

Energy landscape of relaxed amorphous silicon

Francis Valiquette and Normand Mousseau

*Département de Physique and Centre de Recherche en Physique et Technologie des Couches Minces,
Université de Montréal, C.P. 6128, succ. Centre-ville, Montréal, Québec, H3C 3J7, Canada*

(Received 24 January 2003; revised manuscript received 1 July 2003; published 29 September 2003)

We analyze the structure of the energy landscape of a well-relaxed 1000-atom model of amorphous silicon using the activation-relaxation technique (ART nouveau). Generating more than 40 000 events starting from a single minimum, we find that activated mechanisms are local in nature, that they are distributed uniformly throughout the model, and that the activation energy is limited by the cost of breaking one bond, independently of the complexity of the mechanism. The overall shape of the activation-energy-barrier distribution is also insensitive to the exact details of the configuration, indicating that well-relaxed configurations see essentially the same environment. These results underscore the localized nature of relaxation in this material.

DOI: 10.1103/PhysRevB.68.125209

PACS number(s): 61.43.Dq, 66.30.Hs, 02.70.-c

I. INTRODUCTION

The dynamics of many complex materials is dominated by activated jumps over energy barriers generally higher than $k_B T$. For these systems, the energy-landscape picture, which focuses on the topological relation between locally stable states, has proven very valuable. Disconnectivity graphs, for example, introduced by Czerminski and Elber¹ and applied extensively by others,²⁻⁶ have provided a first classification of the dynamics of complex systems based on the structure of their respective energy landscape.

In parallel with these developments, oriented towards reconstructing the topology of the energy landscape, there has been some effort in trying to characterize the energetics and the nature of the events in clusters,^{3-5,7} proteins,^{6,8,9} and in amorphous materials.¹⁰⁻¹³ This extensive sampling is essential in order to try to connect the properties of the landscape with the dynamics measured experimentally, it also serves to build a better understanding of the generic properties of various networks: low vs high connectivity, bulk vs finite-size systems, covalent vs ionic bonding, etc.

In this paper, we present an extensive study of the structure of the energy landscape of amorphous silicon around two minima. Less extensive studies of this material were already presented both by our group¹⁰⁻¹² and Middleton *et al.*,¹³ using two different approaches. Using the activation-relaxation technique (ART nouveau), we generate more than 42 000 activated events around a well-relaxed minimum and analyze the properties of the reaction paths connecting this initial minimum to a nearby saddle point and new minimum. We find that: (i) the activation mechanisms in *a*-Si are local, limiting somewhat the usefulness of the configurational energy landscape picture, (ii) the activation energy is essentially limited by the energy required to break a single bond, (iii) the entropic barrier is almost identical for all events, (iv) more than 20% of all events generated are bond switches of the Wooten-Winer-Weaire type, and (v) the number of events seems to be of the order 30–60 per atom.

II. DETAILS OF THE SIMULATION

We study here the energy landscape around two well-relaxed 1000-atom configurations of *a*-Si. We focus mostly

on the first one, repeating the simulation on a second configuration only to ensure that the results are generic and not dependent on some specific feature of the configuration.

The energy model used is a modified version (mSW) of the empirical Stillinger-Weber (SW) potential^{14,15} developed to ensure that the elastic and structural properties of the amorphous phase of silicon are in agreement with experiment.¹⁵ With its original parameters, the SW potential describes both the liquid and the crystalline phases with good accuracy but fails to reproduce the experimental structure of the amorphous phase.¹⁶⁻¹⁸ Recently, Vink and collaborators proposed a slightly different set of parameters (see Table I), which generates the right amorphous structure in addition to providing the correct vibrational properties. Because of its empirical nature and the fact that it is not optimized for dynamics, the energy barriers computed with this potential must be taken with some care. However, since the structural properties of this material are well described by mSW, the qualitative features of the energy landscape are expected to be correct.

A. ART nouveau

While our previous study of relaxation in *a*-Si used a version of ART that could not identify saddle points,^{12,18,19} the results presented here are obtained using ART nouveau, the latest version of the activation-relaxation technique pre-

TABLE I. Parameters of the original Stillinger-Weber potential as well as those modified by Vink *et al.*, as described in Ref. 15. We use the latter set in this paper.

Parameter	standard SW	modified SW
ϵ (eV)	2.168 26	1.648 33
A	7.049 556 277	7.049 556 277
B	0.602 224 558 4	0.602 224 558 4
σ (Å)	2.0951	2.0951
p	4	4
a	1.80	1.80
λ	21.0	31.5
γ	1.20	1.20




FIG. 1. Solid line: radial distribution function (RDF) of the 1000-atom α -Si model used in the simulation; dashes: experimental RDF measured by Laaziri *et al.* (Ref. 23).

sented in Refs. 7 and 9. Using ideas similar to those proposed by Munro and Wales,²⁰ ART nouveau applies a Lanczos scheme to compute directly the lowest second derivative of the energy during the activation phase and ensures convergence to the saddle point to any desired precision.

