

Systematic errors due to linear congruential random-number generators with the Swendsen-Wang algorithm: A warning

Giovanni Ossola* and Alan D. Sokal†

Department of Physics, New York University, 4 Washington Place, New York, New York 10003, USA

(Received 8 March 2004; published 30 August 2004)

We show that linear congruential pseudo-random-number generators can cause systematic errors in Monte Carlo simulations using the Swendsen–Wang algorithm, if the lattice size is a multiple of a very large power of 2 and one random number is used per bond. These systematic errors arise from correlations within a single bond-update half-sweep. The errors can be eliminated (or at least radically reduced) by updating the bonds in a random order or in an aperiodic manner. It also helps to use a generator of large modulus (e.g., 60 or more bits).

DOI: 10.1103/PhysRevE.70.027701

PACS number(s): 05.10.Ln, 05.50.+q, 02.70.Tt, 02.70.Uu

I. INTRODUCTION

It has been known for about two decades that linear congruential pseudo-random-number generators [1] suffer from strong long-range correlations [2–6]: for instance, generators with modulus $m=2^\beta$ have strong correlations at lags that are multiples of 2^k whenever the ratio k/β is large enough [2–4]. Furthermore, these long-range correlations are known to give rise to systematic errors in Monte Carlo simulations employing local (e.g., Metropolis or heat bath) updates whenever the lattice sites are updated in a fixed order and the number of random numbers used per sweep is a multiple of a large enough power of 2: this happens because one is using strongly correlated random numbers to update the same lattice sites in successive sweeps (within roughly one autocorrelation time) [2,3,7–10]. On the other hand, these systematic errors can be eliminated by the simple expedient of throwing away one random number at the end of each lattice sweep [2,3,7–10].

It has generally been thought that nonlocal algorithms such as the Swendsen–Wang algorithm [11] and Wolff’s single-cluster variant [12] would be immune to these particular defects of linear congruential generators, inasmuch as they employ random numbers in a highly aperiodic way both in “space” and in “time.” We were therefore astonished to find, in our Swendsen–Wang simulation of the three-dimensional Ising model [13], large systematic errors on the 128^3 and 256^3 lattices that we eventually traced (after much wringing of hands) precisely to long-range correlations in the random-number generator.

Recall that one iteration of the Swendsen–Wang (SW) algorithm consists of two steps: first one updates the bond occupation variables at a fixed configuration of the Ising spin variables; then one computes the connected clusters associated to the bond configuration and updates the Ising spin variables by choosing a new spin value independently for each cluster. The second (spin-update) half of the SW algorithm indeed uses random numbers in a thoroughly aperiodic

way, because the cluster sizes and shapes are random. But the first (bond-update) half uses random numbers in a highly structured way: typically one sweeps the bonds of the lattice in some simple fixed order (e.g., lexicographic). Therefore, if the lattice size is very large, the effects of the long-range correlations of the random-number generator can be observed *within a single half-sweep*: the random numbers used in updating the bonds of one part of the lattice will be strongly correlated with those used elsewhere in the lattice. One may expect this correlation to cause systematic errors particularly if (a) the lattice size is commensurate with the lag giving rise to long-range correlations (e.g., a power of 2), and (b) the system’s correlation length is large enough so that the long-range correlations of the random-number generator couple *correlated* parts of the lattice.

The purpose of this Brief Report is, first of all, to provide evidence that such systematic errors can indeed occur and that we have accurately diagnosed their origin; and secondly, to show how the implementation of the Swendsen–Wang algorithm can be modified so as to eliminate (or at least radically reduce) these systematic errors. A more detailed account will be published elsewhere [14].

II. EVIDENCE OF SYSTEMATIC ERRORS

We simulated the nearest-neighbor three-dimensional Ising model on an $L \times L \times L$ simple-cubic lattice with periodic boundary conditions, using the Swendsen–Wang (SW) algorithm [11]. We studied lattice sizes $L=4, 6, 8, 12, 16, 24, 32, 48, 64, 96, 128, 192,$ and 256 , and performed between 10^7 and 10^8 SW iterations for each lattice size. We did all our runs at $\beta=0.221\ 654\ 59$, which is Blöte *et al.*’s best estimate of the critical temperature [15] and is very near to the estimates of other workers [16,17] (see also the review [18]). We measured a large number of observables, including the susceptibility χ , the second-moment correlation length ξ , the energy E , and the specific heat C_H .

