Genetic algorithm optimization of atomic clusters

J. R. Morris, D. M. Deaven, K. M. Ho, C. Z. Wang,

B. C. Pan, J. G. Wacker and D. E. Turner

Ames Laboratory and Department of Physics

Iowa State University

Ames, Iowa 50011-3020

Abstract

We have been using genetic algorithms to study the structures of atomic clusters and related problems. This is a problem where local minima are easy to locate, but barriers between the many minima are large, and the number of minima prohibit a systematic search. We use a novel mating algorithm that preserves some of the geometrical relationship between atoms, in order to ensure that the resultant structures are likely to inherit the best features of the parent clusters. Using this approach, we have been able to find lower energy structures than had been previously obtained. Most recently, we have been able to turn around the "building block" idea, using optimized structures from the GA to learn about systematic structural trends. We believe that an effective GA can help provide such heuristic information, and (conversely) that such information can be introduced back into the algorithm to assist in the search process.

1 Introduction

One of the main interests in genetic algorithms (GA's) is their application to difficult optimization problems. In many such problems, there are many possible solutions, with no clear way to conduct a systematic search. The problem is compounded when there are many optimization parameters, especially if these parameters are strongly coupled. We have recently applied genetic algorithms to one type of these problems, the determination of the lowest energy configurations of a collection of atoms. This is a particularly difficult problem because the number of metastable local energy minima grows exponentially with the number of degrees of freedom available to the system [1, 2]. In many cases, especially in systems with strong directional covalent bonds like carbon or silicon, metastable energy states are separated by large barriers reflecting the high energy cost of breaking bonds to rearrange a cluster's structure. Attempts to use traditional simulated annealing methods [3] to find the global energy minimum usually fail, leaving the system trapped in one of

the numerous metastable configurations. Thus, unless some of the rules are known for constructing the ground state structure of the atomic cluster, we do not have reliable predictions for the global energy minimum for clusters once their sizes go beyond 10 to 20 atoms.

It is clear that an algorithm is needed which can 'hop' from one minimum to another and permit an efficient sampling of phase space. Genetic algorithms, in principle, can do this using an optimization strategy modeled on the Darwinian evolution process. The fitness of a population of candidate solutions to the problem is improved by selecting a group of best-performing candidates to act as parents. A "mating" operation is then performed among the parents to produce children to form the next generation of candidates. This process is repeated until the best solution is located [4, 5].

A key ingredient for the success of a genetic algorithm is the efficiency of the mating operation in producing good solutions to the problem. Traditional implementations of the genetic algorithm, using various forms of string recombination as the mating operator, have been tried previously on the present problem of molecular geometry optimization, but without much success [6].

We feel that the problem of the traditional mating operation is that it does not efficiently transfer favorable properties from the parents to the child. With this in mind, we have constructed a mating operation based upon the physical structure of the parent clusters, in such a way that resulting structures are much more likely to inherit promising characteristics of the parent structures. We have applied our approach to a number of problems, including carbon clusters [7], the Thomson problem of point charges on a sphere [8], and Lennard-Jones clusters [9]. In each case, we have found the genetic algorithm to be helpful in locating optimal solutions. In the most recent (and challenging) application of locating low energy silicon clusters, the genetic algorithm has led to a new understanding of the structural trends - potentially more useful than knowing the lowest energy states of particular clusters.

2 Method

Before discussing the genetic algorithm approach, we should first point out some of the challenges in accurately determining the structures of atomic clusters. (We only discuss theoretical problems; there are a number of experimental difficulties as well.) The atomic arrangement will usually be a low energy one; therefore, our searches are focused on finding the lowest energy configurations. However, calculating the energy of a cluster is not trivial. First-principles total energy calculations are quite accurate, but are computationally intensive. Therefore, searching many different configurations using this approach is not practical. In some cases, reasonably accurate empirical models exist, which produce fairly reliable energies and structures. However, if there are several structures that are close in energy (which often happens), then these empirical predictions must be supported by more accurate techniques. Even when the empirical potentials are reliable, as appears

DISCLAIMER

Portions of this document may be illegible in electronic image products. Images are produced from the best available original document.

to be the case for carbon [10], the calculations of the energies and the local optimization of the structures still require significant computational resources. Given that there are a large number of local optima (as discussed in the results section below), it is therefore important to explore possible low energy structures as efficiently as possible.