Events are generated in the following way: In order to leave the harmonic well, one atom is selected at random. This atom and its neighbors, contained in a shell of radius 3.5 Å, are moved iteratively in a randomly chosen direction, while keeping the energy, projected in the perpendicular directions to a minimum. At each step, a Lanczos scheme is used to compute the lowest eigenvalues and eigenvectors associated with the curvature of the energy landscape. We consider that the configuration has left the harmonic well when the lowest eigenvalue falls below $-50 \text{ eV}/\text{Å}^2$.

The activation process *per se* starts from this point. The configuration is slowly pushed up along the direction of lowest curvature until the modulus of the force falls below $0.5 \text{ eV}/\text{Å}$ —indicating that the configuration has reached a saddle point—or until the lowest eigenvalue become positive—indicating that the trajectory selected is back in the initial harmonic well region and must be rejected. After reaching the saddle point, the configuration is pushed slightly away from it, and is relaxed into a new local energy minimum, called the “final minimum.” This event is stored and a new event is started from the same initial minimum.

B. Properties of the initial configurations

The two initial configurations used here were prepared using ART nouveau and mSW. Starting from a 1000-atom randomly packed unit cell with periodic boundary conditions, ART events were applied until the configurational energy equilibrated. We used a Metropolis accept/reject algorithm, based the energy difference between consecutive local minima, as described in Refs. 18 and 21. The first configuration has an energy per atom of -4.000 eV , compared to -4.137 eV for the crystal, with 20 fivefold and 26 threefold defects. The radial distribution function is in good agreement with recent experimental data (Fig. 1). The model is there-

fore of comparable quality to the models discussed in Refs. 19 and 22. This configuration is used as the origin of all events for the first 42 000-event run.

The second configuration was prepared by further applying ART on the initial configuration, at a Metropolis temperature $T=0.25 \text{ eV}$, for a few thousands of events, until the average displacement per atom reached 1 Å. We generated more than 7000 events around this second minimum in order to ensure that the features of the energy landscape were statistically independent of the minimum selected. All the quantities analyzed here are the same for both minimum, confirming that the properties of the landscape are not affected by the details of the topology. In view of the similarities between the properties of these two sets of data, most of the discussion will focus on minimum 1.

C. Search for activated events

All events presented here are characterized by three configurations: the initial minimum (either minimum 1 or 2), a first-order saddle point, and a second (final) minimum, obtained by relaxing the configurational energy from the saddle point, in a direction opposite to that of the initial minimum. As discussed in Ref. 7, these events are reversible both from the saddle point and the final minimum.

Each event is generated starting from the same initial minimum, but with a different random initial direction, the rest of the procedure being deterministic. It takes about 400 force evaluations to generate a single event. About one third of the random launches lead to a trajectory that brings the configuration back in the harmonic basin; these trajectories are simply rejected and a new random direction selected. The whole search took about 4 weeks on a single processor of a Regatta Power4 IBM.

III. RESULTS

A. Distributions on events

We first discuss the distributions of events generated from minimum 1. 42 581 events were generated following the procedure discussed in the previous section. The distribution of activated ($E_{\text{saddle}} - E_{\text{initial}}$) and asymmetry ($E_{\text{final}} - E_{\text{initial}}$) energies for all these events are given in Fig. 2. The activated energies form a continuous spectrum, from 0 to more than 7 eV, with an average barrier height of 3.0 eV and a distribution width of 1.2 eV. Both the width and average barrier energies are much lower than those reported earlier using the previous version of ART.¹¹ This underlines the limit of the original ART, which cannot converge directly to a saddle point. As shown in Ref. 12, the error on the barrier using the initial version of ART could be as high as 1 eV. As with the method of Munro and Wales,²⁰ on the contrary, ART nouveau makes it possible to identify the transition point to any desired accuracy. Comparing with experiment, we find that the average barrier height is in overall agreement with measurements of Shin and Atwater²⁴ which indicate activation barriers extending from 0.25 eV to about 2.8 eV. The typical activation energies we find correspond also to the isothermal calorimetric data of Roorda *et al.* which indicate high-activation barriers.




FIG. 2. Solid lines: Normalized distribution of energies for the 42 581 events generated around minimum 1. Top: distribution of the activated energy $E_{\text{saddle}} - E_{\text{initial}}$. Bottom: distribution of the asymmetry energy $E_{\text{final}} - E_{\text{initial}}$. Dotted lines: distribution of energies for the 8799 different (unique) saddle points (top) and the 6519 different final configurations.

