

## ARTICLES

# Minima hopping: An efficient search method for the global minimum of the potential energy surface of complex molecular systems

Stefan Goedecker<sup>a)</sup>*Institut für Physik, Universität Basel, Basel, Switzerland*

(Received 19 January 2004; accepted 5 March 2004)

A method is presented that can find the global minimum of very complex condensed matter systems. It is based on the simple principle of exploring the configurational space as fast as possible and of avoiding revisiting known parts of this space. Even though it is not a genetic algorithm, it is not based on thermodynamics. The efficiency of the method depends strongly on the type of moves that are used to hop into new local minima. Moves that find low-barrier escape-paths out of the present minimum generally lead into low energy minima. © 2004 American Institute of Physics.

[DOI: 10.1063/1.1724816]

## I. INTRODUCTION

Finding the global minimum of the potential energy surface of a complex system is a central problem in physics, chemistry, and biology. For a periodic system the global minimum gives the crystalline ground state structure of a solid, for nonperiodic systems it determines the geometric ground state configuration of molecules. Many complex biomolecules exist in nature whose structure is unknown. In the case of proteins the ground state is called the native state and the theoretical determination of this native state is considered to be one of the major challenges of modern biology.

Because the experimental determination of the geometric ground state of complex systems is very difficult, simulation is a promising candidate for the theoretical determination of the structure of such systems. For instance, several unknown alloy structures were recently found by simulation methods.<sup>1</sup> Finding the ground state structure of large systems is computationally very expensive and has only become possible on the latest generation of computers. The fundamental reason why finding the global minimum is so expensive is that the number of local minima increases exponentially with respect to the number of atoms in the system. Initially, during the simulation, the system is trapped in one of these local minima, and in order to reach the global minimum the system has to travel through many intermediate local minima basins and it has to overcome the barriers separating the local minima. A basin is, by the conventional definition, a certain part of the configurational space around a minimum of the potential energy surface. More precisely, a basin contains all the configurations that will relax into this minimum using simple small-step downhill relaxations. We will denote as a super-basin the union of several neighboring basins. If one can arrive from any point in such a super-basin at the lowest minimum of this super-basin without crossing barriers that are very high compared to the average difference in energy

between local minima, it will be called a funnel.<sup>2</sup> The large majority of algorithms for finding the global minimum in physical systems are based on thermodynamic principles. Because of the ubiquitous presence of the Boltzmann factor  $\exp(-\Delta E/k_b T)$  in these algorithms, climbing and crossing over high barriers that separate for instance different funnels is a rare event that frequently does not happen during the available computer time.

## II. FINDING THE GLOBAL MINIMUM WITHOUT THERMODYNAMICS

With the exception of genetic algorithms,<sup>3</sup> most standard algorithms such as simulated annealing,<sup>4</sup> basin hopping,<sup>19</sup> and multicanonical methods<sup>5</sup> are based on thermodynamic principles. A Markov process based on the Metropolis algorithm with a Boltzmann factor leads finally to a thermodynamic distribution. At sufficiently low temperature the ground state configuration will be the dominant configuration and hence the problem is solved in principle. Unfortunately thermodynamics does not tell us anything about how fast the thermodynamic equilibrium distribution is obtained and as a matter of fact it can be extremely slow. Consequently, global minimization strategies based on thermodynamics are of questionable value.

Instead of invoking thermodynamic principles, the basic principle in global optimization should be that one strives to explore as fast as possible the low energy part of the configurational space. Revisiting configurations is obviously detrimental for finding new configurations. All ordinary Monte Carlo like algorithms have the tendency to revisit many times neighboring configurations that are close in energy, if other configurations can only be reached by crossing high barriers. As a matter of fact the simulation can just jump back and forth between two such configurations for a very long time.

