Pseudorandom number generator for massively parallel molecular-dynamics simulations

Brad Lee Holian
Theoretical Division, Los Alamos National Laboratory, Los Alamos, New Mexico 87545

Ora E. Percus
Courant Institute of Mathematical Science, New York University, 251 Mercer Street, New York, New York 10012

Tony T. Warnock
Theoretical Division, Los Alamos National Laboratory, Los Alamos, New Mexico 87545

Paula A. Whitlock
Computer and Information Sciences Department, Brooklyn College, 2900 Bedford Avenue, Brooklyn, New York 11210
(Received 22 November 1993)

A class of uniform pseudorandom number generators is proposed for modeling and simulations on massively parallel computers. The algorithm is simple, nonrecursive, and is easily transported to serial or vector computers. We have tested the procedure for uniformity, independence, and correlations by several methods. Related, less complex sequences passed some of these tests well enough; however, inadequacies were revealed by tests for correlations and in an interesting application, namely, annealing from an initial lattice that is mechanically unstable. In the latter case, initial velocities chosen by a random number generator that is not sufficiently random lead quickly to unphysical regularity in grain structure. The new class of generators passes this dynamical diagnostic for unwanted correlations.

PACS number(s): 02.70.−c, 03.20.+i, 05.40.+j, 64.70.Dv

I. INTRODUCTION

The use of pseudorandom number generators in molecular dynamics (MD), where the Newtonian equations of motion of millions of atoms can be integrated on modern massively parallel computers, is restricted mainly to the initialization of the many-body trajectory. N atoms are placed at, or near, perfect lattice sites and typically given velocities selected from a Maxwellian-Boltzmann distribution (Gaussian). This can be achieved by selecting a pair of random numbers from a uniform distribution and then applying the Box-Muller transformation [1] to obtain two Gaussian-distributed random numbers. The problem of initializing the velocities thus reduces to generating a sequence of uniformly distributed random numbers. Stated more precisely, one needs to generate on the computer "a nonrandom, deterministic sequence of numbers $x_1, x_2, \ldots$, which is supposed to resemble a sequence of independent, random samples from the uniform probability distribution on the interval $0 \le x_i < 1$" [2]. A number of such procedures are described in the excellent book, *Numerical Recipes* [3]. Most of these methods are recursive, employing the relationship [4(a)],

$$x_{i+1} = \{Mx_i + \alpha\} , \tag{1}$$

where $\{y\}$ is the fractional part of y, $M > 1$ is an integer, and α is an "irrational" number in the interval (0,1). (Since real numbers are represented on a computer by a finite number of bits, i.e., *machine precision,* the term *irrational* must be taken with a grain of salt. Also, the proper choice of M and α is an arcane subject in and of itself [3,4(a)]. Recursive procedures are fine for serial computers, but are not very convenient for massively parallel machines. If we associate a different random number with each parallel processor, which in simplified terms could represent an individual atom, then ideally, a parallel pseudorandom number generator should produce a number on each processor simultaneously and independently (leading to a velocity component after transformation). Parallel implementation of Eq. (1) for initialization of velocities would be too time consuming if splitting was used, since all L processors would have to execute Eq. (1) L times (where L might be of the order 10^6), discarding all numbers $x_j, j > i$, where i is the processor number. An alternative would be to use the method proposed by Frederickson *et al.* [5] and further developed by Percus and Kalos [6]. In the latter method, each processor would have its own generator and no random numbers are discarded. Once an array of processors has been given initial random numbers, the sequence given in Eq. (1) could indeed be used, for example, in a parallel Monte Carlo (MC) procedure, where now the index $i \to i+1$ would represent the increment in MC "time" (step number) for a given processor (atom).

Why worry about the initialization of an MD trajectory, when Lyapunov instability [7] guarantees that the initial state is supposedly forgotten within a few collision times? The answer is twofold. First, one would like to be able to specify initial conditions on *any* type of computing environment and be able to check, to within machine accuracy, the different MD simulation algorithms against each other. One would hope, moreover, to achieve this objective of randomness for the initial state with a simple and straightforward (nonrecursive) algorithm. Secondly, we shall give an example in this paper where the initial state is *not* forgotten within a few collision times, and show that a generator that is insufficiently random can

lead to unphysical results.

In this paper, we discuss three increasingly complex candidates for uniform random sequences: (i) the Weyl sequence, which is well known to pass the uniformity test but has serious correlations; (ii) the nested Weyl sequence (NWS), which is uniform and again fails the test for correlations as well as the annealing test; and (iii) an alternative sequence, the shuffled nested Weyl sequence (SNWS), which passes all the tests and which we propose as a portable pseudorandom number generator for parallel computing applications. We give examples of annealing from the unstable square lattice in two dimensions (2D) for the latter two sequences, using the grain map of local bond orientation as a visual diagnostic. For each of the sequences, Weyl, NWS, and SNWS, we give sample results from the Monkey Test used to test for correlations.

II. METHODS

Two obvious tests to be made to see if a proposed algorithm for a pseudorandom number generator accomplishes its objectives are (i) computing the moments of the sequence of supposedly uniform numbers on the interval (0,1) and (ii) computing the correlation function to determine whether or not the sequence is independent. If the sequence is composed of N numbers uniform on the interval (0,1), then the mth moment should be given by

$$\langle x^m \rangle = \frac{1}{N} \sum_{i=1}^{N} x_i^m = \frac{1}{m+1} \left[1 + \frac{\chi_N(m)}{N^{1/2}} \right] , \tag{2}$$

where $-1 < \chi_N(m) < 1$ for all N is indicative of a suitably uniform sequence. Likewise, the correlation function of lag k for N independent numbers should be given by

$$\rho_N(k>0) = \frac{1}{N} \sum_{i=1}^{N} (x_i - \langle x \rangle_N)(x_{i+k} - \langle x \rangle_N) = \frac{\theta_N(k)}{N^{1/2}} , \tag{3}$$

where $-1 < \theta_N(k) < 1$ for all N is indicative of a suitably uncorrelated sequence. In these tests, one is limited only by the amount of work one is willing to do, namely, the number of moments and lags to be computed. In addition, a whole battery of empirical tests are available for testing supposedly random sequences [4]. For detecting correlations among non-neighboring terms in the pseudorandom number sequence, we chose Marsaglia's Monkey Test [8]. In Percus and Whitlock [9], the complete mathematical development of the Monkey Test is described. It can be summarized as follows. Given a sequence of length N consisting of m different symbols, let ξ_r be the number of occurrences of a subsequence of r given symbols. Let $E\{\xi_r\}$ be the expected value of ξ_r, then

