

Device-size atomistic models of amorphous silicon

R. L. C. Vink* and G. T. Barkema†

Institute for Theoretical Physics, Utrecht University, Leuvenlaan 4, 3584 CE Utrecht, The Netherlands

M. A. Stijnman‡ and R. H. Bisseling§

Mathematical Institute, Utrecht University, P.O. Box 80010, 3508 TA Utrecht, The Netherlands

(Received 16 July 2001; published 10 December 2001)

The atomic structure of amorphous materials is believed to be well described by the continuous-random-network model. We present an algorithm for the generation of large, high-quality continuous random networks. The algorithm is a variation of the *sillium* approach introduced by Wooten, Winer, and Weaire [Phys. Rev. Lett. **54**, 1392 (1985)]. By employing local relaxation techniques, local atomic rearrangements can be tried that scale almost independently of system size. This scaling property of the algorithm paves the way for the generation of realistic device-size atomic networks.

DOI: 10.1103/PhysRevB.64.245214

PACS number(s): 61.43.Dg, 71.55.Jv, 61.20.Ja

I. INTRODUCTION

The structure of amorphous semiconductors is believed to be well represented by the continuous-random-network (CRN) model introduced by Zachariasen more than sixty years ago.¹ As a result, the generation of high-quality CRNs has been the subject of investigation for many years. The first CRNs were built by hand, see for instance the work of Polk.² Nowadays, the generation of CRNs is mostly carried out on computers.

The first computer-generated networks, which date back to the sixties and seventies, typically contain a few hundred particles. More advanced algorithms and faster computers have increased the size of the networks that can be handled to a few thousand atoms, with simulation cells of up to $40 \times 40 \times 40 \text{ \AA}$.³ As the simulation cells increase in size, actual devices have decreased in size. For example, the thickness of solar cells based on amorphous silicon has already decreased to 1000 \AA ; and because in-plane periodicity after approximately 30 \AA is expected to be a good approximation of the macroscopic lateral size, a reasonable solar-cell model would require a simulation cell of $30 \times 30 \times 1000 \text{ \AA}^3$, containing approximately 45 000 atoms. This is only one order of magnitude larger than what is currently feasible. For other electronic devices, lithography on $0.1 \mu\text{m}$ ($= 1000 \text{ \AA}$) technology is expected to be reached in the coming decade.

In this work, we present a computational approach to generating large CRNs, and discuss the properties of high-quality networks containing up to 20 000 particles. This achievement shows that the generation of device-size atomic-configuration networks is within reach. Our algorithm is similar in spirit to the algorithm of Wooten, Winer, and Weaire (WWW), which has been the basis of the best CRNs generated to date. Another method to generate CRNs is reverse Monte Carlo, see for example Ref. 3 for results of a CRN consisting of 1728 atoms.

We begin by describing the algorithm of Wooten, Winer, and Weaire. We then move on to describe a number of improvements made to the original WWW algorithm by Barkema and Mousseau in 2000.⁴ These improvements accelerate the relaxation by two orders of magnitude or more.

Both the original and the improved WWW algorithm, however, scale poorly with system size, since the computational effort per attempted local atomic rearrangement increases linearly with system size. In this work, we introduce local force and energy evaluations and improve the scaling of computation time with system size N significantly, namely, to a constant per attempted move plus $O(N)$ per accepted move. We also demonstrate how parallel processing can be used to realize an additional speedup, with parallel efficiencies of over 50%. The significance of these improvements is demonstrated by generating 10 000-atom and 20 000-atom CRNs. We then discuss the structural and electronic properties of these models and conclude with an outlook on future research, aiming towards the generation of device-size atomic networks.

II. THE WWW ALGORITHM

In 1985, Wooten, Winer, and Weaire presented an algorithm for the generation of fourfold coordinated CRNs.⁵ In their approach, a configuration consists of the coordinates of N atoms and a list of the $2N$ bonds between them. The structural evolution consists of a sequence of bond transpositions as illustrated in Fig. 1.