The problem of repeated visits of certain configurations has already been recognized by many researchers. One rem-

<sup>a)</sup>Electronic mail: stefan.goedecker@unibas.ch

edy that was proposed is flooding.<sup>6</sup> The principle of flooding is simple. In basins that were already visited during the simulation, the potential is lifted and consequently it is less likely that a configuration in the same region of the continuous configuration space will be accepted in a future Monte Carlo step. The problem with this approach is that it is very difficult to determine exactly where the potential should be lifted. Ideally it should be done over the whole basin, but the determination of the basin boundaries would be prohibitively expensive computationally. Because of this problem flooding has not found widespread use.

Because it is so difficult to determine the shape of a basin that has to be flooded, another version of flooding has been proposed recently. The flooding is not done in the very high dimensional configurational space, but in a low dimensional space spanned by some suitably defined order parameters of the system.<sup>7</sup> The difficulties of determining the shape of the basin to be flooded are not really eliminated in this method, but are alleviated since the space dimension is much lower.

There is a second more profound problem with flooding. For a simple system, where essentially all basins can be visited during the simulation, energy surface flooding is bound to succeed. Successively all the basins will be flooded until the system finds its way into the global minimum. In a more realistic high dimensional setting with a very large number of basins, things can be different. Imagine a transition basin that is a low energy transition region among several funnels. We postulate that direct transitions between these funnels are unlikely since they are separated by high barriers from the other funnels, with the exception of the transition basin. In order to find the global minimum the simulation has to jump into all the different funnels and each jump requires a visit of the transition basin. Flooding the transition basin will make such jumps more unlikely and will thus slow down the search for the global minimum.

What is needed is a strategy that limits repeated visits, but does not penalize crossings through important transition basins. As will be shown, this can be achieved by doing more violent escape moves out of the current basin if this basin has already been visited. This then gives rise to the minima hopping method.

### III. MINIMA HOPPING

The minima hopping method consists of an inner part that performs jumps into the local minimum of another basin and an outer part that will accept or reject this new local minimum. The acceptance/rejection is done by simple thresholding, i.e., the step is accepted if the energy of the new local minimum  $E_{\text{new}}$  rises by less than  $E_{\text{diff}}$  compared to the current energy  $E_{\text{cur}}$ . The parameter  $E_{\text{diff}}$  is continuously adjusted during the simulation in such a way that half of the moves are accepted and half are rejected. This outer part introduces a preference for steps that go down in energy. However if the inner part proposes only steps that go up in energy, such steps will finally also be accepted after  $E_{\text{diff}}$  has been sufficiently increased after many rejections.

The inner part consists of an escape step away from the current local minimum followed by a geometry relaxation

into the closest local minimum. In the present implementation, the geometry relaxation is done by a combination of standard steepest descent and conjugate gradient methods. The escape step is done by a short molecular dynamics simulation that starts from the current minimum. The atoms have a Boltzmann velocity distribution such that their kinetic energy is equal to  $E_{\text{kinetic}}$ . The system thus has sufficient energy to cross over a barrier of height less than  $E_{\text{kinetic}}$  measured relative to the current minimum. If  $E_{\text{kinetic}}$  is small, one will usually fall back into the current minimum, if it is sufficiently big, one will most likely be ejected from the current basin and end up in a different minimum. The molecular dynamics simulation is stopped as soon as the potential energy has crossed  $m_{\text{dmin}}$  maxima and reached the  $m_{\text{dmin}}$ -th minimum along the trajectory. At this point the geometry relaxation starts. Different molecular dynamics simulations (with different random velocity Boltzmann distributions) are started until the geometry relaxation will give a new local minimum.

Three different cases can be distinguished for the outcome of this inner part. In the first case the geometry relaxation will give back the current local minimum that was used as the starting point for the displacements. The second case is that the new minimum is one that has already previously been visited during the simulation (i.e., that has been accepted in a previous outer acceptance/rejection step). The third case is that the minimum is a new one that has not been visited previously. The third case is the desirable one, since it will result in the exploration of new configurations.