$$E\{\xi_r\} = \frac{N-r+1}{m^r} . \tag{4}$$

Alternatively, these tests for uniformity and correlation can be represented visually by plotting a random sequence vs its kth lag as pairs of points (x_i, x_{i+k}) on the unit square, $(0,1)\times(0,1)$. If a sequence is uniform and uncorrelated, and therefore truly *pseudorandom*, then the unit square will be uniformly covered with no apparent clustering or striping. Color may also be employed by assigning a rainbow of colors to the processor numbers $i=1,\ldots,L$, with deep blue for $i=1$ to light green for $i=L/2$, going up to bright red for $i=L$; in this case, clustering or patterns of colors should not occur either. We will show that even these more sophisticated visual diagnostics are not always sufficiently reliable for detecting hidden correlations. Furthermore, we will show that annealing from an initial unstable crystal structure can reveal otherwise hidden correlations in the "random" number generator used to initialize velocities.

III. WEYL SEQUENCE

The Weyl sequence [10] is obtained for an arbitrary irrational number α (such as $\{2^{1/2}\} = 0.414\ldots$) by taking the fractional parts,

$$X_n = \{n\alpha\} , \tag{5}$$

for $n=1,2,\ldots,N$. This is clearly an easily parallelizable sequence, and Weyl and others have shown [2,10] that it is uniform on the interval (0,1). Empirically, the first hundred moments show that it is uniform. However, it is not uncorrelated [10(b)]; for each lag ($k=1{:}100$), the correlation function as defined in (3) is a constant independent of N. Furthermore, the Monkey Test for any m shows substantial deviations from Eq. (4) [9]. Such obvious correlations introduced into the initial velocities would be unacceptable in MD simulations. Nevertheless, the Weyl sequence leads naturally to the next level of complexity, the nested Weyl sequence.

IV. NESTED WEYL SEQUENCE

The nested Weyl sequence (NWS), a natural extension of the Weyl sequence, is defined for $n=1,\ldots,N$ as

$$Y_n = \{nX_n\} = \{n\{n\alpha\}\} . \tag{6}$$

The first hundred moments of the NWS show that it is uniform, and the first hundred lags of its correlation function reveal no unwanted correlations. The lag $k=1$ correlation plot, with $\alpha = \{2^{1/2}\}$, is plotted as a unit square in Fig. 1. No obvious correlations appear in Fig. 1. When one views this plot on edge, one might imagine seeing some very subtle herringbones, but there is no pattern that is particularly striking.

On a Connection Machine (CM-2), we simulated two-dimensional (2D) annealing of a polycrystalline Lennard-Jones spline [11] (LJ spline) system using MD. The NWS was used as uniformly distributed input to the Box-Muller [1] transformation to obtain the initial Maxwell-Boltzmann velocities. On the CM-2, atoms are laid out on an Eulerian grid, that is, a fixed rectangular parallelepiped array of boxes, with a maximum number of pigeonholes provided in each box for occupancy by the particles. Motivation for this parallel data structure comes from the fact that the serial communications among pigeonholes within a given Eulerian box is very rapid, while neighboring communications among boxes is somewhat

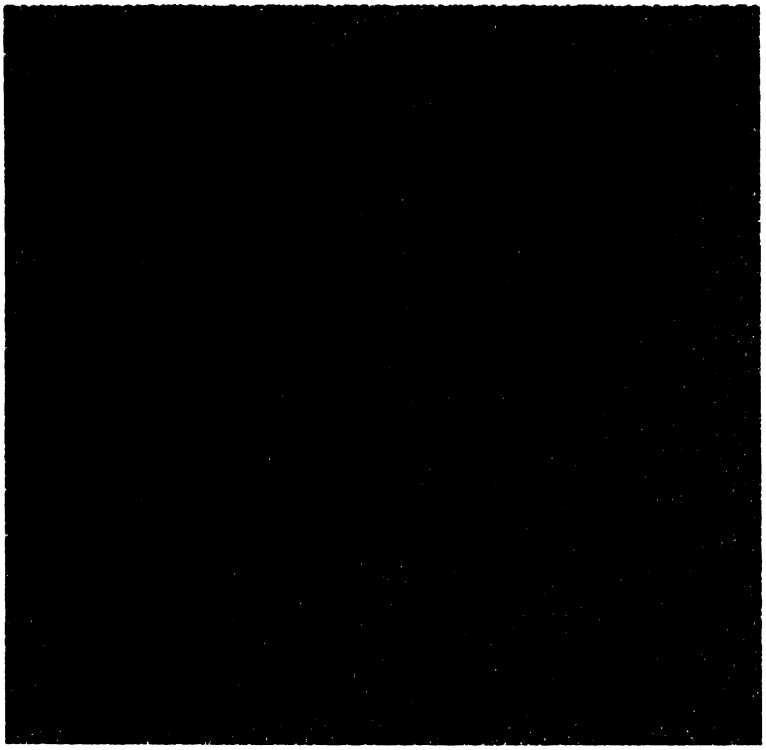

FIG. 1. Lag $k=1$ correlation plot (unit square) for the NWS [Eq. (6)], generated from the seed $\alpha=\{2^{1/2}\}$. The order of the pairs (Y_i, Y_{i+1}) is colored by a rainbow, where deep blue is $i=1$, light green is near $i=32\,768$, up to bright red for $i=L=65\,536$.

FIG. 2. Grain map [Eqs. (7) and (8)] of 262 144 Lennard-Jones spline atoms in two dimensions (about 0.1 μm on a side) at a time of about one and one-half vibrational periods after initialization from a square lattice (density appropriate to zero pressure and one-fourth of the melting temperature), with initial x and y velocities chosen from Box-Muller transformation of two NWS, using $\{2^{1/2}\}$ and $\{3^{1/2}\}$ as seeds. The grain-map rainbow corresponds to deep blue for crystallites with triangular-lattice bond angles (with x axis) near $-30°$, to light green near 0°, up to bright red for $+30°$. Note the unphysical regularity of diagonal stripes of crystallites.

less so, though considerably faster than general router communications. The size of the boxes is such that only nearest-neighbor communication among Eulerian boxes is necessary in order to determine forces between a given atom and another that falls within the range of interaction. (For the LJ spline potential, depending on the density, one and occasionally two neighbor shells are included in 2D.) In general, the data—coordinates, momenta, forces, etc.—are laid out as $(d+1)$-dimensional matrices, where d is the number of Cartesian dimensions (2 or 3). Consequently, the initial atomic coordinates (lattice sites) are strongly correlated with the processor (box and pigeonhole) number.