Within the original WWW approach, the generation of a CRN starts with a cubic diamond structure that is random-




FIG. 1. Diagram depicting the WWW bond transposition. Four atoms A, B, C, and D are selected following the geometry shown left; two bonds, AB and CD , are then broken and atoms A and D are reassigned to C and B, respectively, creating two new bonds, AC and BD , resulting in the geometry shown right.

ized by a large number of such bond transpositions. After randomization, the network is relaxed through a sequence of bond transpositions, accepted with the Metropolis acceptance probability⁶

$$P = \min \left[1, \exp \left(\frac{E_b - E_f}{k_B T} \right) \right], \quad (1)$$

where k_B is the Boltzmann constant, T is the temperature, and E_b and E_f are the total quenched energies of the system before and after the proposed bond transposition.

With an explicit list of neighbors, it is possible to use a simple interaction such as the Keating potential⁷ to calculate energy and forces

$$E = \frac{3}{16} \frac{\alpha}{d^2} \sum_{\langle ij \rangle} (\vec{r}_{ij} \cdot \vec{r}_{ij} - d^2)^2 + \frac{3}{8} \frac{\beta}{d^2} \sum_{\langle jik \rangle} \left(\vec{r}_{ij} \cdot \vec{r}_{ik} + \frac{1}{3} d^2 \right)^2, \quad (2)$$

where α and β are the bond-stretching and bond-bending force constants, and $d = 2.35$ Å is the Si-Si strain-free equilibrium bond length in the diamond structure. Usual values are $\alpha = 2.965$ eV Å⁻² and $\beta = 0.285\alpha$.

With the approach described above, Wooten and Weaire generated a 216-atom model with an angular distribution as low as 109°. A decade later, using the same approach but more computing power, Djordjević, Thorpe, and Wooten produced two large 4096-atom networks of even better quality, with a bond-angle distribution of 11.02° for configurations without four-membered rings and 10.51° when these rings are allowed.⁹

III. THE IMPROVED WWW ALGORITHM

The WWW algorithm in its original form is capable of producing high-quality amorphous networks containing of the order of a thousand atoms; it is not well suited to generate much larger networks. This is mostly due to the fact that for each proposed bond transposition, about one hundred energy and force calculations are required, each scaling as $O(N)$ with system size N . These $O(N)$ operations are the bottleneck of the algorithm.

In 2000, Barkema and Mousseau (BM) presented a number of modifications to the original WWW algorithm, partially aimed at resolving these poor scaling properties.⁴ Their modifications are summarized below:

1. Starting point for the relaxation in this case is a truly random configuration whereby the atoms are placed at random locations in a periodic box at the crystalline density. This guarantees that the resulting network is not contaminated by some memory of the crystalline state.

2. After a bond transposition in the original WWW approach, the structure is always completely quenched, i.e., all the atomic coordinates are fully relaxed. After the quench, the bond transposition is either accepted or rejected based on the Metropolis probability. In contrast, BM determines a threshold energy before quenching. During the quench the final quenched energy is continuously estimated. Relaxation is stopped when it becomes clear that the threshold energy cannot be reached so that the bond transposition will even-

tually have to be rejected. This leads to a large reduction in the number of force evaluations associated with rejected bond transpositions.

3. A local-relaxation procedure is used whenever possible. Immediately after a bond transposition, only a small cluster of atoms in the model experiences a significant force. This cluster consists of the atoms directly involved in the bond transposition (marked A , B , C , and D in Fig. 1) and of nearby atoms, typically up to the fourth neighbor shell of the four transposition atoms. The number of atoms in such a cluster is about 80. It, therefore, suffices to calculate the force *locally* (i.e., only for the 80 or so atoms inside the cluster) rather than *globally* (i.e., for all the atoms in the model).

Calculating the force on a cluster of atoms is an $O(1)$ operation, which means that it is independent of the total system size. Local force calculations are, therefore, much cheaper than global $O(N)$ force calculations. By using a local-relaxation scheme BM increased the efficiency of the algorithm significantly. Still, to make the final accept/reject decision on the proposed move, the total Keating energy of the system has to be calculated, which is again an $O(N)$ operation. In practice, a switch must be made from local to global relaxation, usually after about ten local relaxation steps.

4. The zero-temperature case is treated specifically.

Using the improved WWW algorithm, Barkema and Mousseau generated two 1000-atom models with bond-angle deviations as low as 9.20°. Furthermore, using the same algorithm they generated a 4096-atom model with an angular deviation of 9.89°. Also, Nakhmanson *et al.* generated paracrystalline models.¹⁰ All models show structural and electronic properties in excellent agreement with experiments.