Although there is no experimental information to compare with, it is useful to examine the distribution of the asymmetry energy, i.e., the energy difference between the final and initial minima (see the bottom panel of Fig. 2). The average of the distribution is at 1.7 eV, with a width of 1.4 eV. There are only a few events with a final energy lower than the initial as should be expected, since the initial configuration is already very well relaxed. As shown below, the distribution around minimum 1 is essentially identical to that around minimum 2. This result suggests that the overall shape of the configurational energy landscape is not sensitive to the details of the configuration, contrary to what one could think.

It is also useful to compare our results with those of Middleton and Wales obtained on the same system using the eigenvector following the BFGS approach (EF-BFGS);¹³ ART nouveau is similar in spirit to this method. In Fig. 3, we plot the activated energy calculated from the *initial* minimum and the *final* ($E_{\text{saddle}} - E_{\text{final}}$) for both sets of events. The barrier distribution calculated from the initial minimum show that the sampling of events differs seriously for both methods; while EF-BFGS seems to favor strongly events with a barrier below 2 eV, ART selects events in a more Gaussian way. Since both methods follow closely the direction corresponding to the lowest eigenvalue to the saddle point, this difference is solely due to the algorithm used to leave the harmonic well. Remarkably, however, this selection has very little impact on the barrier distribution calculated from the *final* minimum; in this case, both methods find the same

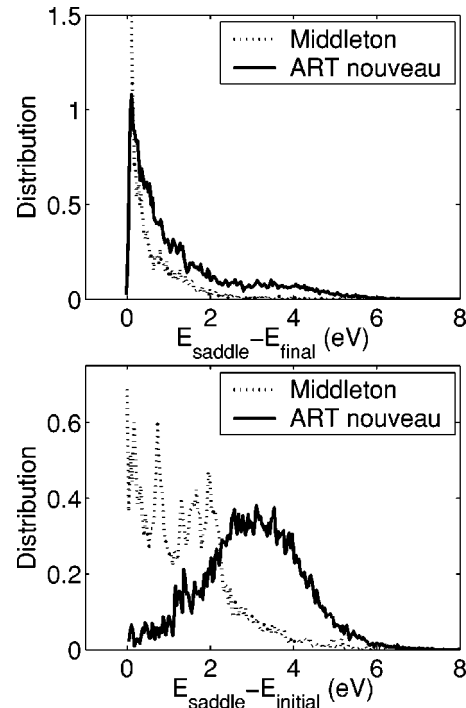


FIG. 3. Comparison between the distributions of activation energy calculated from the final (top) and the initial (bottom) minimum for the set of events computed by Middleton and Wales (Ref. 13) and that generated with ART nouveau from minimum 1.

barrier distribution, which is heavily skewed towards very lower barriers. We discuss the significance of these dissimilarities in Sec. IV.

While the energy distribution inform us about the barriers heights, the distribution of Hamming distances between configurations (Fig. 4) provide some insight as to the rearrangements taking place. As was found previously in *a*-Si and

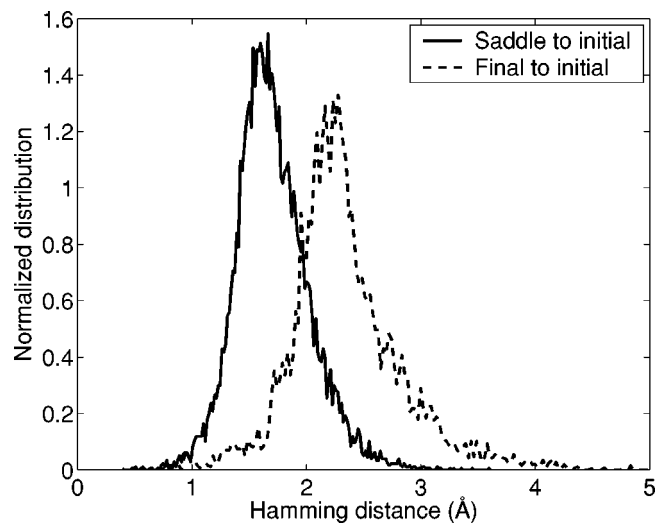


FIG. 4. Distribution of the total displacement (Hamming distance) for the all events starting from minimum 1 at the saddle point (solid line) and at the final minimum (dashed line). The distribution of Hamming displacements for unique events, as defined in the text, is essentially identical to this one.