Since it is desirable to explore new configurations one might think that it is most advantageous to choose  $E_{\text{kinetic}}$  very large. This would be true if the height of the new local minimum “behind” the barrier was independent of the barrier height. If the height of the barrier is correlated with the height of the new local minimum, the situation is entirely different. Since the density of states of the local minima (i.e., the number of local minima per unit energy) grows exponentially with energy, we would have to explore a much larger number of local minima if we allow high barriers to be crossed. As described in detail later, we found a very clear correlation between the barrier height and the height of the new minimum. For this reason one has to make a compromise in the choice of  $E_{\text{kinetic}}$  to obtain the shortest possible overall simulation time. If  $E_{\text{kinetic}}$  is too large, the discrete space of the local minima that has to be explored becomes too large; if it is too small, one has to start a very large number of molecular dynamics simulations until an escape path is possibly found. In practice, it has turned out that adjusting  $E_{\text{kinetic}}$  dynamically during the simulation in such a way that half of all molecular dynamics simulations will bring us into new basins is close to optimal. It has to be stressed that molecular dynamics is very powerful to find low barrier paths into other basins. Out of several alternatives we examined, it gave by far the lowest barriers. If another method were used, it might be necessary to allow for a much large fraction of the escape trials to fall back into the current minimum in order to avoid large values of  $E_{\text{kinetic}}$ .

The value of  $E_{\text{kinetic}}$  is not only increased if one falls back into the current minimum, but also if we get into an-

other minimum that had already been visited. If the simulation starts to “walk” around between minima that were already visited,  $E_{\text{kinetic}}$  will be increased at each molecular dynamics restart until enough energy is available to cross into a new unexplored region of the discrete configurational space. This procedure does not restrict repeated visits to strategic transition basins between several basins. If such a minimum of a transition basin is revisited after extensively and repeatedly exploring some funnel, the simulation will arrive at the transition state with a large value of  $E_{\text{kinetic}}$  and it will therefore go on to other states that are far away from this transition basin. This is a very desirable effect.

In order to decide whether  $E_{\text{kinetic}}$  will be increased or decreased we have to keep track of all the minima visited previously. In the context of force field calculations, where the energy of a configuration can be determined with arbitrary accuracy, the configuration can be identified only by their energy. Introducing a history evidently destroys the Markovian character of the simulation. However, this is not a disadvantage since we do not intend to generate any thermodynamic distribution.

A flowchart of the algorithm is shown below. It contains five parameters.  $\alpha_1$  and  $\alpha_2$  determine how rapidly  $E_{\text{diff}}$  is increased or decreased in the case where a new configuration is rejected or accepted.  $\beta_1$ ,  $\beta_2$ , and  $\beta_3$  determine how rapidly  $E_{\text{kinetic}}$  is modified depending on the outcome of an escape trial.

```

initialize a current minimum 'Mcurrent'
MDstart
ESCAPE TRIAL PART
  start a MD trajectory with kinetic energy  $E_{\text{kinetic}}$ 
  from current minimum 'Mcurrent'. Once the
  potential reaches the  $m_{\text{min}}$ -th minimum
  along the trajectory stop MD and optimize
  geometry to find the closest local minimum 'M'
  if ('M' equals 'Mcurrent') then
     $E_{\text{kinetic}} = E_{\text{kinetic}} * \beta_1$  ( $\beta_1 > 1$ )
    goto MDstart
  else if ('M' equals a minimum visited previously)
  then
     $E_{\text{kinetic}} = E_{\text{kinetic}} * \beta_2$  ( $\beta_2 > 1$ )
    goto MDstart
  else if ('M' equals a new minimum) then
     $E_{\text{kinetic}} = E_{\text{kinetic}} * \beta_3$  ( $\beta_3 < 1$ )
  endif
DOWNWARD PREFERENCE PART
  if (energy('M') - energy('Mcurrent') <  $E_{\text{diff}}$ ) then
    accept new minimum: 'Mcurrent' = 'M'
    add 'Mcurrent' to history list
     $E_{\text{diff}} = E_{\text{diff}} * \alpha_1$  ( $\alpha_1 < 1$ )
  else if rejected
     $E_{\text{diff}} = E_{\text{diff}} * \alpha_2$  ( $\alpha_2 > 1$ )
  endif
  goto MDstart