While annealing can be done by melting a solid, and then quenching the fluid into a glassy state, we chose to accomplish a similar end result by placing atoms initially on square-lattice sites. Since the square lattice is mechanically unstable at any density for central-force interactions (like the LJ spline), we anticipated a rapid, homogeneous nucleation of triangular-lattice crystallites, which would then coalesce into larger and larger grains. Because the excess potential energy of the square lattice is sufficient to melt the sample, even when the initial velocities are nearly zero, we employ a homogeneous, deterministic, feedback mechanism (the so-called Nosé-Hoover thermostated equations of motion [12]) to keep the temperature at about one-fourth the melting temperature for at least 15 or so vibrational periods, before switching to Newtonian, constant-energy equations of motion.

In order to display the annealing process, we color each atom according to its local nearest-neighbor bond environment, assuming hexagonal symmetry, by means of the following algorithm [13]. Each atom i is arbitrarily defined to form a bond with its neighbor j whenever the distance between them (r_{ij}) is less than 1.1 times the zero-temperature, zero-pressure nearest-neighbor distance r_0 (the minimum of the LJ potential). The bond angle with the x axis is then computed according to

$$\theta_{ij}=\cos^{-1}\frac{x_i-x_j}{r_{ij}} , \tag{7}$$

whence the sixfold symmetric grain color is obtained from the sums over neighbors ($w_{ij}=1$ when $r_{ij}\le 1.1r_0$; otherwise, $w_{ij}=0$):

$$\begin{aligned}
\langle\cos 6\theta_i\rangle &= \frac{1}{w}\sum_{j\neq i} w_{ij}\cos 6\theta_{ij} ,\\
\langle\sin 6\theta_i\rangle &= \frac{1}{w}\sum_{j\neq i} w_{ij}\sin 6\theta_{ij} ,\\
w^2 &= \left[\sum_{j\neq i} w_{ij}\cos 6\theta_{ij}\right]^2+\left[\sum_{j\neq i} w_{ij}\sin 6\theta_{ij}\right]^2 ,\\
6\theta_i &= \cos^{-1}\langle\cos 6\theta_i\rangle .
\end{aligned} \tag{8}$$

A color rainbow can be applied such that $\theta_i=-30°$ is deep blue, continuously up to light green for $\theta_i=0°$, and to bright red for $\theta_i=+30°$.

In Fig. 2, the grain map for a NWS-generated velocity distribution is shown at a time of roughly one and a half vibrational periods after initialization of $512^2=262\,144$ atoms on a square lattice. The density was chosen to be close to that of a zero-pressure triangular lattice at one-fourth the melting temperature. Very clear diagonal stripes of correlated triangular crystallites have begun to appear. One can only conclude that since atoms are laid out in regular rows from bottom to top, corresponding to square lattice sites and also parallel processor number, the resulting velocities from the NWS retain some spatial correlation that is clearly undesirable on physical grounds. To give some perspective on the scale of this grain map, usual MD simulations of about 1000 atoms are smaller than the thickness of these stripes, so that the regularity would not be detectable; the sidelength here corresponds to about 0.1 μm.

Subsequent examination of a movie of higher-order correlation plots (i.e., lags $k=1{:}1000$) revealed a small percentage ($<10\%$) of bizarre-looking frames for the NWS, such as is shown in Figs. 3(a) and 3(b) for $k=121$ and 770, respectively. (The particular values of k whose correlation plots are strange depend on the seed, α, which is $\{2^{1/2}\}$ in these plots.) The pairs of random numbers shown in Fig. 3(b) were used in the Box-Muller scheme to generate pairs of Gaussian random numbers. Figure 4 shows that the resulting pairs, plotted against each other, are high correlated. On the other hand, plots for most values of k resemble the unit square in Fig. 1, showing that this occasional needle-in-the-haystack approach for determining correlations in a sequence is not very reliable. We have used the Monkey Test for sequences of length $N=10\,000$, $m=10$, and $r=3$, where the expected value of $\xi_3=9.998$. Table I gives the averaged results for

TABLE I. Results of the Monkey Test for the nested Weyl sequence, $N=10\,000$, $m=10$, $r=3$, and (a) $\alpha=\{33^{1/2}\}$, $E\{\xi_3\}=9.998$ and (b) $\alpha=\{2^{1/2}\}$, $E\{\xi_3\}=9.998$.

				(a)					
Sequences	0,1,7	1,9,8	3,4,8	1,6,2	2,1,1	7,7,2	6,3,9	4,0,1	0,9,5
Observed	49.57	0.00	0.30	0.00	0.00	49.63	0.00	50.11	0.00
Error	0.62	0.00	0.52	0.00	0.00	0.71	0.00	0.63	0.00
				(b)					
Sequences	0,1,7	1,9,8	3,4,8	1,6,2	2,1,1	7,7,2	6,3,9	4,0,1	0,9,5
Observed	0.00	0.00	0.00	0.00	0.00	0.00	36.20	0.00	12.88
Error	0.00	0.00	0.00	0.00	0.00	0.00	0.59	0.00	0.39

