605$, and for the asymmetric two-peak function the error is $\Delta = 0.150347$. Thus, analogous to the continuous case, the more symmetric targets are easier to achieve.

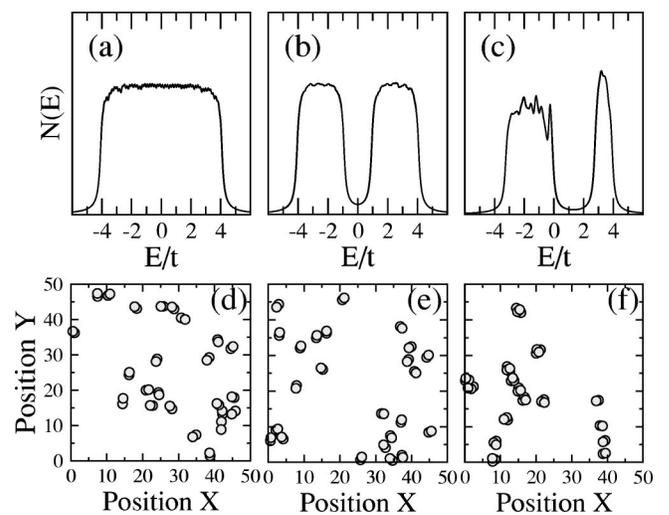


FIG. 7. Adaptive design of clusters with 48 atoms on a discrete lattice. The (a) flat, (b) symmetric two-peak, and (c) asymmetric two-peak densities of states are the same as discussed in the preceding section. The corresponding atomic configurations are shown in (d), (e), and (f).

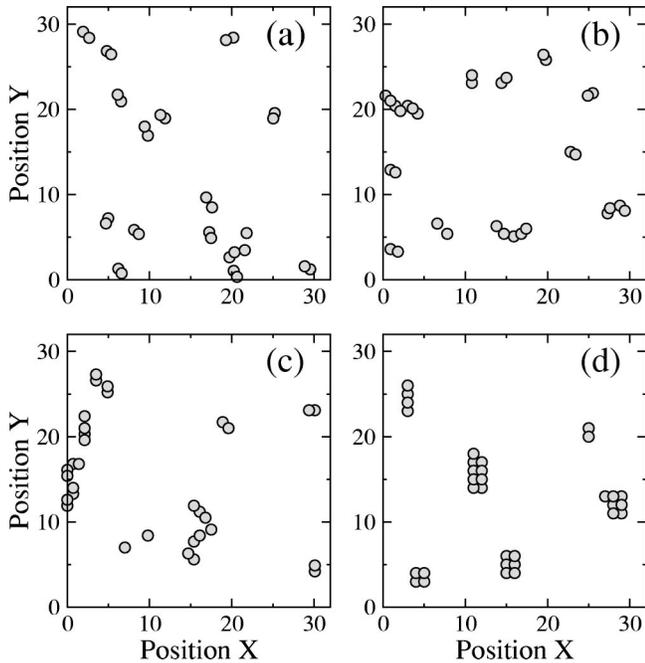


FIG. 8. Effect of the coarseness of the underlying lattice on the formation of building blocks. A total number of 32 atoms are confined to a 32×32 square box. The target function is the top-hat density of states. The grid spacing is varied: (a) 0.01, (b) 0.3, (c) 0.7, and (d) 1.0.

Again, clustering into dimers is observed for the particle-hole symmetric target functions, whereas trimers are the preferred building blocks for the asymmetric target. Since atomic densities in these examples were chosen to be in the dilute limit, boundary effects and interactions between the building blocks are small. The resulting configurations have the character of liquids, governed by weak interactions between the molecular building blocks, and relatively strong confining forces that lead to the formation of dimers and trimers.

