Suppressing Roughness of Virtual Times in Parallel Discrete-Event Simulations

G. Korniss,[1]* M. A. Novotny,[2] H. Guclu,[1] Z. Toroczkai,[3] P. A. Rikvold[4]

In a parallel discrete-event simulation (PDES) scheme, tasks are distributed among processing elements (PEs) whose progress is controlled by a synchronization scheme. For lattice systems with short-range interactions, the progress of the conservative PDES scheme is governed by the Kardar-Parisi-Zhang equation from the theory of nonequilibrium surface growth. Although the simulated (virtual) times of the PEs progress at a nonzero rate, their standard deviation (spread) diverges with the number of PEs, hindering efficient data collection. We show that weak random interactions among the PEs can make this spread nondivergent. The PEs then progress at a nonzero, near-uniform rate without requiring global synchronizations.

Simulating large systems often leaves the programmer with only one option: parallel distributed simulations where parts of the system are allocated and simulated on different processing elements (PEs). A large class of interacting systems, including financial market models, epidemic models, dynamics of magnetic systems, and queuing networks, can be described by a set of local state variables assuming a finite number of possible values. As the system evolves in time, the values of the local state variables change at discrete instants, synchronously or asynchronously depending on the dynamics of the system. Parallel simulation for the former is straightforward (at least conceptually). For the latter—that is, for asynchronous or non-parallel dynamics—one must use some kind of synchronization to ensure causality. The instantaneous changes in the local configuration are also called discrete events, hence the term parallel discrete-event simulation (PDES) (1–3). Examples of PDES applications include dynamic channel allocation in cell phone communication networks (3, 4), models of the spread of diseases (5), battlefield simulations (6), and dynamic phenomena in highly anisotropic magnetic systems (7, 8). Here the discrete events are call arrivals, infections, troop movements, and changes of the orientation of the local magnetic moments, respectively. As the number of PEs on parallel architectures increases to tens of thousands, fundamental questions of the scalability of the underlying algorithms must be addressed. Here, we show a way to construct fully scalable parallel simulations for systems with asynchronous dynamics and short-range interactions. Understanding the effects of the microscopic dynamics (corresponding to the algorithmic synchronization rules) on the global properties of the simulation scheme brings us to the solution. Recently, a similar connection has been made (9) between roll-back-based PDES schemes (10) and self-organized criticality (11).

The two basic ingredients of PDES are the set of local simulated times, often referred to as virtual times (10), and a synchronization scheme (1). For the PDES scheme to be scalable (12), two criteria must be met: (i) The virtual time horizon should progress on average at a nonzero rate, and (ii) the typical spread of the time horizon should be bounded as the number of PEs N_{PE} goes to infinity. The first criterion ensures a nonzero progress rate in the limit of large N_{PE}. It is, however, not sufficient if data are to be collected. Different PEs have progressed to different local simulated times with a possibly large spread among them, making measurement a complex task. Frequent global synchronizations can get costly for large N_{PE}, whereas temporarily storing a large amount of data as a result of the large virtual time spread is limited by the available memory. Therefore, criterion (ii) is crucial for the measurement

part of the algorithm to be scalable. Here, we introduce a PDES scheme in which the PEs make nonzero and close-to-uniform progress without global intervention.

In conservative PDES schemes (13–15), which we focus on, an update is performed by a particular PE only if the resulting change in the local configuration of the simulated system is guaranteed not to violate causality. Otherwise, the PE idles. The efficiency of the scheme depends on the fraction of nonidling PEs. It was shown (16, 17) that the virtual time horizon exhibits kinetic roughening (18, 19) for the basic conservative scheme applied to systems with short-range interactions on regular lattices. In particular, the evolution of the virtual time horizon is governed by the Kardar-Parisi-Zhang (KPZ) equation (20), which plays a central role in nonequilibrium surface growth (18, 19). The above finding has two major implications for the asymptotic scalability of the basic conservative PDES scheme (16, 21): Criterion (i) for the scalability is satisfied because the average progress rate of the virtual time horizon approaches a nonzero value in the limit $N_{PE} \to \infty$. Criterion (ii), however, is violated because the virtual time horizon becomes macroscopically rough.