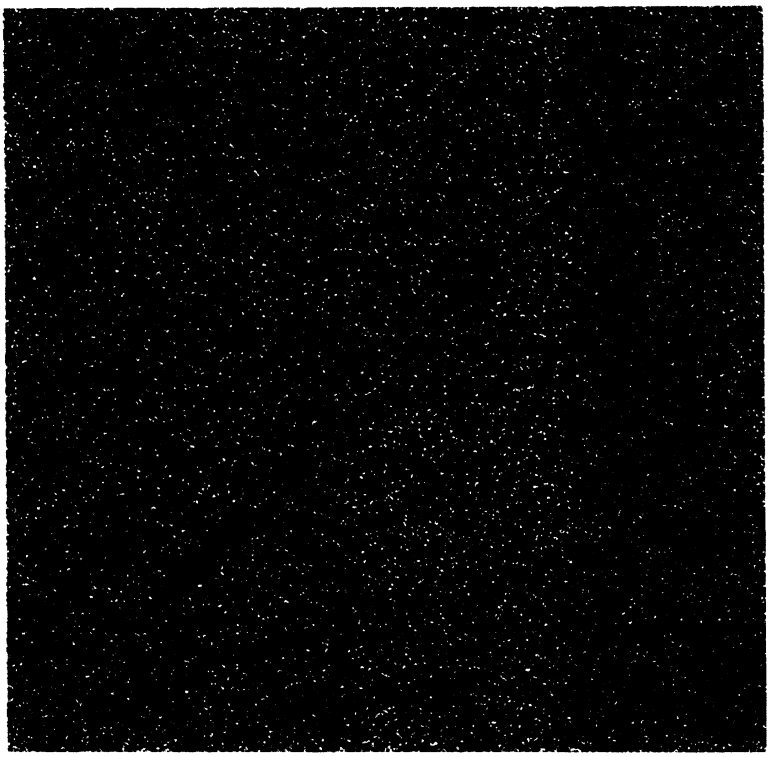

(a)

FIG. 3. Correlation plots (atypical) for NWS [Eq. (6)], generated from seed $\alpha=\{2^{1/2}\}$, $L=65536$: (a) lag $k=121$, (b) lag $k=770$. The strange patterns for these rare lags are in sharp contrast to the usual case shown in Fig. 1.

(b)

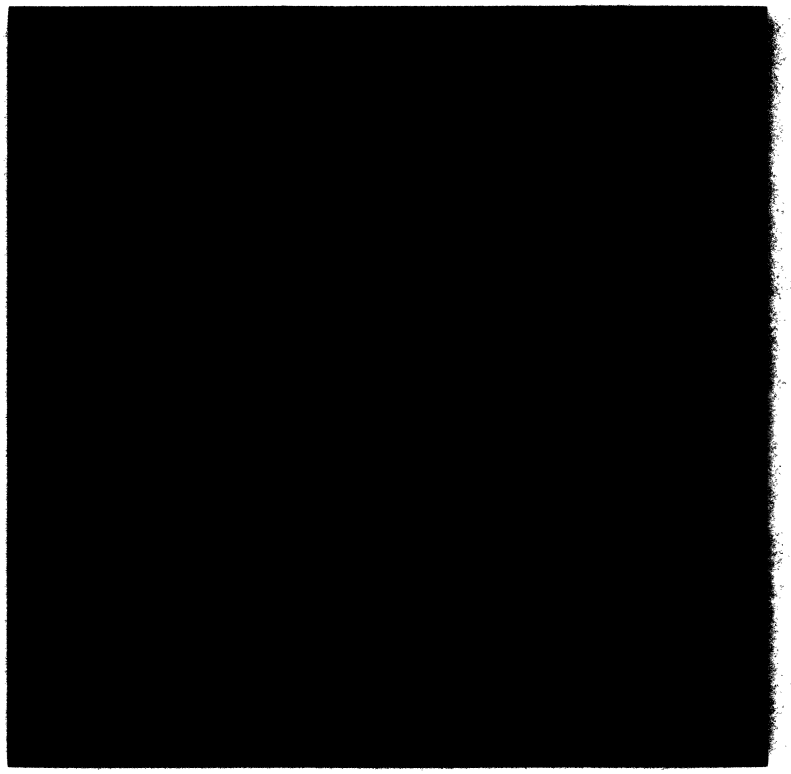

FIG. 4. Plot of y_1 vs y_2 as generated by the Box-Muller scheme [1] with x_i and x_{i+770} used in the calculation shown in Fig. 3(b).

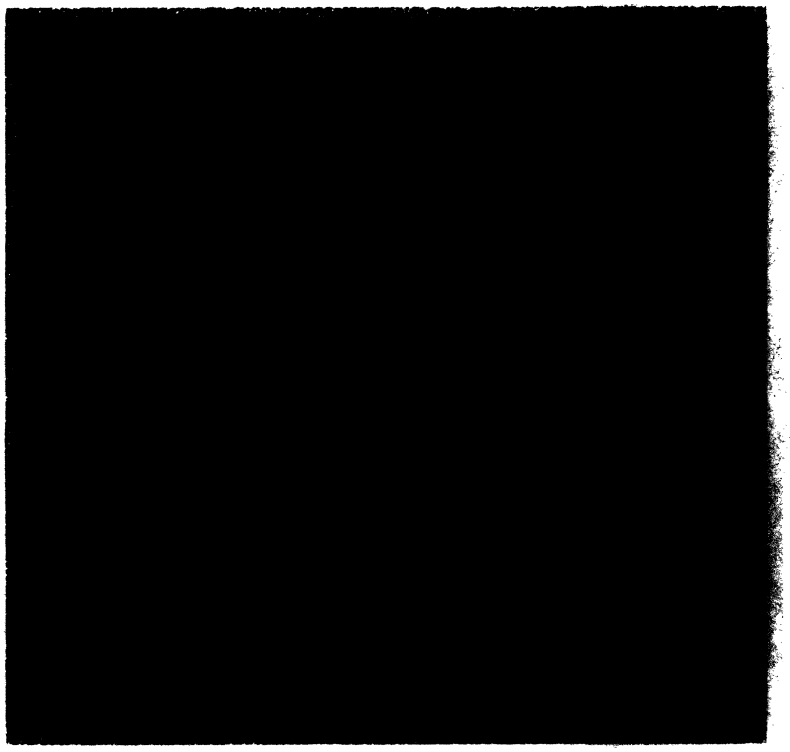

FIG. 5. Lag $k=770$ (not atypical) correlation plot for the SNWS [Eq. (9)], generated from seed $\alpha=\{2^{1/2}\}$, $M=L=65\,536$.

TABLE II. Results of the monkey test for the nested Weyl Sequence, $N=2\times10^6$, $m=100$, $r=3$, and $\alpha=\{2^{1/2}\}$, $E\{\xi_3\}=2.00$.