IV. A SCALABLE WWW ALGORITHM

While the improved WWW algorithm can successfully generate networks containing several thousand particles, it does not deal well with systems of 10 000 particles or more. Each attempted bond transposition still requires one or more $O(N)$ operations. In this section, we present an algorithm for attempting bond transpositions that is completely local, i.e., free of $O(N)$ operations for unsuccessful bond transpositions.

A. Local energy and force evaluations

To exploit the local nature of the bond transpositions, we need to introduce the concept of local energy: we assign to each atom i an energy ϵ_i such that $E = \sum_{i=1}^N \epsilon_i$ with E the Keating energy of the system given by Eq. (2). One way to achieve this is to divide the energy due to two-body interactions equally between the two participating atoms and to assign the energy of three-body interactions to the central atom of the corresponding triple. Thus, we obtain

$$\epsilon_i \equiv \sum_{j=1}^4 \left[\frac{1}{2} \frac{3}{16} \frac{\alpha}{d^2} (\vec{r}_j \cdot \vec{r}_j - d^2)^2 + \sum_{k=j+1}^4 \frac{3}{8} \frac{\beta}{d^2} \left(\vec{r}_j \cdot \vec{r}_k + \frac{1}{3} d^2 \right)^2 \right]. \quad (3)$$

Here, the constants α , β , and d are defined as in Eq. (2); \vec{r}_j represents the vector pointing in the direction of the j th bond away from atom i . The energy E_C of a cluster C of atoms can now be calculated using

$$E_C = \sum_{i \in C} \epsilon_i. \quad (4)$$

The force on the atoms inside the cluster is obtained from the derivative of the Keating energy with respect to the atomic coordinates. Care has to be taken for atoms on the edge of the cluster since these atoms also interact with atoms outside the cluster: due to the two- and three-body terms in the Keating potential, all atoms interact with their first and second nearest neighbors; for atoms located on the edge of the cluster, some of these neighbors are outside the cluster.

B. Local WWW moves

Starting point is a random configuration generated using the method described in Ref. 4. This guarantees that the resulting configurations are not contaminated by some memory of the crystalline state. Assuming that the total Keating energy of the initial configuration is known and equals E , WWW moves can be attempted locally as follows:

1. A threshold energy E_t is determined by using the equation

$$E_t = E - k_B T \ln(1 - r), \quad (5)$$

where r is a random number uniformly drawn from the interval $[0,1)$. The move is accepted if the attempted bond transposition leads to a configurational energy below the threshold energy; otherwise it is rejected.

2. The four atoms involved in the attempted bond transposition and all atoms up to the fourth neighbor shells of these four atoms are grouped into a cluster. These atoms, about 80, are allowed to move. Also the atoms in the fifth neighbor shell are included in the cluster since their local energies may change. These atoms, about 70, are not allowed to move.

3. A list is constructed of all the bonds that contribute to the force on the atoms inside the cluster. As was explained above, some of these bonds involve atoms outside the cluster. For each bond we store the labels of the two atoms constituting the bond, the x , y , and z components of the bond vector (taking care of the periodic boundary conditions), and the square of the bond length. We then calculate the cluster energy E_C using Eq. (4). In the calculation of the cluster energy most bonds are encountered more than once. To increase efficiency, a bond (i.e., its set of three components) is calculated only once during an energy or force evaluation; once a bond has been calculated it is time stamped with an integer flag and the bond information is stored. Later refer-

ences to the same bond are then retrieved from memory. We also store the energy of the atoms that remain outside the cluster: $E_r \equiv E - E_C$. We then perform the bond transposition to obtain the geometry shown in the right frame of Fig. 1.

4. The system is relaxed locally, i.e., only atoms inside the fourth shell allowed to move. At each relaxation step we use Eq. (4) to calculate the energy of the cluster E_C and the atomic forces, again making sure each bond is calculated only once, and perform structural relaxation as in the original and improved WWW algorithms. At each relaxation step the total energy of the system is equal to $E = E_r + E_C$. Local relaxation is continued until the energy has converged or until it becomes clear that the threshold energy cannot be reached.