In the first version of our program, the random numbers were supplied by a linear congruential generator with modulus $m=2^{48}$, increment $c=1$, and multiplier $a=31\ 167\ 285, 10\ 430\ 376\ 854\ 301,$ $77\ 596\ 615\ 844\ 045,$ or $181\ 465\ 474\ 592\ 829$. All these multipliers give good results

*Electronic address: giovanni.ossola@physics.nyu.edu

†Electronic address: sokal@nyu.edu

TABLE I. Results of our Swendsen–Wang simulations on the three-dimensional Ising model at criticality, using a linear congruential generator with modulus $m=2^{48}$. Error bar (one standard deviation) is shown in parentheses. The row marked $L=\infty$ indicates our best estimate of the asymptotic value x^* . The last two columns indicate the deviation of each point from the fit curve (2), in percent and in standard deviations. Points deviating by more than 3σ are marked in boldface.

| L | ξ/L | deviation (%) | deviation (σ) |
|----------|-------------|---------------|-------------------------|
| 4 | 0.63000(9) | −0.566% | − 17.48 σ |
| 6 | 0.63576(10) | −0.076% | − 3.24 σ |
| 8 | 0.63769(10) | 0.004% | 0.44 σ |
| 12 | 0.63909(11) | −0.013% | −0.83 σ |
| 16 | 0.63998(8) | 0.002% | 0.15 σ |
| 24 | 0.64093(13) | 0.015% | 0.83 σ |
| 32 | 0.64122(13) | −0.010% | −0.51 σ |
| 48 | 0.64172(10) | −0.007% | −0.51 σ |
| 64 | 0.64223(16) | 0.032% | 1.44 σ |
| 96 | 0.64219(16) | −0.017% | −0.76 σ |
| 128 | 0.62215(25) | −3.159% | − 78.94 σ |
| 192 | 0.64383(38) | 0.192% | 3.23 σ |
| 256 | 0.77798(79) | 21.052% | 169.69 σ |
| ∞ | 0.64299(8) | | |

on the spectral test in low dimensions, compared to other multipliers for the same modulus [1,19]. We verified that the runs with the four different multipliers gave results that are consistent within error bars for all the major observables; after making this verification, we averaged all the runs for each L .

The results for the correlation length ξ are reported in the first two columns of Table I. Finite-size-scaling theory predicts that ξ/L should behave for large L (if indeed we are at the critical temperature) as

$$\xi/L = x^* + AL^{-\omega} + \dots, \quad (1)$$

where x^* is a universal amplitude ratio characteristic of the given system with periodic boundary conditions, ω is a correction-to-scaling exponent, A is a nonuniversal correction-to-scaling amplitude, and the dots indicate higher-order corrections to scaling. The data in Table I are qualitatively consistent with Eq. (1), *except for the points at $L=128$ and $L=256$, which show extremely large deviations.*

A closer examination of the data in Table I reveals that the point at $L=192$ may also exhibit a small but statistically significant deviation from the fitting curve. To make all these observations more quantitative, let us perform a weighted least-squares fit to Eq. (1) with $\omega=0.82$ (the best estimate from [15]), using all the data with $L_{\min} \leq L \leq 96$ and varying L_{\min} while checking the goodness of fit. A good fit ($\chi^2=3.85$, 6 DF, confidence level = 70%) can be obtained already with $L_{\min}=8$, yielding

$$\xi/L \approx 0.64299(8) - 0.02931(79)L^{-0.82}. \quad (2)$$

Not surprisingly, the points $L=4$ and $L=6$ show significant deviations from the fit curve, due to higher-order corrections to scaling. More surprising are the points $L=128$, which lies roughly 3% (≈ 79 standard deviations) below the fit curve, and $L=256$, which lies a whopping 21% (≈ 170 standard deviations) above the fit curve. Clearly, there are large systematic errors on the lattices $L=128$ and $L=256$. Finally, the point $L=192$ lies approximately 0.2% (≈ 3 standard deviations) above the fit curve: this *may* indicate the presence of a small systematic error also for this lattice.

At first we worried whether we had made a programming error that might lead to incorrect results on large lattices (e.g., due to integer overflow). We checked the program carefully and were unable to find any such mistakes. Moreover, the fact that the systematic discrepancy is much smaller (if it exists at all) at $L=192$ than at $L=128$ suggests that the problem—whatever its cause—does not arise *solely* from the lattice being large.

Intrigued by the fact that these large discrepancies might be arising only at lattice sizes that are large powers of 2 (or perhaps multiples of large powers of 2), we made shorter runs (between 3×10^4 and 10^6 SW iterations) at many other lattice sizes—all multiples of 2 from 4 through 140, and all multiples of 10 through 250—in order to check whether any other deviant points could be found. The upshot is that—within the statistical error of these shorter runs, which ranges from 0.2% on small lattices to 1% at $L \approx 128$ to an admittedly rather crude 2.3% on the largest lattices—there are no detectable discrepancies except at $L=128$ and 256.