Sequences	34,80,90	3,57,27	56,1,72	87,87,48	62,47,90	91,17,85	68,33,67	96,12,48
Observed	2.42	0.57	0.37	1.72	1.94	0.28	1.02	0.58
Error	0.86	0.33	0.23	0.72	0.76	0.17	0.41	0.32

100 replications of the Monkey Test for several triplets of numbers in the range (0,9). Table I(a) shows the results for the NWS with the seed $\{33^{1/2}\}$ and Table I(b) shows the results for the seed $\{2^{1/2}\}$. The results are startling in both cases, because only a few, favored triplets of numbers are observed at all in the sequence of 10 000 numbers. The Monkey Test was also done for sequences of length $N=2\times10^6$, $m=100$, and $r=3$. In the case, the expected value of the ξ_3 is 2.00. The averaged results of 100 replications are shown in Table II. At first glance the results do not look completely unacceptable; however, the large errors are suspicious and further investigation shows very correlated results. The first 20 replications of the test found *no* occurrences of any of the triplets studied. When one of the triplets did occur in an individual sequence, it occurred repeatedly, 10 to 60 times. This highly unrandom behavior is reflected in the large errors. Since the NWS shows severe correlation even for $r=3$, it is not surprising that spatial correlations of velocity and processor number are observed in the annealing simulation. We therefore propose a new sequence for generating uniform pseudorandom numbers.

V. SHUFFLED NESTED WEYL SEQUENCE

The unwanted correlation can be removed by first "randomizing" the processor number ν_n, for $n=1,\ldots,L$ so that ν_n lies in the interval $(\frac{1}{2},M+\frac{1}{2})$, where $M\gg1$ is a large integer. multiplier (not smaller than the number of processor L by more than an order of magnitude, however). Then Eq. (6) is applied:

$$\nu_n=MY_n+\tfrac{1}{2}=M\{n\{n\alpha\}\}+\tfrac{1}{2}\ ,$$
$$Z_n=\{\nu_n\{\nu_n\alpha\}\}\ . \tag{9}$$

Such a procedure is referred to as "shuffling" and has been shown [14] to improve the statistical properties of a generator in some cases. The new generator is called, therefore, the shuffled nested Weyl sequence (SNWS).

A lag $k=770$ correlation plot for the SNWS, unlike that for the NWS in Fig. 3(b), is virtually indistinguishable (at least in its lack of features) from all other lags, and is shown in Fig. 5. Furthermore, the Monkey Test was carried out for the same set of symbols as the NWS, but the results are significantly different (see Table III). For the first case, with $N=10\,000$, $m=10$, and $r=3$, the observed values of ξ_3 after 100 replications agree reasonably well with the expected value of 9.998. In the second case, where $N=2\times10^6$, $m=100$, and $r=3$, the agreement with the expected value of $\xi_3=2.00$ is excellent. The SNWS eliminated the spatial correlation of velocities in annealing from the unstable square lattice, as shown in Fig. 6; had the Monkey Test been performed first, these results might have been anticipated. The qualitative uniform features of homogeneous nucleation are in stark contrast with Fig. 2. The annealing results using SNWS are indistinguishable from those obtained from the random number generator implemented by Thinking Machines Corporation on the CM-2 (Wolfram's "Rule 30" cellular automaton [15]). When the new sequence is used for an even larger system (2.4 million atoms) in an annealing simulation, still no undesirable stripping is seen.

In our experience, the apparent length of the sequence is limited only by the floating-point word length. The multiplier is some large positive integer greater than the number of atoms (or the number of random numbers being sampled at a time). The seed need only be the fractional part of such numbers as the square root of any integer that is not a perfect square (2, 3, 5, etc.) or an irrational number (i.e., having a nontrivial fractional part) such as the base of the natural logarithms or π. No sensitivity to the seed, within these guidelines at least, has been observed in practice. Finally, to illustrate the algorithm, we exhibit below some generic pseudocode (RANDY) for the SNWS:

```
mult = 1234567
seed = mod (sqrt (2.), 1.)
do n= 1, number
  x (n)=mod(n*mod(n*seed,1.) ,1.)* mult+0.5
  x (n)=mod(x(n)*mod(x(n)*seed,1.),1.)
end do
```

TABLE III. Results of the Monkey Test for the shuffled nested Weyl sequence, (a) $N=10\,000$, $m=10$, $r=3$, and $\alpha=\{2^{1/2}\}$, $E\{\xi_3\}=9.998$ and (b) $N=2\times10^6$, $m=100$, $r=3$, and $\alpha=\{2^{1/2}\}$, $E\{\xi_3\}=2.00$.

(a)

Sequences	0,1,7	1,9,8	3,4,8	1,6,2	2,1,1	7,7,2	6,3,9	4,0,1	0,9,5
Observed	10.66	9.57	10.34	10.29	9.74	10.37	10.52	9.92	10.02
Error	0.31	0.30	0.30	0.35	0.32	0.32	0.35	0.32	0.32

(b)

Sequences	34,80,90	3,57,27	56,1,72	87,87,48	62,47,90	91,17,85	68,33,67	96,12,48
Observed	2.07	1.67	2.25	2.20	2.17	2.14	2.00	2.19
Error	0.15	0.13	0.14	0.15	0.15	0.16	0.14	0.15

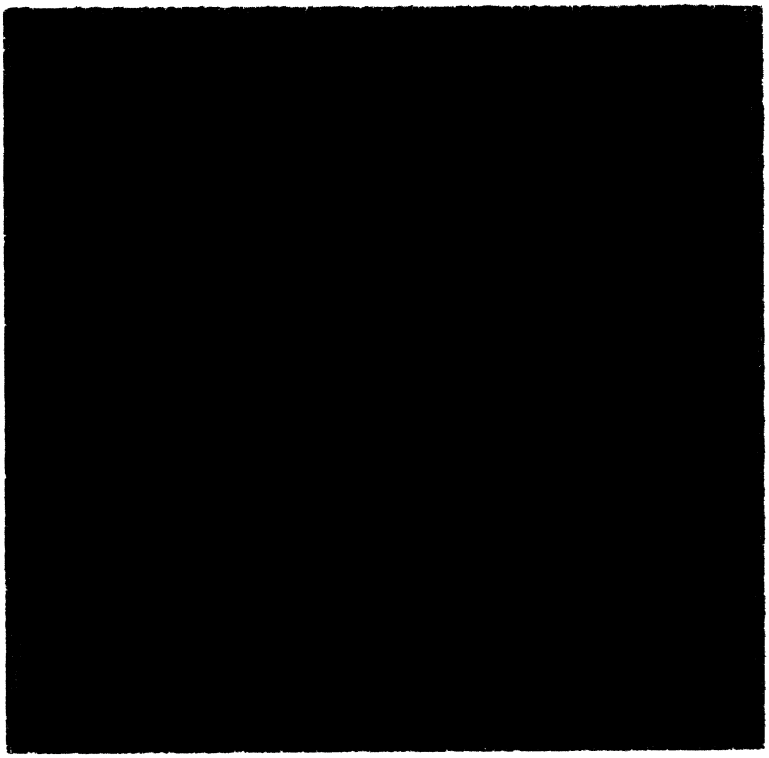

FIG. 6. Grain map for annealing from a square lattice (as in Fig. 2), using initial velocities obtained from the SNWS [Eqs. (9)], using $\{2^{1/2}\}$ and $\{3^{1/2}\}$ as seeds and $M=L=262\,144$. Note the absence of regularity in the grain map, in contrast with Fig. 2.