In the local-relaxation procedure above, the computational effort per attempted bond transposition does not grow with the system size. Local relaxation alone, however, is not sufficient and we also have to relax globally to relieve any strain that may have built up between atoms on the edge of a cluster and noncluster atoms. For clusters extending up to the fifth neighbor shell around the atoms directly involved in the bond transposition we find that global relaxation can lower the configurational energy typically by less than 0.1 eV. We, therefore, switch from local to global relaxation when, during local relaxation, the energy comes to within 0.1 eV of the threshold energy. In most cases, this leads to the move being accepted.

V. PARALLEL PROCESSING

We have developed a parallel version of our algorithm in bulk synchronous parallel (BSP) style,¹¹ with alternating phases of computation and communication separated by a barrier synchronization. The parallel algorithm has two main parts, local relaxation and global relaxation.

The local relaxation is done in parallel by letting every processor try a sequence of randomly chosen bond transpositions, until one of the processors finds an acceptable transposition. The processors work independently but synchronize at regular intervals to communicate their success or failure to the others. If several processors succeed, an arbitrary one is chosen as the winner. This approach requires the replication of all the atomic data. Fortunately, the memory storage needed is limited to a few arrays of size N , which usually fits in every processor. Furthermore, this approach requires refreshing the atomic data when the positions change, causing communication, but this only happens after a success, which is a relatively rare event (about once every thousand attempts). Therefore, we replicate the data instead of distributing them, and base our parallel local-relaxation algorithm on replicated data.

A suitable time interval must be chosen between successive synchronizations. If the interval is too short, synchronization time or work-load fluctuations will become significant; if it is too long, an accepted move will be found in almost every time interval and much time will be wasted: the work done by other processors after a move has been accepted is unnecessary. We set the time interval on empirical grounds at 50 relaxation steps performed during bond-

TABLE I. Energetic and structural properties of models relaxed with the Keating potential. The first two models, DTW4096a and DTW4096b, are the 4096-atom models prepared in Ref. 9 and refer, respectively, to a model with and without four-membered rings. Configurations BM1000a and BM1000b are 1000-atom configurations prepared by Barkema and Mousseau using the improved WWW algorithm and BM4096 is a 4096-atom model prepared in the same way.⁴ Configurations ‘‘10k’’ and ‘‘20k’’ represent, respectively, 10 000-atom and 20 000-atom models prepared using the scalable WWW algorithm described in the text. The ring statistics are for irreducible rings and ρ_0 is based on $d=2.35$ Å.

	DTW4096a	DTW4096b	BM1000a	BM1000b	BM4096	10k	20k
E (eV/atom)	0.336	0.367	0.267	0.264	0.304	0.301	0.286
ρ/ρ_0	1.000	1.000	1.043	1.040	1.051	1.054	1.042
$\langle r \rangle/d$	0.996	0.997	0.982	0.982	0.980	0.980	0.981
$\langle \theta \rangle$	109.24	109.25	109.30	109.27	109.28	109.28	109.25
$\Delta\theta$	10.51	11.02	9.21	9.20	9.89	9.88	9.63
rings/atom							
4	0.015	0.000	0.000	0.000	0.000	0.000	0.020
5	0.491	0.523	0.472	0.480	0.490	0.480	0.456
6	0.698	0.676	0.761	0.750	0.739	0.742	0.759
7	0.484	0.462	0.507	0.515	0.467	0.512	0.501
8	0.156	0.164	0.125	0.116	0.148	0.142	0.149
9			0.034	0.033	0.035	0.034	0.039

transposition attempts. (Each relaxation step requires about 10^5 floating point operations.) For a large number of processors, this choice becomes critical.

The global relaxation is done in parallel by partitioning the simulation cell over the p processors of the parallel computer and letting every processor compute the energies ϵ_i , forces, and displacements for the atoms in its own part of the cell. In contrast to the local relaxation, it is now justified to have all the processors participate in one relaxation step: the amount of work, $O(N)$, in a relaxation step is much larger and there is no other useful work to do anyway.