```

In the simulations, which will be presented later, the values  $\beta_1 = \beta_2 = 1/\beta_3 = 1.05$  and  $\alpha_1 = 1/\alpha_2 = 1.02$  were used. Increasing  $\beta_2$  or  $\beta_3$  leads to a less thorough search. With the above-mentioned values the simulations never failed to find




FIG. 1. Illustration of the Bell–Evans–Polanyi principle. The potential energy is described by two parabolas. The transition state is the intersection of both parabolas. Evidently the transition state is raised if the local minima on the right-hand side is raised (dashed line) and vice versa.

the minimum. With increased values the global minimum is found faster in the cases where it is found, but there are also cases where it can be missed. After the system has explored the low energy configurations it starts to explore higher energy regions. This is a consequence of the fact that  $E_{\text{kinetic}}$  is increased whenever a known low energy configuration is revisited. This increase is a necessary condition for the success of the minima hopping method and it is encountered if  $\beta_2^{\alpha_1 \alpha_2} \beta_3 \geq 1$ . Once  $E_{\text{kinetic}}$  has increased significantly, the simulation should be stopped because the global minimum was most likely found if the search was thorough enough. If the simulation is not stopped a nonperiodic system will finally explode.

#### IV. SIGNIFICANCE OF THE BELL–EVANS–POLANYI PRINCIPLE FOR GLOBAL OPTIMIZATION

The Bell–Evans–Polanyi principle<sup>8</sup> states that highly exothermic chemical reactions have a low activation energy. In our language this means that it is more likely to find a low energy local minimum if one crosses from the current basin over a low barrier into a new basin than if one overcomes a high barrier. As a matter of fact, the Bell–Evans–Polanyi principle is more likely to be well satisfied in the context of global optimization than in the context of single chemical reactions. In the context of global optimization it has to hold only in an average sense and not for each individual barrier crossing. A difference is that the Bell–Evans–Polanyi principle is a statement about the true physical transition state, i.e., about the lowest barrier separating two basins. In the context of global optimization one crosses usually not from one basin to another via the exact transition state but over a higher energy barrier. Nevertheless the Bell–Evans–Polanyi principle should hold equally well for reaction paths that are not going through the exact physical transition state. The basic assumption of the Bell–Evans–Polanyi principle, namely that the potential energy landscape along the reaction path can be approximated by two parabolas centered at the two minima (Fig. 1), is valid for any reasonable reaction path that does not involve enormous barriers.

TABLE I. Correlation between the average barrier height  $\langle E_{\text{kinetic}} \rangle$  and the number of minima visited before the global minimum.  $\langle E_{\text{kinetic}} \rangle$  is the average of the fluctuating quantity  $E_{\text{kinetic}}$  over the entire simulation.

Method	$\langle E_{\text{kinetic}} \rangle$	Number of minima visited before global minimum
Random displacement	$\geq 300$	34 000
MD, $\beta_1 = 1.05$ , $mdmin = 1$	4.7	830
MD, $\beta_1 = 1.02$ , $mdmin = 1$	3.6	510
MD, $\beta_1 = 1.01$ , $mdmin = 1$	3.2	430
MD, $\beta_1 = 1.01$ , $mdmin = 2$	.5	410

Numerical results clearly show that the Bell–Evans–Polanyi principle holds for the reaction paths of ordinary escape trials. Table I shows results for a system that has been studied extensively, the 38-atom Lennard-Jones cluster.<sup>9</sup> These studies have shown that the system has two funnels. The funnel containing the ground state contains only a relatively small number of other minima. The second funnel, on the other hand, contains a lot of minima. Therefore practically all runs first fall into this broad funnel. For this reason we have actually chosen the deepest minimum of this funnel, which is the second deepest overall minimum, as the starting point of our simulation.