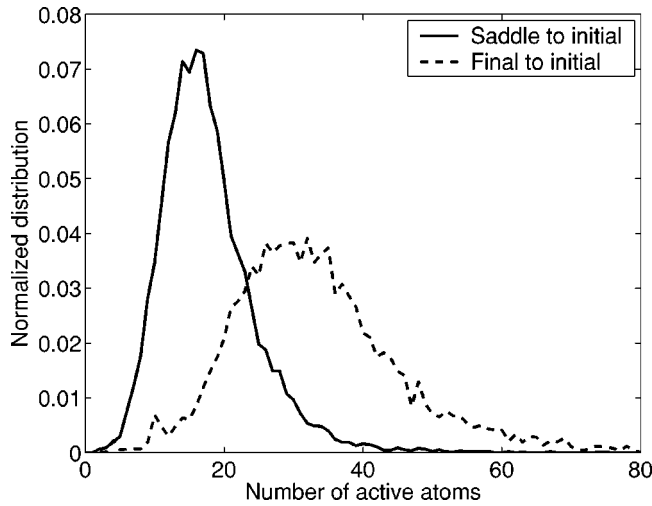


FIG. 5. Distribution of the number of activated atoms at the saddle point (solid line) and the final minimum (dashed line) for all events starting at minimum 1. The distribution is almost identical when only unique events, as defined in the text, are taken into account.

v -SiO₂ most saddle points are somewhere at midway between the initial and final stage.^{11,21} This is somewhat unexpected in the energy landscape picture: in a high-dimensional space, any two randomly selected direction are orthogonal. For truly high-dimensional events, therefore, one would expect to find little correlation between the displacement at the saddle point and that at the final minimum. This is not the case, however, if events are taking place in a much restricted subspace. The distribution of displacements indicates therefore that even though the simulation system is embedded in a 3000-dimensional space, the effective subspace in which each event takes place is much smaller and local in nature; one must therefore be cautious when interpreting the dynamics of a material based solely on the energy landscape picture.

The local nature of the dynamics is also found plotting the number of active atoms participating in each event (Fig. 5). An active atom is defined as one that has moved by more than 0.1 Å from its initial state. As discussed in Refs. 19 and 21, the value of the threshold is chosen to be close to the typical atomic displacement due to thermal vibration at room temperature. We see that most events involve a maximum of 25 atoms at the saddle point, and 40 at the final minimum. If the threshold is increased to 0.4 or 0.5 Å, the number of active events drops typically to between 2 and 4 per event. The relatively narrow distribution of active atoms is not an artifact due to some bias in the sampling of the cell. Figure 6 shows the probability of each atom to participate into an event. The remarkably homogeneous figure confirms that the rearrangements can take place anywhere in the network, with a participation probability varying by at most a factor 3.

B. Properties of the landscape

As the sampling of events is random, the total data set contains redundant events that might bias the analysis of the

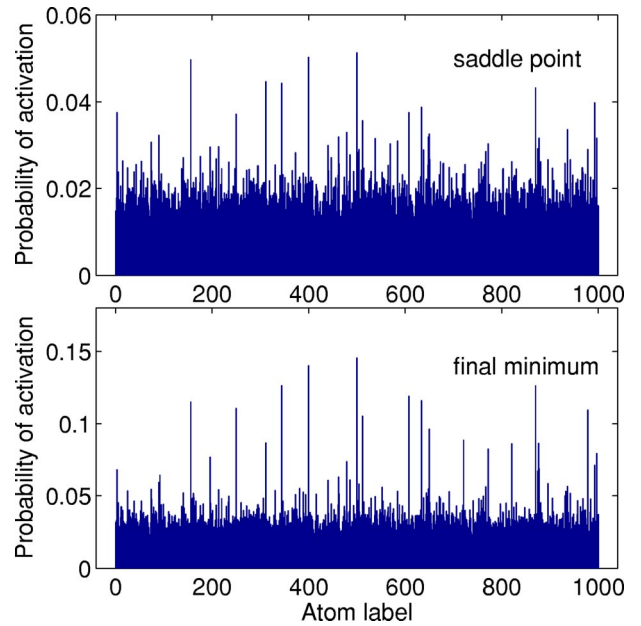


FIG. 6. Probability of participating in a event for each atom of the model, computed using on all generated events. Top: probability of being an active atom at the saddle point; bottom: probability of being an active atom at the final minimum.

landscape. It is therefore useful to look at the properties of “unique” events, i.e., to analyze the set of single copies of all events. We identify each event based on the identity of the participating atoms as well as on the energy at the barrier and at the new minimum. In order to decrease the impact of the finite precision in the convergence at the saddle and the minimum, we define the participating atoms as those moving by at least $\delta r_{\text{threshold}} = 0.4$ and 0.2 Å, respectively, at the saddle point and the final minimum. Two events are considered identical if the active atoms are the same and the energy barrier differs by less than 0.2 eV. We verified that the precise values of the various thresholds do not affect the qualitative results although the precise number of different events obviously depends on the threshold. The number of unique events, as a function of the number of events already sampled, is shown in Fig. 7. We find 6519 different minima, 8799 different saddles, and 11 014 different events generated within the 42 000+ sequence. On average, each event is therefore visited almost four times.