Communication arises in the global relaxation because processors need data from other processors concerning atoms near interprocessor boundaries. Thus at the end of an iteration, a processor has to communicate the positional changes of its boundary atoms (i.e., atoms within two bonds from an atom on another processor). To reduce the size of the boundary region, we use three types of partitioning:¹² standard cubic (SC), which splits the simulation cell into $p=k^3$ sub-cubes; body-centred cubic (BCC), which splits it into $p=2k^3$ truncated octahedra centered at the sites of the BCC lattice; and face-centered cubic (FCC), which splits the cell into $p=4k^3$ rhombic dodecahedra centered at the sites of the FCC lattice. The BCC and FCC partitionings generate about 10% less communication than the commonly used SC partitioning. With these three partitionings, we can choose from a wide range of processor numbers p , including all powers of two.

We have implemented the parallel algorithm using the BSPlib communications library¹³ on a Cray T3E (with 300 MHz Dec Alpha 21164 processors). In the local relaxation, we have achieved a speedup of 18.6 on 32 processors for the 20 000-atom model, i.e., an efficiency of 58%. The efficiency loss is mainly due to work load fluctuations and to unnecessary work at the end of successful intervals. In the global

relaxation, we have achieved a speedup of 19.3, i.e., an efficiency of 60%, see Ref. 12 for more details. Here, the efficiency loss is mainly due to redundant computations for boundary regions and communication of boundary data. The total speedup of our parallel version depends on the mixture of local and global relaxations needed. This mixture is influenced by a variety of parameters such as the temperature T in the Metropolis acceptance criterion (1) and the expected energy reduction due to the global relaxation. In our simulations, the amounts of CPU time spent on local and global relaxation were nearly equal.

The parallel and sequential versions of our program are interchangeable: both can be used to relax a given CRN, or to create a new one. The choice between them can be made depending on the computer available. A rough estimate of the total computation time needed for the 10 000-atom model is 20 processor weeks on a 300 MHz Dec Alpha 21164 processor; for the 20 000-atom model, 7 processor weeks on a 667 MHz Dec Alpha 21264 processor. Using the parallel program reduces the elapsed time, although by less than a factor p because of efficiency loss.

VI. RESULTS

Using the scalable WWW algorithm we have generated one 10 000-atom amorphous silicon network and one 20 000-atom network. In this section, we discuss the structural and electronic properties of these networks. In Table I, we compare our configurations relaxed with the Keating potential with those of Djordjević, Thorpe, and Wooten (DTW) (Ref. 9) and with models generated by Barkema and Mousseau using the improved WWW algorithm.⁴ We also provide irreducible ring statistics.

Table I shows that the strain per atom for the 10 000-atom and 20 000-atom models is significantly lower than that of

TABLE II. Structural properties of configurations after relaxation with the modified Stillinger-Weber (mSW) potential. The total ring number per atom (including reducible rings) is also reported, as well as the energy after relaxation with the original Stillinger-Weber (SW) potential.

	BM1000a	BM1000b	BM4096	10k	20k
E (eV/atom, mSW)	-4.026	-4.034	-3.990	-3.994	-4.008
E (eV/atom, SW)	-4.126	-4.133	-4.106	-4.109	-4.116
ρ/ρ_0	0.947	0.950	0.936	0.938	0.933
$\langle r \rangle/d$	1.018	1.017	1.020	1.021	1.020
$\langle \theta \rangle$	109.25	109.24	109.20	109.19	109.20
$\Delta \theta$	9.77	9.70	10.51	10.54	10.18
rings/atom					
4	0.000	0.000	0.001	0.003	0.020
5	0.472	0.480	0.489	0.481	0.456
6	0.840	0.847	0.830	0.844	0.843
7	1.011	1.023	0.979	1.034	1.020
8	2.025	2.002	2.064	2.038	2.018

the DTW models. Compared to the 1000-atom models prepared with the improved WWW algorithm (BM1000a and BM1000b) we find that the strain per atom in our 10 000-atom and 20 000-atom models is only slightly higher, thus clearly demonstrating the efficiency of the scalable WWW approach.

An important quantity that can be compared with experiment is the width of the bond-angle distribution $\Delta \theta$. Experimentally, this quantity can be extracted from the radial

distribution function (RDF) (Ref. 14) or the Raman spectrum.^{15,16} The most recent measurement, obtained from the RDF, yields 10.45° for as-implanted samples and 9.63° for annealed samples.¹⁴ The bond-angle distributions of the 10 000-atom and 20 000-atom models generated by us are in good agreement with these experimental values.