Even though molecular dynamics (MD) gives always reasonably low barriers, the parameters  $\beta_1$  and  $mdmin$  can be tuned to vary the barrier height. By reducing the value of  $\beta_1$  the system is permitted to perform more unsuccessful escape trials and will thus find on average lower barriers. By increasing the value of  $mdmin$  the system can oscillate more frequently within the basin from which it wants to escape or even jump over two barriers and is therefore also more likely to find lower barriers. The last four lines of Table I show these MD results. The first line of Table I was obtained in a run where the escape trials were not done by MD, but by random displacements of the atoms. The step size of the random displacements was also adjusted such that half of all escape trials will lead into new basins. Random displacements can lead to some astronomically large barriers. For this reason  $E_{\text{kinetic}}$  is in this case (Fig. 1) not an average barrier, but a typical low barrier height.

All the results show a clear correlation between the average value  $\langle E_{\text{kinetic}} \rangle$  and the number of minima that are visited by the simulation before the global minimum is found. This confirms the validity of the Bell–Evans–Polanyi principle. Lower barriers lead on average into lower new basins. Hence higher energy regions that contain a very large number of minima are avoided and the global minimum is found by traversing a smaller number of local minima. Note that our test is particularly stringent because we can not only go down in energy. In order to leave the initial funnel we have to go up in energy as well, but we do not go higher than needed.

The number of minima that are traversed before finding the ground state is subject to large fluctuations. As shown in Fig. 2 this number has a nearly perfect exponential distribution. For this reason the values in Table I were averaged over 1000 runs in the case of MD and over 100 runs in the case of random displacements.

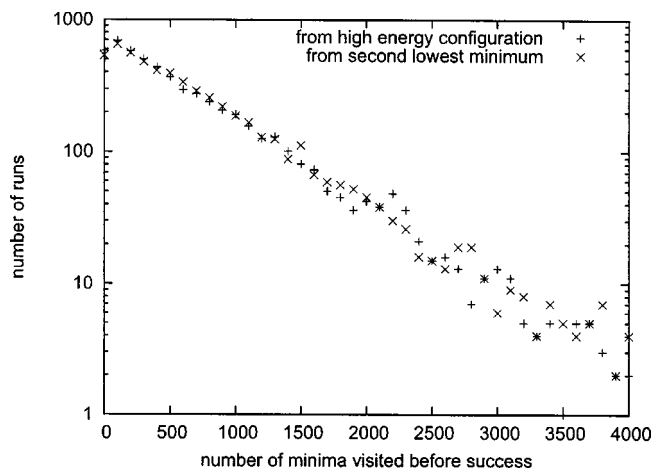


FIG. 2. Number of runs that found the global minimum after traversing a certain number of other minima. The results are given both for the case where the initial configuration is the second deepest minimum and for a high energy starting configuration.

Even lower barriers could be obtained by eigenvector following methods.<sup>10</sup> However, this introduces some preferred search direction and thus practically excludes certain transitions into neighboring basins. In agreement with other studies<sup>11</sup> we therefore found that such choices, which greatly reduce the randomness of the search direction, slow down the search for the global minimum.

What counts in practice is not so much the number of minima that are found before the global minimum, but the overall computing time. Since each step of the minima hopping algorithm becomes more expensive as  $\beta_1$  is decreased or  $mdmin$  increased, one has to find a set of values for  $\beta_1$  and  $mdmin$  that represents a good compromise between the number of steps and the cost of each step.

## V. COMPARING WITH SIMILAR METHODS

### A. Basin hopping

Basin hopping<sup>19</sup> has turned out to be a powerful Monte Carlo method for the determination of the global minimum of complex molecular systems. It has been applied to various realistic systems such as silicon clusters.<sup>12</sup> Basin hopping is a method that transforms the potential energy surface in an advantageous way. The value of the potential energy surface within one basin is replaced by the value of the potential energy at the associated minimum as shown in Fig. 3. The common explanation for the success of the basin hopping method is that it eliminates the barriers between the basins of different minima. This is true, but nevertheless, it has to be stressed that the basin hopping method does not eliminate the barriers between super-basins or funnels. The transformed piecewise constant potential energy surface of the basin hopping method still exhibits barriers that have to be overcome by Monte Carlo steps. If the height of these remaining barriers of the transformed surface between super-basins is small compared to the height of the original barriers of the untransformed surface between the basins (upper panel of Fig. 3), the basin hopping method is expected to offer a significant advantage, otherwise (lower panel of Fig. 3), the




FIG. 3. In the upper part the barriers separating basins are very high compared to the energy differences of the various local minima, in the lower part both are comparable.

advantage will be marginal. As we have seen, the barrier height encountered in a simulation depends strongly on the kind of moves. Random displacements lead to particularly high barriers.