Using the distribution of unique events, it is possible to assess the biases in the ART sampling. Figure 8 shows the ratio of all saddle points or minima generated over the list of all unique ones. As was the case for the Lennard-Jones clusters, a system with a totally different energy landscape, ART sampling seems to select events with an exponential bias, $\exp(-\Delta E/4)$.

These results are obtained using a subset of all events surrounding minimum 1. Even after 42 000 events, the number of different events around minimum is not yet converged. It is possible to obtain some very crude estimate of this number. We first set the full distribution of saddle points identical to that shown in Fig. 7. Supposing a random selection of events, with an exponential bias on the barrier height such as

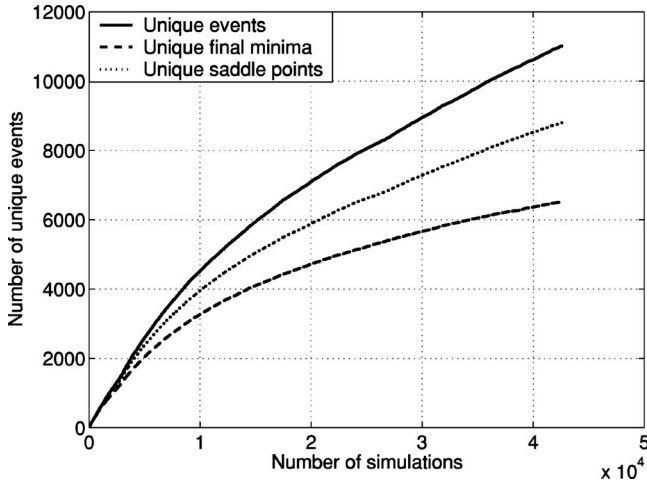


FIG. 7. Solid line: Number of unique events as a function of the number of events sampled. The dashed and the dotted lines show the number of unique saddle points and final minima as a function of the number of events sampled.

that of the previous paragraph, taken from the distribution of unique events, we find that growth of the total of unique events as a function of trial can be reproduced with an exponential bias $\exp(-\Delta E/A)$ with A between 0.40 and 0.60, and a total number of events somewhere between 30 000 and 60 000 or 30–60 different saddle points per atom.




FIG. 8. Top: Log-normal ratio of all saddles points generated from minimum 1 over the unique ones only. A histogram for both distributions is first constructed, as a function of the energy, and the ratio is taken over this histogram. Bottom: same but the configurations at the final minimum. The insets in both figures shows the same distributions plotted linearly.

C. Entropy

Having computed the distribution of activation energy barriers around a minimum, it would be possible to consider associating a time scale with the events. This can be done within the framework provided by the transition state theory.²⁵ For activated mechanisms, the diffusion rate is given by

$$D = \frac{g}{2\alpha} l^2 \nu_0 e^{-\Delta F/k_b T}, \quad (1)$$

where g is the number of equivalent diffusion paths, α the spatial dimension, l the length of the jump, ν_0 a phonon frequency, and $\Delta F = \Delta E - T\Delta S$, the variation in the Helmholtz free energy between the transition state and the initial state.

In order to use the above equation, we need to measure the change in entropy from the minimum to the transition state. We can do that using the harmonic approximation:

$$\Delta S = k_b \ln \left[\frac{\prod_{i=1}^{3N} \nu_i^{(i)}}{\nu_0 \prod_{i=1}^{3N-1} \nu_i^{(s)}} \right], \quad (2)$$

where $\nu_i^{(i)}$ and $\nu_i^{(s)}$ are the real phonon frequencies at the minimum and at the saddle point, respectively. Since there are $3N-1$ real frequencies at the transition state, we replace the imaginary frequency by a typical phonon frequency ν_0 .

Assessing the time scale associated with leaving a given minimum remains a very expensive task even within the harmonic approximation as it would be necessary to compute the entropy associated with each of the 30 000–60 000 saddle points. A simpler approach poses that the entropic prefactor is independent of the specific event.

We can ascertain the validity of such assumption by computing the distribution of the entropic difference from the initial minimum to the saddle point. Due to the high cost of diagonalizing a 3000×3000 matrix, we have applied this procedure to 50 randomly selected events. As seen in Fig. 9, the contribution of the entropy to the activation is of the order of 0.0043 eV/K, with a variation of about 2.5×10^{-4} eV/K. The fluctuations in the entropic contributions to the barrier can therefore affect the attempt frequency by about a factor 25. To a first approximation, it is therefore reasonable to consider that the entropy is only a multiplicative constant in the dynamics of the system.