To achieve this goal, we devised the following genetic algorithm approach. Starting with a population of candidate structures, we relax these candidates to the nearest local minimum using a conjugate-gradient minimization or molecular dynamics quenching. Using the relaxed energies as the criteria of fitness, a fraction of the population is selected as parents. The mating operation is performed by a "cut-and-paste" procedure: First, we choose a random plane passing through the center of mass of each parent cluster. We then cut the parent clusters in this plane, and assemble the child from the atoms of one parent which lie above the plane, and the atoms of the other parent which lie below the plane. If the child generated in this manner does not contain the correct number of atoms, the parent clusters are translated an equal distance in opposing directions normal to the cut plane so as to produce a child which contains the correct number of atoms. The resulting child structure is then relaxed to the nearest local minimum.

3 Results

3.1 Carbon clusters

We first illustrate the above algorithm by an application to the case of C_{60} . In this simulation, the interatomic interaction in the carbon clusters is described by a tight-binding potential previously shown to be very accurate for description of fullerene structures [10, 11]. Mating operations are performed every 30 time steps: 4 structures with the lowest energies are chosen as parents from a population of 16 candidate structures. We adopted an 'elitist' strategy in which the 4 parents are included in the next generation together with 12 new structures formed by mating operations between different parents. Each new structure is relaxed using molecular dynamics (MD) quenching. After 30 time steps, the energy of the new structure is examined. If the new structure is not within a certain energy range (~ 0.2 eV) of the energy of the parent structures, it is discarded and replaced by another child. Otherwise, it is fully relaxed to its lowest energy state (this usually takes about 200 MD steps). If the energy of the fully relaxed structure is lower than any of the parent structures, it is incorporated into the parent population and the highest energy parent structure discarded.

Figure 1 shows the results of the C_{60} simulations. The energy per atom is plotted for the lowest energy (solid line) and highest energy (dashed line) structure in the population as a function of the number of MD steps expended by the simulation. Several generic features of the algorithm are apparent in these results: In the initial stage, the energy drops very quickly and the population soon consists of reasonable candidates. A sampling of the structures of the population during the initial period indicate a sequence of structures of unfinished cages very similar to that observed in simulated annealing.

Figure 1: Running the genetic algorithm on C₆₀.

The initial stage usually occupies only a small fraction of the total time spent by the algorithm. The results from the initial stage resembles what would be obtained by the simulated annealing method, except that the genetic algorithm obtains it more efficiently by quenching without going through long annealing procedures. The rest of the time is spent in an end game, where the remaining defects in the structure are removed until the ground state structure is found.

3.2 Thomson problem

In view of the success of the genetic algorithm in solving the structure of carbon clusters, we decided to test the robustness of the scheme in applications to other kinds of cluster problems. The first is the well-studied historical Thomson problem of finding the arrangement of N charges constrained on a sphere [2]. We studied this problem to observe the effect of long range forces on our algorithm. The problem is also of interest in the study of symmetry-breaking in the ground state and metastable structures due to the constrained finite number of particles. The similarity of the spherical geometry to the cage structures of the C₆₀ clusters also encouraged us to apply our approach to this problem.

This problem has been used as a test case for optimization, and there are a number of results for systems less than 100 particles [2, 16, 17]. From these studies, the most reliable results had come from *random* searches, rather than simulated annealing approaches.[2] The difficulty apparently arises in the fact that most simulated annealing studies would use one starting structure and try to optimize that, rather than using multiple initial structures.