Next, let us explore the dependence of these solutions on the coarseness of the underlying lattice. In Fig. 8, the grid spacing is varied over two orders of magnitude from 0.01 up to 1.0. As expected, the convergence to the target top-hat function deteriorates dramatically as the substrate is made coarser. For the smallest spacing of 0.01 one converges to a final error of $\Delta = 0.28457$, for a spacing of 0.3 the error is $\Delta = 2.45047$, for a spacing of 0.7 the error becomes $\Delta = 3.81779$, and for the coarsest case with spacing 1.0 the error is $\Delta = 36.3721$, indicating failure of convergence. The corresponding configurations in Fig. 8 show a strong dependence of the clustering sizes as the system is trying to cope with less available phase space to match the target function. For the finest grid spacing [Fig. 8(a)] one observes almost entirely isolated dimers. As the spacing is increased [Fig. 8(b)], a few small strings and groups are formed. At grid spacing 0.7 [Fig. 8(c)], the solution is made up mostly of long strings and dimers in close proximity to each other. Ultimately, for the coarsest grid spacing of 1.0 [Fig. 8(d)], the final configuration consists of square and rectangular blocks of atoms. This result demonstrates that there is a hi-

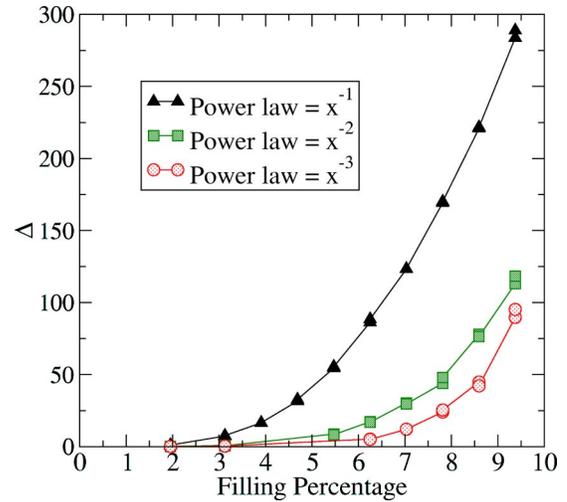


FIG. 9. Effect of the density of atoms and the power-law dependence of the tight-binding overlap integral on the convergence of adaptive quantum design.

erarchy of building blocks. The most suitable building blocks for the given target function are dimers. When these become less available due to lattice constraints, the adaptive method selects higher order solutions, i.e., larger size clusters, in order to cope with the more restricted phase space of possible configurations. Hence, for each level of coarseness, adaptive design discovers solutions that enable—up to a given degree of accuracy—a targeted system response.

By widening the grid spacings and keeping the box size constant, the density of atoms in the system is effectively increased, and simultaneously the available energy levels are spaced further apart. “Density” here is meant in terms of particles per available lattice sites, not particles per square unit. Let us examine these dependencies by varying the power law governing the atomic overlaps [Eq. (2)], and by increasing the number of atoms, i.e., the filling fraction, on a fixed lattice. Results for the achieved convergence are shown in Fig. 9. These demonstrate that excellent target matches can be achieved in the dilute limit with filling densities of a few percent. For larger coverages, the numerical search becomes exponentially less effective, indicating increasing frustration effects that need to be addressed by global updating schemes. Higher power laws imply effective shorter-range atomic overlaps, thus rendering the system more dilute. This is reflected in Fig. 9, where the departure from the regime of negligible matching errors is pushed toward higher filling percentages as α in Eq. (2) is increased.

Some aspects of the long-range tight-binding model on a substrate have already been confirmed experimentally using scanning tunneling microscopy (STM) to precisely position gold atoms on the surface of a nickel-aluminum crystal. In a recent study performed at UCI by Dr. Ho and co-workers,³ STM measurements show that the splitting in the value of eigenenergies for Au dimers on NiAl depends inversely on Au atom separation corresponding to $\alpha = 1$. Here, the grid spacings are dictated by the 0.29 nm lattice periodicity of available add-atom sites on the NiAl substrate. Remarkably, the power-law dependence of the effective overlaps t_{ij} be-

tween the deposited atoms takes into account interactions with the substrate which are typically difficult to model by first-principles computations.

VI. CONCLUSIONS

In this work, adaptive quantum design techniques were applied to tailor the quasiparticle density of states of atomic clusters, modeled by the long-range tight-binding Hamiltonian. Broken-symmetry spatial configurations of atoms were optimized to match target spectra. By applying adaptive search algorithms, it was shown that matches to target responses can be achieved by forming hierarchies of molecular building blocks that depend on system constraints. For example, symmetric top-hat and two-peak target densities of states can be achieved by forming lattices of weakly interacting dimers. While these are the elementary building blocks for particle-hole symmetric case target functions, more complex molecules, such as the trimer and the quadrumer, are found to dominate solutions for asymmetric target functions. Implementation of this approach on discrete lattices, corresponding to finite configuration space, introduces frustration effects that destabilize the elementary building blocks in the limit of coarse grid spacings.