Without introducing a history list of all the previously visited minima, the basin hopping algorithm and the minima hopping algorithm have a rather similar behavior, even though their theoretical principles are different. In contrast to basin hopping, minima hopping is not a Monte Carlo method. What makes the significant difference in practice is the feedback introduced by the history. As a consequence the minima hopping method can climb out of a “wrong” funnel much faster than the basin hopping method. It is thus super-

rior to the basin hopping method for systems that have a deep “wrong” funnel. Wrong means in this context just that the funnel does not contain the global minimum.

Being a thermodynamic method, basin hopping also has the inconvenience that a temperature has to be chosen. Choosing a nonoptimal value can prevent finding the global minimum.

## B. Activation relaxation technique

The minima hopping method is essentially a sequence of activation relaxation events followed by an acceptance/rejection step. Whereas in the ART method<sup>13</sup> one tries to find very low barriers and in the ART nouveau method<sup>11</sup> even the exact transition state, the molecular dynamics scheme that we have proposed will in general go over higher barriers. The findings on the validity of the Polyani principle would suggest that ART is even better suited than MD for minima hopping. The reason why we have chosen MD instead of ART is that it is not obvious how to introduce a feedback parameter such as  $E_{\text{kinetic}}$  into the ART scheme. The existence of such a parameter is essential for a history based scheme. Only such a parameter allows us to make more violent jumps that explore new regions of the configuration space when the simulation has the tendency to repeatedly revisit certain configurations of the space.

## C. Temperature accelerated molecular dynamics

Temperature accelerated molecular dynamics<sup>14</sup> is a scheme that also uses molecular dynamics to escape from a basin. The temperature is chosen sufficiently high, so that the system will not remain very long in one basin. A modified version of temperature accelerated molecular dynamics<sup>15</sup> also introduces a history that allows for some simplifications if a basin is revisited during the simulation. However, the primary purpose of temperature accelerated molecular dynamics is not to find the global minimum, but to follow a system over very long time scales. Therefore temperature accelerated molecular dynamics does not use the history to make more violent moves, which would not have a physical counterpart. Temperature accelerated molecular dynamics




FIG. 4. The history of all the minima visited in the search for the ground state of a 38 atom Lennard-Jones cluster. The energy was shifted such that the ground state has energy 0.




FIG. 5. A crystal that is rotated slightly with respect to the ground state of a silicon crystal (left panel) together with the true ground state (right panel).

also does not have an acceptance/rejection step that gives a preference toward lower energy configurations.

## VI. NUMERICAL TESTS

Doye and Wales have studied Lennard-Jones clusters in great detail and mapped out the structure of their local minima configurations.<sup>9</sup> Their disconnectivity graphs for these systems allow one to predict very well whether the global minimum of a system is difficult or easy to find. As a consequence of these detailed investigations the Lennard-Jones cluster can be considered as a benchmark system for any new global minimization algorithm. The minima hopping method was therefore also applied to these clusters and the putative global minimum found by Doye and co-workers was easily rediscovered in all cases as shown in Fig. 4 for the 38-atom cluster. On a Pentium based PC the average run time varied between a few minutes for the 38-atom cluster to a few hours for the 75-atom cluster. The run time depends of course sensitively on implementation details and in particular on the quality of the local minimization algorithm. When comparing with similar methods, other more objective measures will therefore be used in the following.

The performance of the minima hopping method was also compared with the performance of the basin hopping method using exactly the same type of moves, namely random displacements.<sup>16</sup> For a simple funnel structure, as in the 19-atom Lennard-Jones cluster, the performance of both methods is quite similar. For a system with two deep funnels minima hopping is superior since it climbs out of the wrong funnel much faster. Averaging again over 100 runs for the 38-atom cluster, basin hopping visited some 75 000 minima before falling into the global minimum, compared to 34 000

for minima hopping (Table I). For the basin hopping run we choose the temperature that gives the fastest success and which corresponds to 1.2 energy units.