At $L=128$ and $L=256$, we found discrepancies not only for the correlation length but also for the susceptibility, the energy, and the specific heat. It is a curious fact, however, that all the Fortuin-Kasteleyn identities [Ref. [20], Eqs. (3.20)–(3.23)] are verified perfectly (to within statistical error). This contrasts with the systematic errors found by Damgaard and Heller [8] in a Metropolis Monte Carlo simulation of the $U(1)$ Higgs model, where a Ward identity was violated by up to 10 standard deviations, and those found by Ballesteros and Martín-Mayor [21] in a Wolff single-cluster simulation of the two- and three-dimensional Ising models, in which Schwinger–Dyson identities were violated by up to eight standard deviations.

III. VARIANT SIMULATIONS

In the Introduction, we have argued that long-range correlations in the random-number generator can cause undesired correlations within a single bond-update half-sweep, leading to systematic errors. In order to test whether our proposed explanation for the observed systematic errors is the correct one, we ran variant simulations in which two aspects of the simulation were systematically altered: the modulus $m=2^\beta$ of the random-number generator ($\beta=16, 20, 24, 28, 32, 40, 48, 60, 63$, and 64), and the manner in which the random numbers are used within the bond-update subroutine. The latter test is essential if we are to prove not only that the trouble comes from the random-number generator,

but more specifically that it comes from the way that the random numbers are used *in the bond-update subroutine*.

All of the multipliers used here give good results on the spectral test in low dimensions compared to other multipliers for the same modulus.¹ The purpose of trying random-number generators with fewer than 48 bits was to induce systematic errors on small lattices where they could be studied quantitatively to high precision and compared with those observed with the 48-bit generator on larger lattices. The purpose of trying random-number generators with 60/63/64 bits was, of course, to provide a standard of comparison in which the systematic error is eliminated or at least radically reduced.

We also tried three variants of the bond-update subroutine:

Standard. This is our original program, in which the bonds are updated in lexicographic order, and one random number is used per bond.

Aperiodic. Here the bonds are again updated in lexicographic order, but a random number is used only if the two spins are equal. (If the two spins are unequal, the corresponding bond is automatically left unoccupied, so no random number is needed.) If our explanation of the cause of the systematic errors is correct, this strategem should eliminate the systematic errors on lattices that are multiples of large powers of 2, though it may conceivably shift those systematic errors to other lattice sizes (namely, those for which the lattice size, multiplied by the fraction of nearest-neighbor spins that are equal, yields a suitable “resonance”).

Shuffle. The bonds are updated in a random order.² If our explanation of the cause of the systematic errors is correct, this strategem should entirely eliminate the systematic errors, even with a relatively poor (e.g., 32-bit) random-number generator.

Our first version of the “shuffle” subroutine permuted the array containing the bond indices. Unfortunately, this program ran very slowly—about a factor of 2 slower than the “standard” version at $L=16$, growing to a factor ≈ 8 at $L=256$ —probably because the highly nonlocal access to the bond array caused a large number of cache misses. Our second version permuted instead the array of random numbers;³ this is statistically equivalent but allows the bond array to be accessed in sequential order. This program ran less slowly, once again about a factor of 2 slower than the “standard” version at $L=16$, but growing only to a factor ≈ 4 at $L=256$.

¹For moduli 2^{32} , 2^{40} , 2^{48} , 2^{60} , 2^{63} , and 2^{64} , these multipliers can be found in published tables of multipliers that perform well on the spectral test [1,19]. We double-checked these computations, and performed analogous computations from scratch for the smaller moduli.

²A uniform random permutation of n elements can easily be generated, in a time of order n , using $n-1$ random numbers random numbers [Ref. [1], pp. 139–140].

³More precisely, it permuted a LOGICAL array containing the results of the comparisons of the random numbers against $p=1-e^{-2\beta}$. This requires only one byte storage per bond, rather than eight bytes for storing the random number itself, thereby reducing both memory usage and cache misses during the generation of the random permutation.

The results of all these variant simulations, carried out on lattice sizes $L=8, 16, 32, 64, 96, 128, 192$, and 256 will be reported elsewhere [14]; here we provide only a brief summary. We find that the 60/63/64-bit generators give consistent results (within statistical error) for all three variants of the bond-update subroutine, confirming our expectation that they exhibit negligible systematic error on lattices $L \leq 256$.⁴ By contrast, each “standard” algorithm with ≤ 48 bits exhibits detectable systematic errors whenever the lattice size L is a multiple of a sufficiently large power of 2; how large depends on the modulus. More precisely, the 16-bit (resp. 20-bit, 24-bit, 28-bit, 32-bit, 40-bit, 48-bit) standard algorithm exhibits detectable systematic errors whenever L is a multiple of 8 (resp. 8, 8, 16, 32, 64, 128). In addition, the 48-bit “standard” algorithm at $L=192$ shows a discrepancy of almost 3σ , which *may* indicate a systematic error. No other statistically significant discrepancies are observed.