The core task of adaptive quantum design is the numerical search for global minima in typically shallow landscapes of configurations with many local minima. Since this procedure typically requires many function calls, an efficient implementation on parallel computers is necessary. For the purposes of this study, the complexity of the physical model was minimized in order to limit the computational expense. While the long-range tight-binding model can be viewed as a semirealistic testing ground for adaptive quantum design techniques, it is crucial to apply these algorithms to more sophisticated models that include, among other ingredients, orbital directionality, spin degrees of freedom, and electronic correlations. Furthermore, adaptive design is applicable to related areas in nanotechnology, including the design of electro-optic components⁶ and RF systems.⁸

Our enthusiasm for adaptive quantum design is, in part, driven by the potential for scientific and technological discovery built into the methodology. We find it appealing to perform machine-based searches for new configurations, new components, and new subsystems, as we believe these will inevitably create new understanding, heuristics, and intuition. The fact that searches are performed in an ultralarge configuration space virtually guarantees the discovery of completely new designs and a much more thorough exploration of the possibilities and capabilities of nanotechnology.

ACKNOWLEDGMENTS

We are grateful to Ioan Gheorma, Weifei Li, Peter Littlewood, Omid Nohadani, Tommaso Roscilde, and Rong Yu for discussions. We acknowledge financial support by DARPA and the U.S. Department of Energy, Grant No. DE-FG03-01ER45908. Computational support was provided by the USC Center for High Performance Computing and Communications and by the NERSCC.

APPENDIX: OPTIMIZATION AND SEARCH ALGORITHMS

The optimization algorithms used in this work are based on local updates of atomic positions in order to minimize the error Δ defined in Eq. (6). Each atom in the system is visited periodically, and a trial change of its position is attempted. Depending on the response in Δ and on the specific algorithm, this trial step is either accepted or rejected. For the case of continuous configuration spaces, these local updates are random shifts of positions within a given radius. A hardcore constraint is implemented which forbids atoms to be placed on top of each other. For discrete configuration space, a stochastic distribution function is used to decide which sites in the neighborhood of the original position of an atom should be visited in a trial step. While this function is naturally peaked at nearest-neighbor sites, it has to include a finite probability of longer-range updates in order to avoid getting stuck in local minima of search space. The particular results discussed in this appendix are obtained for the continuous case.

The Broydn method is a multidimensional extension of the one-dimensional secant method.⁵ Unfortunately, its global convergence is rather poor for more than 20 variable parameters. Therefore, it is only applicable to the smallest system sizes, and is not useful for the nonlinear multiparameter searches required for adaptive quantum design.

In the guided random walk or “downhill method,” each random step that results in a smaller target error Δ is accepted.⁶ Random steps are trial spatial variations about a particle’s position. This guided random walk technique is a quickly implemented power horse. However, especially for shallow landscapes of solutions it gets easily stuck in local minima.

Simulated annealing uses an effective temperature representing the likelihood of accepting a step that does not minimize the function. This temperature is lowered slowly, at a rate of 10%, with each iteration. The initial temperature is taken as $T_{init} = 5t$. This method is better at avoiding local minima than the previous techniques. However, it takes a relatively longer time to converge.

The triggered annealing method is a hybrid of the downhill and the simulated annealing method. It implements the downhill method until minimizing steps become hard to find, at which point the simulated annealing method is used to escape from local minima. Parameters are chosen to be the same as for the simulated annealing method. Triggered annealing tends to converge relatively quickly if there are only a few local minima.

The particle replacement method uses the simple downhill method for guided random updates. In addition, it identifies particles that have not been updated for an extended period, because of being stuck in a local minimum, and assigns them to a new random position within the lattice boundaries. In our implementation, a particle is replaced if there are ten idle iterations without a successful downhill update for that particular particle. Note that this particle replacement update is different from random step updates because it is independent of the previous position of the particle, and thus it is particu-