To test our algorithms for more difficult systems with more atoms we turned to silicon as described by the EDIP interatomic potential.<sup>17</sup> This system has the advantage that the global minimum is known for any number of atoms. The ground state is the perfect crystal. The simulations were started using an entirely wrong simple cubic crystal structure, where each silicon atom has six nearest neighbors instead of four. With fixed cell size the silicon systems have wrong funnels that contain a very large number of minima. Physically the local minimum at the bottom of a wrong funnel corresponds to a perfect silicon crystal that is rotated with respect to the correct lattice vectors as shown in Fig. 5 for the case of a 64-atom crystal. Rotating the crystal in the correct position is not possible with periodic boundary conditions. So the crystal has to become amorphous before it can it can recrystallize in the correct position. This obviously requires overcoming very high barriers.

We applied the minima hopping method to 64- and 216-atom crystals and always found the ground state. These system sizes are nearly two orders of magnitude larger than the crystalline systems for which up to now the global minimum was found by simulation.<sup>18</sup> For the silicon systems it turned out that it is impossible to find the global ground state without our history feedback mechanism. Both basin hopping as well as simulated annealing methods fail<sup>20</sup> for this system. They all get stuck in the astronomically large number of amorphous local minima. Figure 6 illustrates how minima hopping succeeds in finding the global minimum. The algorithm first goes down fairly smoothly. Since it is not difficult to find new lower energy minima both  $E_{\text{kinetic}}$  and  $E_{\text{diff}}$  are relatively small. Once the algorithm has arrived at the bottom of a wrong funnel (after some 2.2 million minima), configurations begin to be revisited frequently and  $E_{\text{kinetic}}$  starts to increase. Because of the violent moves the new configurations are on average higher in energy and  $E_{\text{diff}}$  has to increase as well. This is clearly visible as some broadening of the energy history curve. At some point the algorithm has finally climbed out of the wrong funnel and has arrived at the correct funnel. The system then collapses very rapidly into the global ground state after having visited some 4 million minima. Note that this number is very small compared to the




FIG. 6. The history of all the minima visited in the search for the ground state of a 216 atom silicon crystal using the minima hopping method. The energy was shifted such that the ground state has energy 0.

total number of minima. Considering only Wooten–Winer–Weaire processes,<sup>21</sup> the estimated number of local minima is  $2^{216} \approx 10^{65}$ .

Many different physical processes can be invoked to cross from one local minimum into another one.<sup>22</sup> The results show that all these processes are activated by the escape trials using molecular dynamics.

## VII. CONCLUSIONS

To find the global minimum of a potential energy surface that has multifunnel structure requires an algorithm that can rapidly climb out of wrong funnels. This feature can only be obtained by abandoning the standard Markov-based Monte Carlo methods and by introducing a feedback mechanism that, based on the whole simulation history, enforces the exploration of new regions of the configuration space. The minima hopping method, presented in this article, contains such a feedback mechanism.

A preference toward low energy configurations is obtained by an acceptance/rejection step. Imposing this preference is simpler if the proposed moves already have a tendency to go into other low energy minima. As expected from the Bell–Evans–Polanyi principle, low energy barriers lead on average to low energy minima. The molecular dynamics scheme that we use for the escape part gives reasonably low barriers and is applicable to any atomistic system.

## ACKNOWLEDGMENTS

The author would like to thank Christoph Bruder, Wladimir Hellmann, and Cyrus Umrigar for useful comments on the manuscript, the Swiss National Science Foundation for its financial support, and the staff of the computing center at the University of Basel for their technical support.