D. Topological classification

In the previous subsections, we have analyzed the properties of the energy landscape itself. We now turn to the classification of the events generated around a single minimum. There are many properties that could in principle serve for the classification. However, most of them, including the local volume of the event, the change in bond angles and even the asymmetric energy of the event do not show enough structure to be helpful. The topological classification used here is described in detail in Ref. 19. Briefly, it only consid-

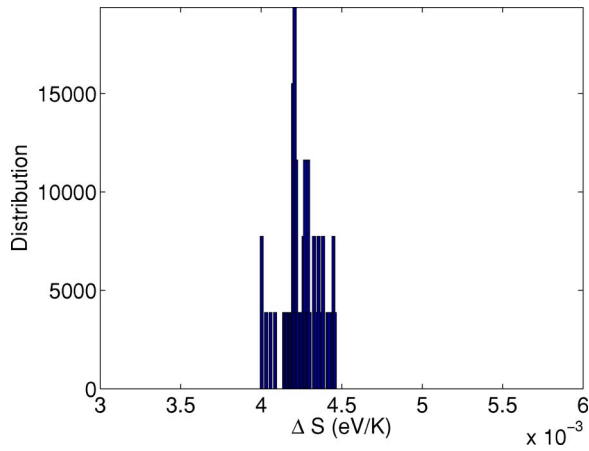


FIG. 9. Distribution of the entropy barrier at the saddle point evaluated in the harmonic limit, for 50 different events selected at random.

ers changes in the coordination of the atoms, obtained through bond breaking and binding. Analyzing the more than 40 000 events generated here, we find more than 2000 different types of topological events. Of those, we discuss below only the events that occurred more than 1% of the time; about 400 times in our sequence. Only five types of event satisfy this criterion.

As in previous simulations, the most common event we find, representing about 20% of all generated events, is the Wooten-Winer-Weaire (WWW) bond exchange mechanism (or *abacbd*, in our notation), introduced almost 20 years ago in the sillium model (see Fig. 10). Since then, this mechanism has been seen both in crystalline and amorphous materials.^{14,26–30}

Figure 11 shows the energy and Hamming distance distribution for this mechanism. Because of the frequency of this mechanism, it might not be surprising to find that both the energy and Hamming distance distributions follow closely those obtained for the whole set of events. These distributions underline once more, nevertheless, the rôle of strain in determining the activation barrier for a given mechanism: the same topological jump can lead to a new configuration with an energy varying by as much as 6 eV.




FIG. 10. A typical WWW-type event. The dark atoms change their coordination during the event, the medium gray move by more than 0.1 Å during the event and the light gray atoms are near neighbors of the atoms involved.

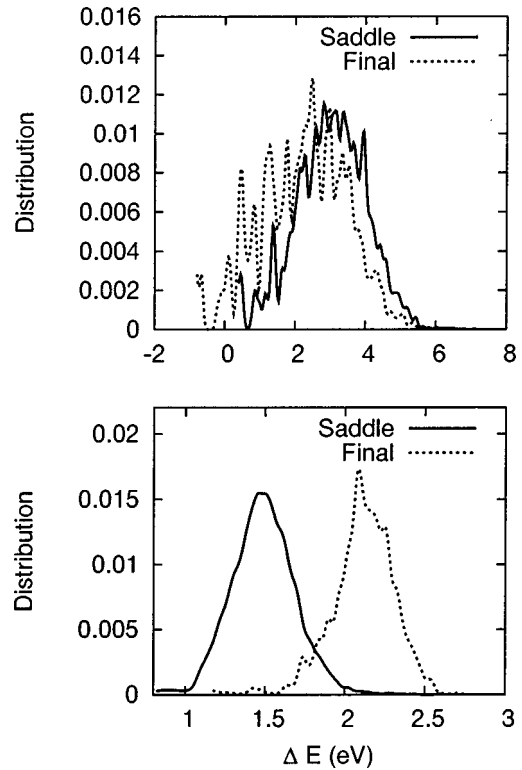


FIG. 11. Top: Energy distribution for the 8000 WWW-type events. Solid line: distribution at the saddle point; dotted line: distribution of the asymmetry energy. Bottom: Distribution of the Hamming distance to the initial minimum at the saddle point (solid line) and the final minimum (dashed line).

The topological nature of the other four dominant events is also shown in Fig. 12. The class *abc* is associated with a bond jumping from one pair of atoms to the next. *ab* represents a bond breaking at the saddle point and reforming. It does not involve any topological change and is responsible for the low-energy peak in the asymmetry energy distribution. As such it is not very interesting. Finally, the two other classes are modifications on the three most common types.