It has been pointed out that the Weyl sequence is characterized by constant first differences, that is,

$$X_n - X_{n-1} = \alpha - j \ , \quad n = 1,2,\ldots , \quad j = 0,1 \ ,$$

while the nested Weyl sequence is characterized by constant *second* differences,

$$Y_{n+1} - 2Y_n - Y_{n-1} = \{2\alpha\} - j \ ,$$
$$n = 1,2,\ldots , \quad j = -1,0,+1,+2 \ .$$

As we have noted, the correlation function for the latter shows no peculiarity, and the correlation plots show patterns only at infrequent intervals.

VI. CONCLUSIONS

The shuffled nested Weyl sequence (SNWS), which we propose for generating uniform pseudorandom numbers, a natural extension of the Weyl sequence and the nested Weyl sequence (NWS), passes both uniformity and some correlation tests. Moreover, it alone of the three passes an even more stringent dynamical test, namely annealing from an initial unstable crystal lattice. The result is a simple two-line computer algorithm.

ACKNOWLEDGMENTS

B.L.H. and T.T.W. acknowledge support from Los Alamos National Laboratory, which is operated by the University of California for the Department of Energy under Contract No. W-7405-36. Computing was done on the Connection Machine (CM-2) at the Advanced Computing Laboratory at Los Alamos. This work was also supported in part by NSF Grant No. ASC-912 1428 for O.E.P. and P.A.W. and in part by NATO Grant No. CRG 900661 for O.E.P. In addition, B.L.H. thanks Bill Wood, Art Voter, Bill Hoover, Ramon Ravelo, Doug Kothe, and Terry Dontje for the comments and suggestions that were most helpful to him.

[1] G. E. P. Box and M. E. Muller, Ann. Math. Stat. **29**, 610 (1958). If x_1 and x_2 are uniform random numbers on (0,1), then the Box-Muller transformation gives $y_1 = (-2\ln x_1)^{1/2}\cos 2\pi x_2$ and $y_2 = (-2\ln x_1)^{1/2}\sin 2\pi x_2$ as normally distributed (Gaussian) random numbers. However, it should be noted that the Box-Muller scheme amplifies correlation between pairs of pseudorandom numbers.

[2] J. N. Franklin, Math. Comp. **17**, 28 (1963).

[3] W. H. Press, B. P. Flannery, S. A. Teukolsky, and W. T.

Vetterling, *Numerical Recipes* (Cambridge Unviersity, Cambridge, 1989).

[4] (a) D. E. Knuth, *The Art of Computer Programming* (Addison-Wesley, Reading, MA, 1969), Vol. 2; (b) E. J. Dudewicz and T. G. Ralley, *The Handbook of Random Number Generation and Testing with TESTRAND Computer Code* (American Science, Columbus, OH, 1981).

[5] P. Frederickson, R. Hiromoto, T. L. Jordan, B. Smith, and T. Warnock, Parallel Comput. **1**, 175 (1984).

[6] O. E. Percus and M. H. Kalos, J. Parallel Dist. Comput. **6**, 477 (1989).

[7] For a discussion of Lyapunov instability in molecular dynamics simulations, see B. L. Holian, W. G. Hoover, and H. A. Posch, Phys. Rev. Lett. **59**, 10 (1987).

[8] G. Marsaglia (unpublished).

[9] O. E. Percus and P. A. Whitlock (unpublished).

[10] (a) H. Weyl, Math. Ann. **77**, 313 (1916); (b) P. Diaconis, Ann. Prob. **5**, 72 (1977); V. T. Sos, *Studies in Pure Mathematics (To the memory of Paul Turan)* (Birkhauser, Basel, 1983), pp. 685–700.

[11] The Lennard-Jones spline pair potential is described in detail in B. L. Holian, A. F. Voter, N. J. Wagner, R. J. Ravelo, S. P. Chen, W. G. Hoover, C. G. Hoover, J. E. Hammerberg, and T. D. Dontje, Phys. Rev. A **43**, 2655 (1991). (The LJ spline pair potential is obtained by modifying the usual Lennard-Jones 6–12 potential for separation distances r beyond the inflection point, whence a cubic spline in r^2 brings both potential and its first derivative smoothly to zero at r_{max}.)

[12] S. Nosé, J. Chem. Phys. **81**, 511 (1984); W. G. Hoover, Phys. Rev. A **31**, 1695 (1985).

[13] A. F. Voter, B. L. Holian, and N. J. Wagner (unpublished).

[14] M. D. McLaren and G. Marsaglia, J. Assoc. Comput. Mach. **12**, 83 (1965); B. D. Ripley, *Stochastic Simulation* (Wiley, New York, 1987), pp. 42 and 43.

[15] S. Wolfram, Adv. Appl. Math. **7**, 123 (1986).