- <sup>1</sup>G. Johansson, T. Bligaard, A. Ruban, H. Skriver, K. Jacobsen, and J. Nørskov, *Phys. Rev. Lett.* **88**, 255506 (2002).
- <sup>2</sup>S. Chan and K. Dill, *Proteins: Struct., Function, Genetics* **30**, 2 (1998); <http://laplace.compbio.ucsf.edu/danny/Protein/fig-gallery.html>
- <sup>3</sup>D. Goldberg, *Genetic Algorithms in Search, Optimization, and Machine Learning* (Addison–Wesley, New York, 1989).
- <sup>4</sup>P. Salomon, P. Sibani and R. Frost, *Facts, Conjectures and Improvements for Simulated Annealing* (SIAM, Philadelphia, 2002).
- <sup>5</sup>B. Berg and T. Neuhaus, *Phys. Lett. B* **267**, 249 (1991).
- <sup>6</sup>H. Grubmueller, *Phys. Rev. E* **52**, 2893 (1995); A. Voter, *J. Chem. Phys.* **106**, 4665 (1997).
- <sup>7</sup>A. Laio and M. Parrinello, *Proc. Natl. Acad. Sci. U.S.A.* **99**, 12562 (2002).
- <sup>8</sup>F. Jensen, *Computational Chemistry* (Wiley, New York, 1999).
- <sup>9</sup>D. Wales, M. Miller, and J. Doye, *J. Chem. Phys.* **110**, 6896 (1999); J. Doye, M. Miller, and D. Wales, *ibid.* **111**, 8417 (1999).
- <sup>10</sup>L. Munro and D. Wales, *Phys. Rev. B* **59**, 3969 (1999).
- <sup>11</sup>R. Malek and N. Mousseau, *Phys. Rev. E* **62**, 7723 (2000).
- <sup>12</sup>S. Yoo and X. Zeng, *J. Chem. Phys.* **119**, 1442 (2003).
- <sup>13</sup>G. T. Barkema and N. Mousseau, *Phys. Rev. Lett.* **77**, 4358 (1996); N. Mousseau and G. T. Barkema, *Phys. Rev. E* **57**, 2419 (1998).
- <sup>14</sup>M. Sorensen and A. Voter, *J. Chem. Phys.* **112**, 9599 (2000).
- <sup>15</sup>F. Montalenti and A. Voter, *J. Chem. Phys.* **116**, 4819 (2002).
- <sup>16</sup>The basin hopping program that can be obtained from <http://www-wales.ch.cam.ac.uk/software.html> contains special angular random moves that are more efficient than ordinary three-dimensional random moves.
- <sup>17</sup>M. Z. Bazant and E. Kaxiras, *Phys. Rev. Lett.* **77**, 4370 (1996); M. Z. Bazant, E. Kaxiras, and J. F. Justo, *Phys. Rev. B* **56**, 8542 (1997); J. F. Justo, M. Z. Bazant, E. Kaxiras, V. V. Bulatov, and S. Yip, *ibid.* **58**, 2539 (1998); S. Goedecker, *Comput. Phys. Commun.* **148**, 124 (2002).
- <sup>18</sup>D. Williams, *Acta Crystallogr., Sect. A: Found. Crystallogr.* **52**, 326 (1996); R. Wawak, J. Pillardy, A. Liwo, K. Gibson, and H. Scheraga, *J. Phys. Chem. A* **102**, 2904 (1998).
- <sup>19</sup>D. Wales and J. Doye, *J. Phys. Chem. A* **101**, 5111 (1997); J. Doye and D. Wales, *Phys. Rev. Lett.* **80**, 1357 (1998); J. Doye, D. Wales, and M. Miller, *J. Chem. Phys.* **109**, 8143 (1998).
- <sup>20</sup>Several runs were terminated after two weeks of running time due to no visible progress toward finding the global minimum. The basin hopping method was done with molecular dynamics moves, which are more efficient than random moves. In the simulated annealing runs the temperature was decreased exponentially.
- <sup>21</sup>F. Wooten, K. Winer, and D. Weaire, *Phys. Rev. Lett.* **54**, 1392 (1985); S. Goedecker, T. Deutsch, and L. Billard, *ibid.* **88**, 235501 (2002).
- <sup>22</sup>F. Valiquette and N. Mousseau, *Phys. Rev. B* **68**, 125209 (2003).