Our approach was quickly able to locate the known results (for $N \leq 132$), and to extend these results up to 200 particles [8]. For less than about 80 particles, the problem is fairly trivial, as there are only a handful of possible solutions. However, the number of solutions rapidly grows, with an estimated 270 possible local minima for $N = 132$ and nearly 8000 for $N = 200$ [2]. The fact that we (apparently) could locate the minimum for most of these cases demonstrates the efficiency of our approach.

3.3 Lennard-Jones clusters

We have also studied the case of Lennard-Jones clusters to observe the behavior of the algorithm on compact clusters (as opposed to the cage-like structures that occur in the previous cases). This is a simple interaction, modeling the interaction between noble gas atoms such as argon. There are many studies of these clusters for $N \leq 100$ [18, 19, 20], and the number of known local minima grows much more rapidly than that of the Thomson problem – for $N = 55$, the estimate is on the order of 10^{21} minima. Again, we find that our algorithm efficiently reproduces most of the ground state structures reported (with exceptions for the $N = 75, 76$ and 77 Lennard-Jones clusters; see Ref. [20]). In several cases, we were even able to locate structures with lower energies than those previously reported. (In one case, we found two different structures with lower energies than had been found previously.)

3.4 Silicon clusters

Most recently, we have been studying Si clusters [21]. Unlike the previous cases, there are few known results except for the smallest of clusters ($N \leq 10$), and even structural trends are not well understood. Further, there have not been any empirical potentials whose results were confirmed by more accurate calculations. Even well-accepted potentials which describe bulk behavior reasonably well produce non-physical results for clusters. Most empirical models predict that the $N = 20$ cluster would be higher in energy than two $N = 10$ clusters, an unphysical result. The lack of success in determining the structures of these clusters demonstrates the challenge of this problem.

Recently, our group has produced a more accurate Si potential, and we have begun applying it to Si clusters. The accuracy of the potential is very important: if the potential leads to unphysical clusters, or over-estimates the energies of clusters that are actually low energy, then even a sophisticated search algorithm will not lead to important results. On the other hand, if the potential can be used to identify low energy candidates correctly,

Figure 2: Structures based upon the $N = 9$ silicon sub-unit.

then these clusters may be checked using more accurate calculations. Rather than use the genetic algorithm to find the lowest energy structure (within our empirical model), we have used it to generate a set of low energy solutions, producing a smaller search space that was then studied using more accurate treatments.

Using a combination of a new Si potential and the genetic algorithm, we have identified new low-energy structures in the range from $N = 10$ to $N = 20$. The genetic algorithm approach is essentially the same as that of the carbon clusters. To ensure a good global search, 10–20 ecologies were run for each value of N . Also, we performed simulated annealing to produce other candidate structures, and to compare the effectiveness of this approach with that of the GA. We found that for clusters larger than 13 atoms, the GA was able to produce lower energy candidate structures, indicating that this is a better search strategy.

For $N < 13$, our results are consistent with previous predictions. However, the $N = 13$ structure that we found is significantly lower in energy than previous results. Unlike previous calculations, we find that this structure is built upon that of a nine-atom sub-

Figure 3: Cage-like low energy structures for Si_{17} , Si_{19} and Si_{20} .

unit. Although this sub-unit (shown in fig. 2) is not the lowest energy $N = 9$ structure, there is a family of clusters based upon this sub-unit (or several such units, for $N \geq 18$) for all clusters in this study. These structures are shown in fig. 2; with the exceptions of $N = 9, 17, 19$ and 20 , these are the lowest energy structures that we have found.

For $N \geq 17$, we have found another family of structures that are more cage-like than the previous structures. These are the lowest energy structures that we have found for $N = 17, 19$ and 20 . These structures, shown in fig. 3, have one central atom, surrounded by a shell of atoms. Such “stuffed fullerene” clusters have been suggested previously for larger clusters [22, 23, 24, 25].