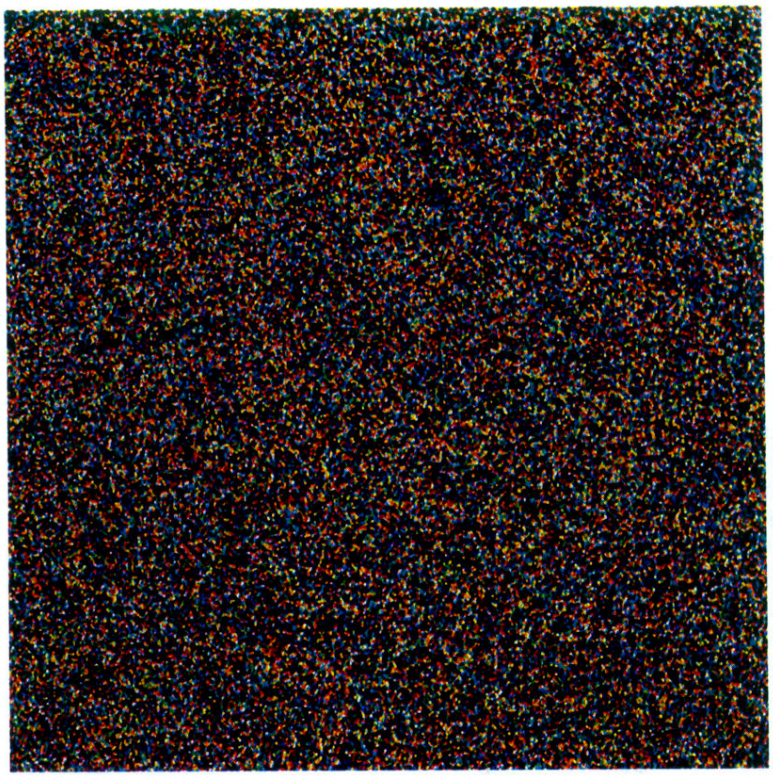

FIG. 1. Lag $k=1$ correlation plot (unit square) for the NWS [Eq. (6)], generated from the seed $\alpha=\{2^{1/2}\}$. The order of the pairs (Y_i, Y_{i+1}) is colored by a rainbow, where deep blue is $i=1$, light green is near $i=32\,768$, up to bright red for $i=L=65\,536$.

FIG. 2. Grain map [Eqs. (7) and (8)] of 262 144 Lennard-Jones spline atoms in two dimensions (about 0.1 μm on a side) at a time of about one and one-half vibrational periods after initialization from a square lattice (density appropriate to zero pressure and one-fourth of the melting temperature), with initial x and y velocities chosen from Box-Muller transformation of two NWS, using $\{2^{1/2}\}$ and $\{3^{1/2}\}$ as seeds. The grain-map rainbow corresponds to deep blue for crystallites with triangular-lattice bond angles (with x axis) near $-30°$, to light green near 0°, up to bright red for $+30°$. Note the unphysical regularity of diagonal stripes of crystallites.

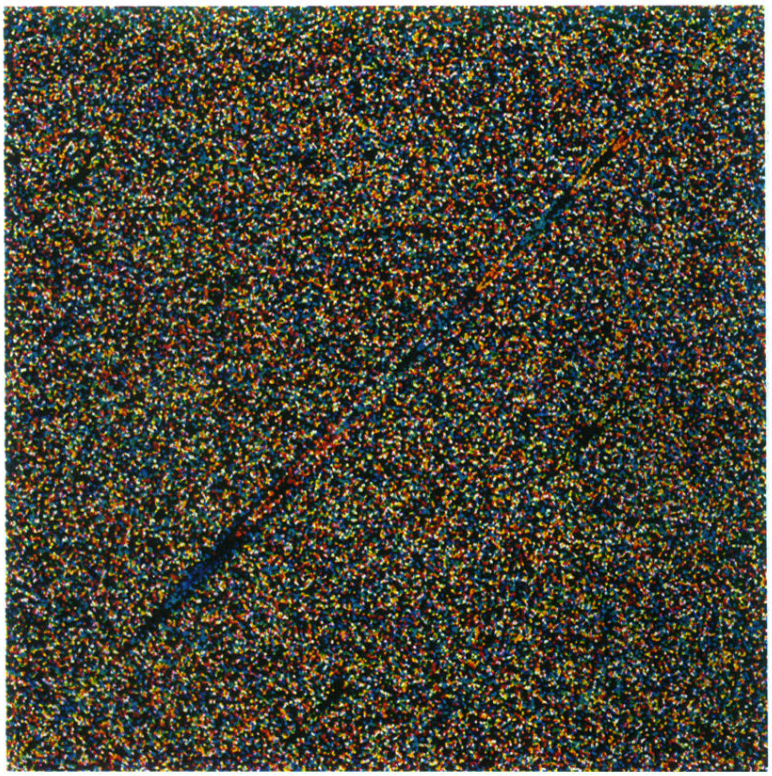

(a)

FIG. 3. Correlation plots (atypical) for NWS [Eq. (6)], generated from seed $\alpha = \{2^{1/2}\}$, $L = 65536$: (a) lag $k = 121$, (b) lag $k = 770$. The strange patterns for these rare lags are in sharp contrast to the usual case shown in Fig. 1.

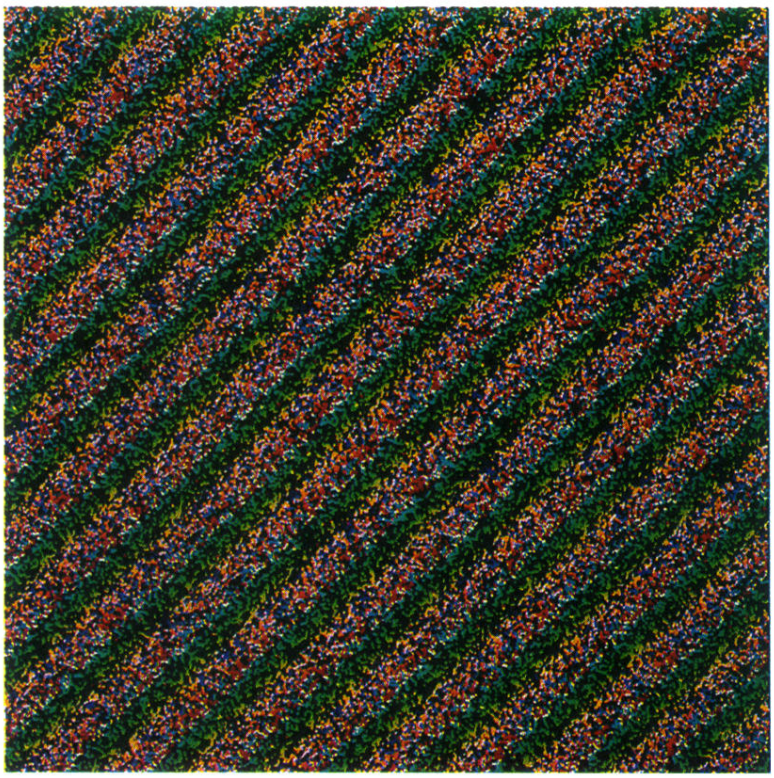

(b)