We conclude that, if one wants to use a linear congruential generator with the Swendsen–Wang algorithm, the safest approach is to use a generator of 64 bits (or more) together with the “shuffle” bond update. Unfortunately, the shuffle method is somewhat slow. A much faster—and, as far as we can tell, also safe—method is to use a 64-bit generator together with the “aperiodic” bond update.

Despite the known problems of linear congruential generators arising from long-range correlations, there are still several advantages in using them. First, they are relatively cheap in terms of CPU time, and are convenient for use in a series of successive runs because the complete state of the generator can be saved in a single computer word. More importantly, they are well understood theoretically, as regards both short-range [1] and long-range [2–6,14] correlations; in particular, excellent equidistribution of t -tuples of successive random numbers for small t can be achieved by careful choice of the multiplier. By contrast, for more exotic random-number generators (e.g., combination generators), the problems may not be absent, but simply hidden.

A more detailed analysis of these simulations will be published elsewhere [14], along with a discussion of the advantages and disadvantages of linear congruential versus other types of pseudo-random-number generators (see also [21–24]).

ACKNOWLEDGMENTS

We wish to thank Henk Blöte, Luís Antonio Fernández, Werner Kerler, Juan Ruiz-Lorenzo, Lev Shchur, and Jian-Sheng Wang for valuable correspondence and for sharing their unpublished data; Mulin Ding and Raj Sivanandarajah for efficient assistance with numerous aspects of the Physics Department computing system; and Román Scoccimarro for generously letting us do some test runs on his DEC Alpha computer. This research was supported in part by U.S. National Science Foundation Grants No. PHY-0099393 and No. PHY-0116590.

⁴We are continuing runs on the lattice $L=256$ in an effort to detect very small systematic errors. These results will be reported later [14].

- [1] D. E. Knuth, *The Art of Computer Programming*, 2nd ed. (Addison-Wesley, Reading, MA 1981), Vol. 2.
- [2] C. Kalle and S. Wansleben, *Comput. Phys. Commun.* **33**, 343 (1984).
- [3] T. Filk, M. Marcu, and K. Fredenhagen, *Phys. Lett.* **165B**, 125 (1985).
- [4] O. E. Percus and J. K. Percus, *J. Comput. Phys.* **77**, 267 (1988).
- [5] A. DeMatteis and S. Pagnutti, *Numer. Math.* **53**, 595 (1988).
- [6] J. Eichenauer-Herrmann and H. Grothe, *Numer. Math.* **56**, 609 (1989).
- [7] I. Montvay and P. Weisz, *Nucl. Phys. B* **290**, 327 (1987), see p. 340.
- [8] P. H. Damgaard and U. M. Heller, *Nucl. Phys. B* **309**, 625 (1988), see p. 639.
- [9] B. A. Berg and A. H. Billoire, *Phys. Rev. D* **40**, 550 (1989), see Appendix B.
- [10] M. Baig, H. Fort, J. B. Kogut, and S. Kim, *Phys. Rev. D* **51**, 5216 (1995), see p. 5221.
- [11] R. H. Swendsen and J.-S. Wang, *Phys. Rev. Lett.* **58**, 86 (1987).
- [12] U. Wolff, *Phys. Rev. Lett.* **62**, 361 (1989).
- [13] G. Ossola and A. D. Sokal, *Nucl. Phys. B* **691**, 259 (2004), and unpublished.
- [14] G. Ossola and A. D. Sokal (unpublished).
- [15] H. W. J. Blöte, L. N. Shchur, and A. L. Talapov, *Int. J. Mod. Phys. C* **10**, 1137 (1999).
- [16] H. G. Ballesteros, L. A. Fernández, V. Martín-Mayor, A. Muñoz Sudupe, G. Parisi, and J. J. Ruiz-Lorenzo, *J. Phys. A* **32**, 1 (1999).
- [17] M. Hasenbusch, K. Pinn, and S. Vinti, *Phys. Rev. B* **59**, 11471 (1999).
- [18] M. Hasenbusch, *Int. J. Mod. Phys. C* **12**, 911 (2001).
- [19] P. L'Ecuyer, *Math. Comput.* **68**, 249 (1999).
- [20] J. Salas and A. D. Sokal, *J. Stat. Phys.* **98**, 551 (2000).
- [21] H. G. Ballesteros and V. Martín-Mayor, *Phys. Rev. E* **58**, 6787 (1998).
- [22] A. M. Ferrenberg, D. P. Landau, and Y. J. Wong, *Phys. Rev. Lett.* **69**, 3382 (1992).
- [23] I. Vattulainen, T. Ala-Nissila, and K. Kankaala, *Phys. Rev. Lett.* **73**, 2513 (1994).
- [24] P. D. Coddington, *Int. J. Mod. Phys. C* **5**, 547 (1994).