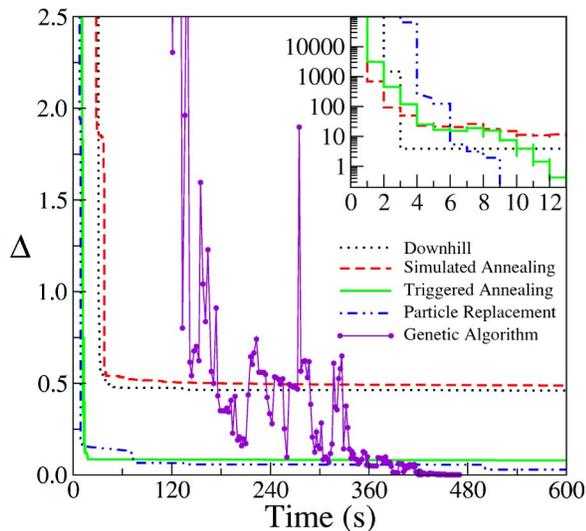


FIG. 10. Comparison of the convergence of several optimization algorithms. The error Δ is shown as a function of computer run time. The inset uses a logarithmic scale for Δ .

larly powerful in avoiding physical blocking of updates in lower dimensional systems.

In genetic algorithms a population of possible solutions is created. Those that best minimize the function are allowed to take part in creating a new generation of possible solutions.⁷ These methods are generally good at avoiding local minima, and are also easily implemented on parallel computers. They typically require more function calls than other search algorithms.

In order to illustrate the efficiency of these various approaches each method is used to match a flat top-hat target function on a one-dimensional lattice with 24 tight-binding atoms and a box size of 96. Particles exiting the box from one end enter it from the other side via periodic boundary

conditions. The target density of states is chosen to be the symmetric top-hat function, with poles evenly spaced between $E = -4t$ and $4t$. All of these benchmark runs are started with atoms placed randomly along a chain. The computer is a Pentium III 1 GHz with 2GB pc133 memory, and the additional seven nodes used by the genetic algorithm method are Pentium III 850 MHz processors with 1 GB memory. Each method is run for 600 s.

As shown in Fig. 10, all methods, with the exception of the genetic algorithm, converge rapidly within the first minute of run time, and show only relatively small corrections afterwards. The two techniques with the fastest convergence (inset of Fig. 10) are the downhill and the simulated annealing methods. However, their asymptotic error functions remain relatively large, indicating that they get easily stuck in local minima. In general, simulated annealing is expected to give better convergence than the simple downhill method. On longer time scales than shown in Fig. 10 this is indeed the case. The reason for the rather poor convergence within the first 10 min shown in this figure is that simulated annealing pushes atoms to explore metastable, frustrated configurations. However, for the one-dimensional case that was used for this benchmarking it is particularly difficult for this method to find its way out of such metastable states. In higher dimensions (less amenable to benchmarking) there is no such problem.

In contrast, the triggered annealing and particle replacement methods yield much better matches to the target function, while still converging relatively fast. As shown in the inset, the annealing methods sometimes accept trial steps in the “wrong” direction in order to avoid local minima. Finally, the genetic algorithm takes a relatively long time to converge. However, it yields by far the best match to the target density of states after about 7 min run time. In order to ensure best matches to the target, this last method is therefore used whenever the computational effort allows it.

¹W.A. Harrison, *Electronic Structure and the Properties of Solids* (W. H. Freeman, San Francisco, 1980).

²G.V. Nazin, X.H. Oiu, and W. Ho, *Science* **308**, 77 (2003); *Phys. Rev. Lett.* **90**, 216110 (2003).

³N. Nilius, T.M. Wallis, M. Persson, and W. Ho, *Phys. Rev. Lett.* **90**, 196103 (2003); N. Nilius, T.M. Wallis, and W. Ho, *Science* **297**, 1853 (2002).

⁴T.M. Wallis, N. Nilius, and W. Ho, *Phys. Rev. Lett.* **89**, 236802 (2002).

⁵W.H. Press, S.A. Teukolsky, W.T. Vetterling, and B.P. Flannery, *Numerical Recipes: The Art of Scientific Computing* (Cambridge University Press, Cambridge, 1986).

⁶Y. Chen, R. Yu, W. Li, O. Nohadani, S. Haas, and A.F.J. Levi, *J. Appl. Phys.* **94**, 6065 (2003).

⁷See, e.g., D. Whitley, *Stat. Comput.* **4**, 65 (1994).

⁸I.L. Gheorma, S. Haas, and A.F.J. Levi, *J. Appl. Phys.* **95**, 1420 (2004).