Although the Keating potential already produces high-quality networks by itself, it is important to check the stability of these networks when relaxed with a more realistic interaction potential that does not require a pre-set list of

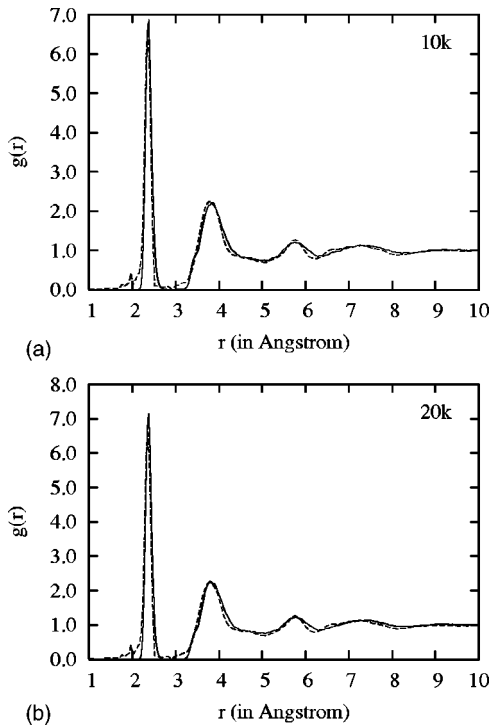


FIG. 2. Pair-correlation function for the 10 000-atom model (top) and the 20 000-atom model (bottom) after relaxation with the modified Stillinger-Weber potential (solid line). The dashed line shows the experimental result from Ref. 14; distances are in Å.




FIG. 3. Electronic density of states for the 10 000-atom model (top) and the 20 000-atom model (bottom) as obtained from *ab initio* tight binding.²³

neighbors. For this purpose we use the Stillinger-Weber (SW) potential¹⁷ but with an enhanced angular force: the three-body term is increased by 50% with respect to the two-body term. This *ad hoc* modification was shown to produce good structural properties for amorphous silicon.^{18–22}

The properties of the networks after relaxation with the (modified) SW potential are reported in Table II. For all configurations, the bond-angle distribution widens and the density decreases.

Figure 2 shows the pair-correlation function for the 10 000-atom and 20 000-atom models compared to the experimental pair-correlation function obtained by Laaziri *et al.* on annealed *a*-Si samples prepared by ion bombardment.¹⁴ Agreement is excellent. However, configurations differing widely in topology can easily produce similar pair-correlation functions. Agreement with the experimental pair-correlation function must, therefore, be regarded as a minimum demand on a high-quality CRN.

A more stringent criterion that can be used to evaluate the quality of a model is the coordination number of the atoms. Using the minimum of the pair-correlation function between the first and second neighbor peak as the nearest-neighbor cut-off distance (at $r \approx 3.10$ Å) and after relaxation with the modified Stillinger-Weber potential, the 10 000-atom model develops 1.08% of atoms with three or five neighbors. For the 20 000-atom model this percentage is 0.24%. The ratio of the number of threefold coordinated atoms to the number of fivefold coordinated atoms is sensitive to the precise value of the nearest-neighbor cut off; a smaller cut-off decreases the number of fivefolds, while increasing the number of threefolds. With the cut-off used here, the 10 000-atom model contains 0.13% of threefold coordinated atoms and 0.95% of fivefold coordinated atoms. For the 20 000-atom model these percentages are 0.07% and 0.17%, respectively.

While structural averages provide good insight into the overall quality of a model, they do not say much regarding local environments. It is, therefore, also important to look at the electronic properties of our models: even small densities of highly strained geometries or defects will be picked up as states in the gap of the electronic density of states (EDOS). In Fig. 3 we show the EDOS of the 10 000-atom and 20 000-atom models. The FIREBALL local-basis *ab initio* code²³ was used to obtain the EDOS. A remarkable feature of the state densities shown here is the absence of states in the gap, leading to a perfect gap of 1.3 eV for both models.

VII. SUMMARY AND CONCLUSION

We have presented here a scalable version of the WWW algorithm that allows for local atomic rearrangements to be tried using only $O(1)$ operations. We have developed an efficient parallel version that achieves good load balance and limits communication. The scalable performance of the algorithm has been demonstrated by generating one 10 000-atom and one 20 000-atom model. Structural and electronic properties of these models are excellent and they compare well to experiments.