4 Discussion

The work presented here demonstrates that our approach is successful for a number of applications. The fact that one mating approach is applicable to these reasonably disparate problems is striking, especially given the strong differences in geometries. These are problems that other approaches or traditional GA’s have not been able to address. The difficulty is not only that we must rapidly explore many minima; with the large number of possible structures available, we must also ensure that favorable “traits” of the parents be preserved through the generations. In our case, that corresponds to good local structures. Using string operations to generate new structures will not do this.

In most applications, a successful genetic algorithm must be able to preserve the relationship between strongly coupled parameters. This will only be successful for traditional linear operations (such as crossover) if the coupling is extremely simple and is also known in advance. We feel that any GA application using such operations will be inefficient, if the parameters are strongly interdependent. (Of course, if the search space is sufficiently simple, then such operations combined with local optimization may be successful. In

such cases, however, a genetic algorithm will not be preferable to a random search or to simulated annealing.) It is clear that mating algorithms that respect important relationships between related parameters may be much more efficient. For our applications, the approach we described preserves most local atomic environments, while searching the complicated solution space in an unbiased way. The arbitrary numbering or ordering of the atoms is not taken into account in the "mating" process; rather, it is the physical arrangement that is important, and that is used to generate new structures.

In many applications, searching for optimal solutions is aided by using "heuristic" information. If some factor is known about good solutions, then building this into the search can greatly speed up the search. However, if this "known factor" is not necessarily a part of the solution, forcing solutions to include such information may miss other good solutions. In the case of the Lennard-Jones particles, most results have been obtained by growth rules, searching for a structure based upon the results of smaller structures. By avoiding such biases, we were able to locate the optimal structure for $N = 38$, which is not similar to any other optimal clusters for $N < 100$.

In the case of the Si clusters, we were able to turn around the "building block" idea. The genetic algorithm generated a number of structures which contained a "structural unit" in the optimal structure. It is not surprising that some core arrangement of atoms exist, but in our case, it was the genetic algorithm that led to the particular common sub-unit. By knowing what sub-unit may appear, we may also try to extend results by building structures by hand, using the identified sub-units. This might be able to extend our predictive abilities, even when the problem is sufficiently complex that our GA breaks down. Furthermore, this suggests the idea of "seeding" the GA with some possible sub-units. If we put likely sub-units into the initial population, then the search may more efficiently locate low energy structures. If there are several possible sub-units, then such seeding may help broaden the search, preventing one class of structures from dominating the population.

In general, we believe that searches through complex solution spaces may be aided by genetic algorithms. However, blind application of bit-string mating operations seems unlikely to efficiently improve solutions, unless the parameters may be optimized more or less independently. Our application demonstrates that genetic algorithms may be constructed in such a way that complex interrelationships between parameters may be preserved by the mating process, leading to new solutions that might otherwise be missed. Our approach might be made more optimal in a number of ways; some variations and additions are currently being explored. We feel, however, that these applications are an interesting and useful demonstration of an alternative approach to genetic algorithms, and may provide some insight towards other challenging optimization problems.

5 Acknowledgments

Ames Laboratory is operated for the U.S. Department of Energy by Iowa State University under contract no. W-7405-Eng-82. This work was supported by the Director for Energy Research, Office of Basic Energy Sciences, and the High Performance Computing and Communications Initiative. Part of this work was made possible by the Scalable Computing Laboratory, which is funded by Iowa State University and Ames Laboratory.