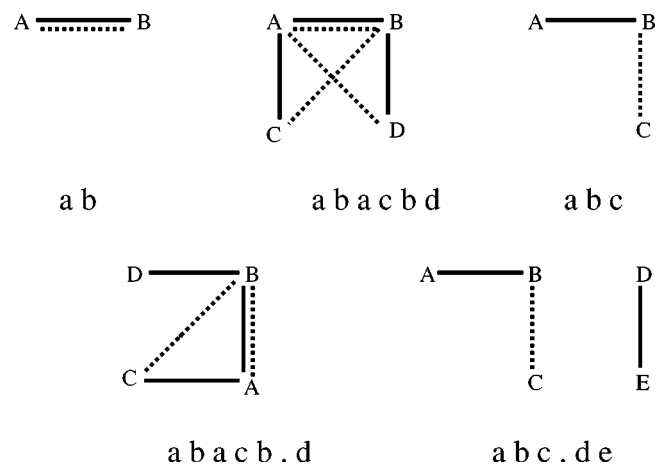


FIG. 12. The five most common types of events, according to our topological classification. The full lines represent a bond present in the initial configuration and the dotted lines a bond present in the final configuration.

TABLE II. Frequency and energies of the most frequent events found in the simulation.

Class	Event Number	ΔE_{saddle}	ΔE_{final}
<i>abacbd</i>	8478	3.04	2.26
<i>abc</i>	759	1.64	1.00
<i>ab</i>	735	3.04	0.21
<i>abacb.d</i>	640	2.47	1.91
<i>abc.de</i>	404	1.64	1.00

We can also analyze the events in terms of change of coordination, effects that could be measured experimentally. During the event, defects are introduced and removed leading mostly, as we have shown above, to metastable states with a rather low escape barrier.

We find about 100 different types of coordination changes as the configuration moves from the initial state to the final minimum (see Table II). Of these, three types of changes account for more than 30% of all events and 15 for more than 75%. The most commons involve no change the overall properties of the network (representing 14% of all events) and the creation of, respectively, one and two threefold defects by the removal of one and two fivefold defects (representing 16% and 12% of all events). As is often the case for these disordered networks, there is no correlation between the type of coordination change and the bond angles or the asymmetry energy; all the subsets that we have examined give the same distributions. We note, however, that the general tendency, as the system is activated from a well-relaxed configuration, is to decrease the average coordination by breaking bonds from four- and fivefold atoms, leading to the creation, on average, of one threefold coordinated atom per event.

IV. DISCUSSION AND CONCLUSIONS

The purpose of this study was to characterize in detail the energy landscape around a single minimum in a well-relaxed model of *a*-Si. Previous work has examined much shorter sequences of events as the configuration relaxed and the atoms diffused. Here, we performed two extensive simulations, always starting from the same two minima. Our results indicate that events are essentially local with a barrier height limited by the cost of breaking one single bond. We also found that there is very little variation in the entropic barrier in spite of the wide spread in the energy part. This is likely related to the fact that the overall shape of the energy landscape is not finely dependent on the specifics of a configuration. Finally, the bias found in ART nouveau indicates that it might be possible to contemplate setting up a kinetic Monte Carlo simulation on this system.

It is interesting to note the strong dissimilarities between the activation barrier distribution as computed from the initial and the final minimum. In particular, the barrier measured from the latter are much lower than those measured from the former configuration and new coordination defects are introduced. This difference can be explained by the fact that the initial configuration selected is very well relaxed; most low-energy barriers, associated with a very unstable direction, would have been crossed during the relaxation, in previous events. These states are likely not contributing to diffusion or the relaxation of the network since they would rapidly relax back into the initial minimum. The physically relevant barrier distribution for defining the evolution of the system is therefore that measured from the initial minimum.

In spite of the 40 000+ events generated, the total number of events around a minimum remains unknown. We can nevertheless estimate that the number of different paths is somewhere between 30 and 60 paths/atom; as shown here, this number should be independent of the size of the network as all events are local in nature.

Finally the stability of the distribution from one minimum to another, as well as the narrow distribution of entropic barriers suggests that it would be possible to develop an accelerated algorithm for *a*-Si, similar to that proposed by Henkelman and Jónsson for the diffusion of Cu on Cu,³¹ and that of Hernandez-Rojas and Wales for LJ glasses.³² This is made even easier by the exponential bias of ART, also seen in LJ clusters. The origin of this bias is not understood but it simplifies significantly the statistical analysis of an ART sampling.