These high-quality models have the long-term goal of accurately modeling devices such as solar cells. At this point, using periodic-boundary conditions in the two extended directions, we are able to simulate *a*-Si films with a thickness of about 1000 Å. Once such atomic configurations become available, the role of various structural and electronic defects can be studied.

ACKNOWLEDGMENTS

We thank Dave Drabold for communicating to us the EDOS of the 10 000-atom and 20 000-atom models.

*Electronic address: vink@phys.uu.nl; URL: <http://www.phys.uu.nl/~vink>

†Electronic address: barkema@phys.uu.nl

‡Electronic address: mastijnm@math.uu.nl

§Electronic address: Rob.Bisseling@math.uu.nl; URL: <http://www.math.uu.nl/people/bisseling>

¹W.H. Zachariasen, *J. Am. Chem. Soc.* **54**, 3841 (1932).

²D.E. Polk, *J. Non-Cryst. Solids* **5**, 365 (1971).

³S. Kugler, L. Pusztai, L. Rosta, P. Chieux, and R. Bellissent, *Phys. Rev. B* **48**, 7685 (1993).

⁴G.T. Barkema and N. Mousseau, *Phys. Rev. B* **62**, 4985 (2000).

⁵F. Wooten, K. Winer, and D. Weaire, *Phys. Rev. Lett.* **54**, 1392 (1985).

⁶N. Metropolis, A.W. Rosenbluth, M.N. Rosenbluth, A.H. Teller, and E. Teller, *J. Chem. Phys.* **21**, 1087 (1953).

⁷P.N. Keating, *Phys. Rev.* **145**, 637 (1966).

⁸F. Wooten and D. Weaire, *Solid State Phys.* **40**, 1 (1987).

⁹B.R. Djordjević, M.F. Thorpe, and F. Wooten, *Phys. Rev. B* **52**, 5685 (1995).

¹⁰S. Nakhmanson, P.M. Voyles, N. Mousseau, G.T. Barkema, and D.A. Drabold, *Phys. Rev. B* **63**, 235207 (2001).

¹¹L.G. Valiant, *Commun. ACM* **33**, 103 (1990).

¹²M.A. Stijnman, R.H. Bisseling, and G.T. Barkema (unpublished).

¹³J.M.D. Hill, B. McColl, D.C. Stefanescu, M.W. Goudreau, K. Lang, S.B. Rao, T. Suel, T. Tsantilas, and R.H. Bisseling, *Parallel Comput.* **24**, 1947 (1998).

¹⁴K. Laaziri, S. Kycia, S. Roorda, M. Chicoine, J.L. Robertson, J. Wang, and S.C. Moss, *Phys. Rev. B* **60**, 13 520 (1999).

¹⁵D. Beeman, R. Tsu, and M.F. Thorpe, *Phys. Rev. B* **32**, 874 (1985).

¹⁶R.L.C. Vink, G.T. Barkema, and W.F. van der Weg, *Phys. Rev. B* **63**, 115210 (2001).

¹⁷F.H. Stillinger and T.A. Weber, *Phys. Rev. B* **31**, 5262 (1985).

¹⁸K. Ding and H.C. Andersen, *Phys. Rev. B* **34**, 6987 (1986).

¹⁹G.T. Barkema and N. Mousseau, *Phys. Rev. Lett.* **77**, 4358 (1996); N. Mousseau and G.T. Barkema, *Phys. Rev. E* **57**, 2419 (1998).

²⁰J.M. Holender and G.J. Morgan, *J. Phys.: Condens. Matter* **3**, 7241 (1991).

²¹N. Mousseau and G.T. Barkema, *Phys. Rev. B* **61**, 1898 (2000).

²²R.L.C. Vink, G.T. Barkema, W.F. van der Weg, and N. Mousseau, *J. Non-Cryst. Solids* **282**, 248 (2001).

²³O.F. Sankey and D.J. Niklewski, *Phys. Rev. B* **40**, 3979 (1989); O.F. Sankey, D.A. Drabold, and G.B. Adams, *Bull. Am. Phys. Soc.* **36**, 924 (1991).