References

- [1] L. T. Wille and J. Vennik, *J. Phys. A* **18**, L419 (1985).
- [2] T. Erber and G. M. Hockney, *Phys. Rev. Lett.* **74**, 1482 (1995); review article in *Adv. Chem. Phys.* **98**, 495 (1997).
- [3] S. Kirkpatrick, *Science* **220**, 671 (1983); D. Vanderbilt, *J. Comput. Phys.* **56**, 259 (1984).
- [4] J. H. Holland, *Adaptation in natural and artificial systems* (Ann Arbor: The University of Michigan Press, ©1975).
- [5] D. E. Goldberg, in *Genetic Algorithms in Search, Optimization, and Machine Learning*, Addison-Wesley ©1989.
- [6] For some examples, see Y. Xiao and D. E. Williams, *Chem. Phys. Lett.* **215**, 17 (1993); T. Brodmeier and E. Pretsch, *J. Comput. Chem.* **15**, 588 (1994); D. B. McGarrah and R. S. Judson, *J. Comput. Chem.* **14**, 1385 (1993); J. A. Niesse and H. R. Mayne, *Chem. Phys. Lett.* **261**, 576 (1996).
- [7] D. M. Deaven and K. M. Ho, *Phys. Rev. Lett.* **75**, 288 (1995).
- [8] J. R. Morris, D. M. Deaven and K. M. Ho, *Phys. Rev. B* **53**, 1740 (1996). Recent updates may be found at <http://cmp.ameslab.gov/~jrmorris/thomson.html>.
- [9] D. M. Deaven, N. Tit, J. R. Morris and K. M. Ho, *Chem. Phys. Letters* **256**, 195 (1996).
- [10] C. H. Xu, C. Z. Wang, C. T. Chan and K. M. Ho, *J. Phys. Cond. Mat.* **4**, 6047 (1992).
- [11] B. L. Zhang, C. Z. Wang, and K. M. Ho, *Chem. Phys. Lett.* **193**, 225 (1992); C. Z. Wang, B. L. Zhang, K. M. Ho, and X. Q. Wang, *Intn. J. Mod. Phys. B* **7**, 4305 (1993).
- [12] J. C. Grossman, L. Mitas, and K. Raghavachari, *Phys. Rev. Lett.* **75**, 3870 (1995).

- [13] E. A. Rohlfing, D. M. Cox, and A. Kaldor, *J. Chem. Phys.* **81**, 3322 (1984).
- [14] D. Tománek and M. A. Schluter, *Phys. Rev. Lett.* **67**, 2331 (1991).
- [15] See for example P. Ballone and P. Milani, *Phys. Rev. B* **42**, 3201 (1990).
- [16] J. T. Wille, *Nature* **324**, 46 (1986).
- [17] E. L. Altschuler, T. J. Williams, E. R. Ratner, F. Dowla and F. Wooten, *Phys. Rev. Lett.* **72**, 2671 (1994); **74**, 1483 (1995).
- [18] J. A. Northby, *J. Chem. Phys.* **87**, 6166 (1987).
- [19] N. J. A. Sloane, R. H. Hardin, T. D. S. Duff, and J. H. Conway, *Discrete Comput. Geom.* **14**, 237 (1995).
- [20] J. P. K. Doye and D. J. Wales, *J. Chem. Phys.* **103**, 4324 (1995).
- [21] K. M. Ho, B. C. Pan, C. Z. Wang, J. G. Wacker, D. E. Turner and D. M. Deaven, submitted to *Phys. Rev. Letters*.
- [22] E. Kaxiras, *Chem. Phys. Lett.* **163**, 323 (1989).
- [23] J. L. Elkind, J. M. Alford, F. D. Weiss, R. T. Laaksonen, and R. E. Smalley, *J. Chem. Phys.* **87**, 2397 (1987); L. R. Anderson, S. Maruyama, and R. E. Smalley, *Chem. Phys. Lett.* **176**, 348 (1991).
- [24] U. Rothlisberger, W. Andreoni, and M. Parrinello, *Phys. Rev. Lett.*, **72**, 665 (1994).
- [25] J. Pan and M. V. Ramakrishna, *Phys. Rev. B* **50**, 15431 (1994); M. V. Ramakrishna and A. Bahel, *J. Chem. Phys.*, **104**, 9833 (1996).

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.