At this point, however, the main limitation of this type of simulation is the absence of detailed experimental results. Although the overall energy scale of the barriers is in agreement with experimental numbers, it is difficult to assess whether or not the theoretical description is correct: there is no information available experimentally on the shape of the barrier distribution. It is currently impossible to generate a better comparison with experiment. Further experimental work is clearly needed to expand on the current available data and help ensure that the advances in numerical results are on solid ground.

ACKNOWLEDGMENT

We thank F. van Wijland for providing us with the solution to the classical occupancy problem. We also thank D. Wales for useful discussions and for providing their original data. F.V. is grateful to the Natural Sciences and Engineering Research Council of Canada (NSERC) for financial assistance. This work was partially supported by NSERC and the Fonds Nature et Technologie du Québec (NATEQ). N.M. received financial assistance from the Research Corporation.

¹R. Cerminski and R. Elber, J. Chem. Phys. **92**, 5580 (1990).

²O.M. Becker and M. Karplus, J. Chem. Phys. **106**, 1495 (1997).

³D.J. Wales, M.A. Miller, and T.R. Walsh, Nature (London) **394**, 758 (1998).

⁴J.P.K. Doye, M.A. Miller, and D.J. Wales, J. Chem. Phys. **110**, 6896 (1999).

⁵J.P.K. Doye, M.A. Miller, and D.J. Wales, J. Chem. Phys. **111**, 8417 (1999).

- ⁶P.N. Mortenson and D.J. Wales, *J. Chem. Phys.* **114**, 6443 (2001).
⁷R. Malek and N. Mousseau, *Phys. Rev. E* **62**, 7723 (2000).
⁸N. Mousseau, P. Derreumaux, G.T. Barkema, and R. Malek, *J. Mol. Graphics Modell.* **19**, 78 (2001).
⁹N.M.G.H. Wei and P. Derreumaux, *J. Chem. Phys.* (to be published).
¹⁰G.T. Barkema and N. Mousseau, *Phys. Rev. Lett.* **81**, 1865 (1998).
¹¹N. Mousseau and G.T. Barkema, *Phys. Rev. B* **61**, 1898 (2000).
¹²Y. Song, R. Malek, and N. Mousseau, *Phys. Rev. B* **62**, 15 680 (2000).
¹³T.F. Middleton and D.J. Wales, *Phys. Rev. B* **64**, 024205 (2001).
¹⁴F.H. Stillinger and T.A. Weber, *Phys. Rev. B* **31**, 5262 (1985).
¹⁵R.L.C. Vink, G.T. Barkema, W.F. van der Weg, and N. Mousseau, *J. Non-Cryst. Solids* **282**, 248 (2001).
¹⁶K. Ding and H.C. Andersen, *Phys. Rev. B* **34**, 6987 (1986).
¹⁷S.J. Cook and P. Clancy, *Phys. Rev. B* **47**, 7686 (1993).
¹⁸G.T. Barkema and N. Mousseau, *Phys. Rev. Lett.* **77**, 4358 (1996).
¹⁹N. Mousseau and G.T. Barkema, *Phys. Rev. E* **57**, 2419 (1998).
²⁰L.J. Munro and D.J. Wales, *Phys. Rev. B* **59**, 3969 (1999).
²¹N. Mousseau, G.T. Barkema, and S.W. de Leeuw, *J. Chem. Phys.* **112**, 960 (2000).
²²G.T. Barkema and N. Mousseau, *Phys. Rev. B* **62**, 4985 (2000).
²³K. Laaziri, S. Kycia, S. Roorda, M. Chicoine, J.L. Robertson, J. Wang, and S.C. Moss, *Phys. Rev. Lett.* **82**, 3460 (1999).
²⁴J.H. Shin and H.A. Atwater, *Phys. Rev. B* **48**, 5964 (1993).
²⁵S. Glasstone, K.J. Laidler, and H. Eyring, *Theory of Rate Processes* (McGraw-Hill, New York, 1941).
²⁶F. Wooten, K. Winer, and D. Weaire, *Phys. Rev. Lett.* **54**, 1392 (1985).
²⁷T. Motooka, *Phys. Rev. B* **49**, 16 367 (1996).
²⁸M. Tang, L. Colombo, J. Zhu, and T. Diaz de la Rubia, *Phys. Rev. B* **55**, 14 279 (1997).
²⁹N. Bernstein, M.J. Aziz, and E. Kaxiras, *Phys. Rev. B* **61**, 6696 (2000).
³⁰S. Goedecker, T. Deutsch, and L. Billard, *Phys. Rev. Lett.* **88**, 235501 (2002).
³¹G. Henkelman and H. Jónsson, *J. Chem. Phys.* **115**, 9657 (2001).
³²J. Hernandez-Rojas and D.J. Wales, *J. Non-Cryst. Solids* (to be published).