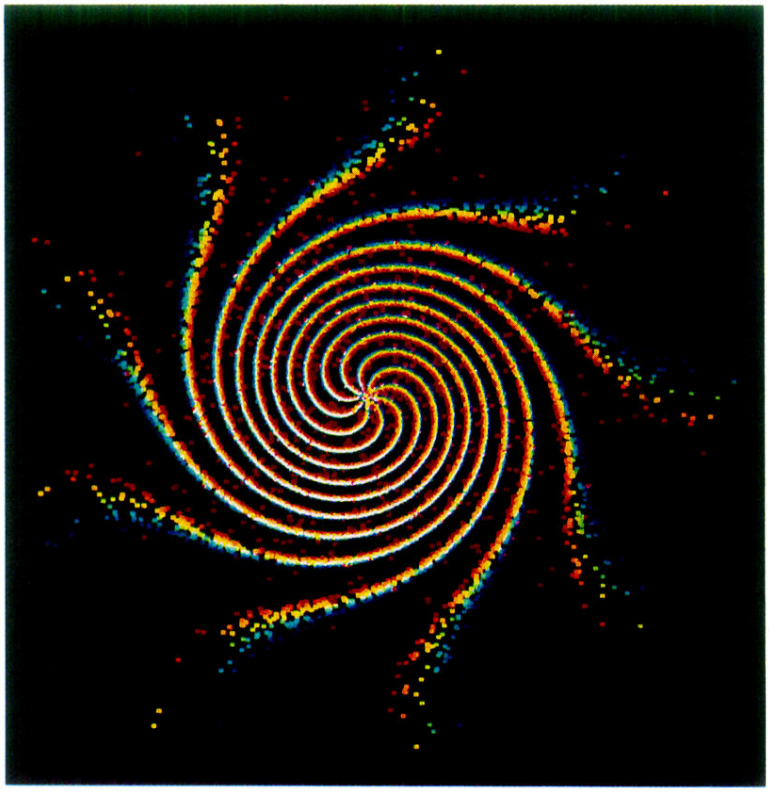

FIG. 4. Plot of y_1 vs y_2 as generated by the Box-Muller scheme [1] with x_i and x_{i+770} used in the calculation shown in Fig. 3(b).

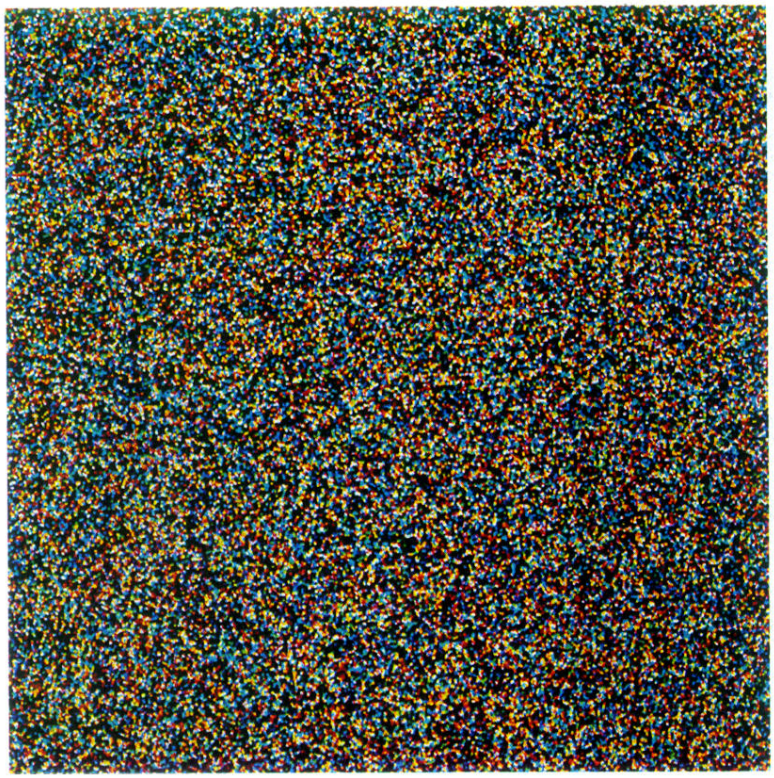

FIG. 5. Lag $k=770$ (not atypical) correlation plot for the SNWS [Eq. (9)], generated from seed $\alpha=\{2^{1/2}\}$, $M=L=65\,536$.

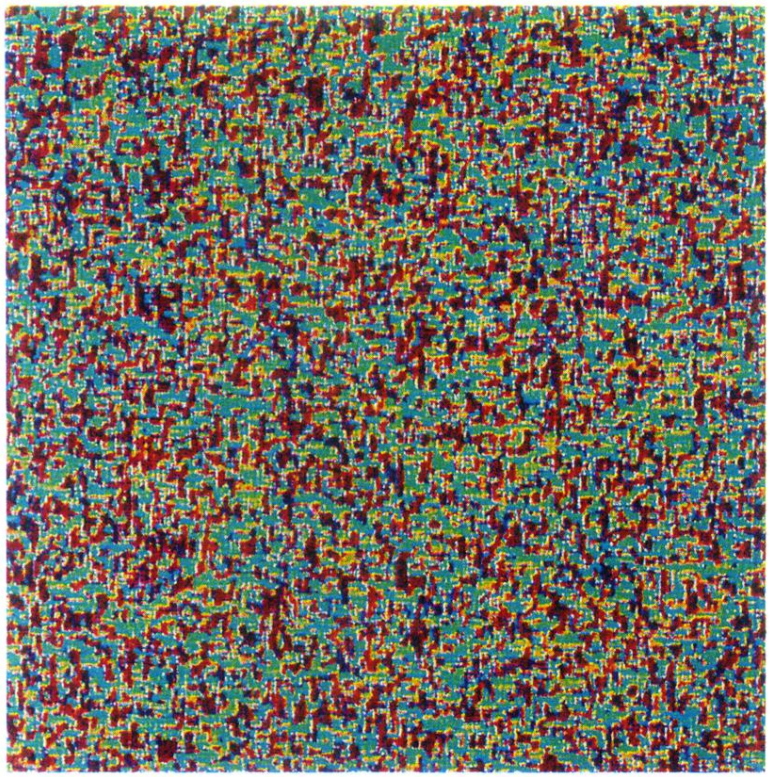

FIG. 6. Grain map for annealing from a square lattice (as in Fig. 2), using initial velocities obtained from the SNWS [Eqs. (9)], using $\{2^{1/2}\}$ and $\{3^{1/2}\}$ as seeds and $M=L=262\,144$. Note the absence of regularity in the grain map, in contrast with Fig. 2.