For illustration, we consider a general one-dimensional system with nearest-neighbor interactions, in which the discrete events exhibit Poisson asynchrony. In the one-site-per-PE scenario, each site has its own local simulated time, constituting the virtual time horizon $\{\tau_i(t)\}_{i=1}^{N_{PE}}$, where t is the discrete number of parallel steps executed by all PEs (which is proportional to the wall-clock time). According to the basic conservative synchronization scheme (14, 15), at each parallel step t, only those PEs for which the local simulated time is not greater than the local simulated times of their neighbors can increment their local time by an exponentially distributed random amount. [Without loss of generality, we assume that the mean of the local time increment is 1 in simulated time units (stu).] Thus, if $\tau_i(t) \leq \min\{\tau_{i-1}(t), \tau_{i+1}(t)\}$, PE i can update the configuration of the underlying site it carries and determine the time of the next event. Otherwise, it idles. Despite its simplicity, this rule preserves unaltered the asynchronous causal dynamics of the underlying system (14, 15).

The progress rate of the simulation $\langle u(t) \rangle_{N_{PE}}$ (the density of local minima of the virtual time horizon) approaches a nonzero constant in the asymptotic long-time, large-N_{PE} limit (16, 21). The average width of the virtual time horizon, however, diverges as $N_{PE} \to \infty$ (16, 17). Specifically, the average width is defined as

[1]Department of Physics, Applied Physics, and Astronomy, Rensselaer Polytechnic Institute, 110 8th Street, Troy, NY 12180, USA. [2]Department of Physics and Astronomy and ERC Center for Computational Science, Mississippi State University, Post Office Box 5167, Mississippi State, MS 39762, USA. [3]Complex Systems Group, Theoretical Division, Los Alamos National Laboratory, Mail Stop B-213, Los Alamos, NM 87545, USA. [4]Department of Physics, Center for Materials Research and Technology, and School of Computational Science and Information Technology, Florida State University, Tallahassee, FL 32306, USA.

*To whom correspondence should be addressed. E-mail: korniss@rpi.edu

$$\langle w^2(t)\rangle_{N_{PE}} = (1/N_{PE}) \left\langle \sum_{i=1}^{N_{PE}} [\tau_i(t) - \bar{\tau}(t)]^2 \right\rangle \quad (1)$$

where

$$\bar{\tau}(t) = (1/N_{PE}) \sum_{i=1}^{N_{PE}} \tau_i(t) \quad (2)$$

For any finite number of PEs, the width grows for early times as $\langle w^2(t)\rangle_{N_{PE}} \sim t^{2\beta}$, where β is the growth exponent and is equal to $\frac{1}{3}$. After a system size–dependent cross-over time $t_\times \sim N_{PE}^z$ (where z is the dynamic exponent and is equal to $\frac{3}{2}$), the width reaches its steady state with $\langle w^2(\infty)\rangle_{N_{PE}} \sim N_{PE}^{2\alpha}$, where α is the roughness exponent—equal to $\frac{1}{2}$ in our example—which quantifies how the average width of the surface diverges in the large-N_{PE} limit. The reason for this divergence is that the local interaction topology of the PEs mimics that of the underlying system. Under the basic conservative synchronization rules, the PEs form a strongly interacting system in which the correlation length reaches the system size N_{PE} (16). This correlation length is responsible for the diverging surface fluctuations captured by the average width.

To reduce the width, one must decorrelate the fluctuations of the virtual time surface. We achieve this by introducing quenched random communication links between the PEs in addition to the regular nearest-neighbor interactions of the underlying physical system. At the beginning of the simulation, we connect each PE to another one, chosen randomly from the rest. These random links, in addition to the regular lattice connections, result in a small-world network (22) where, beyond the nearest neighbors, each PE is connected with a randomly chosen one (23). In the modified conservative PDES scheme, at every parallel step, each PE with probability p compares its local simulated time with its full virtual neighborhood and can only advance if it is a local minimum—that is, if $\tau_i(t) \leq \min\{\tau_{i-1}(t), \tau_{i+1}(t), \tau_{r(i)}(t)\}$, where $r(i)$ is the random connection of PE i. Each PE follows the original scheme with probability $(1 - p)$. Note that the occasional extra checking of the simulated time of the random neighbor is not needed for the faithfulness of the simulation. It is merely introduced to control the width of the time horizon.

Using a coarse-graining procedure analogous to that used in (16), we find that the large-scale properties of the virtual time horizon of our modified scheme are governed by the equation

$$\frac{\partial\hat{\tau}}{\partial t} = -\gamma\hat{\tau} + \frac{\partial^2\hat{\tau}}{\partial x^2} - \lambda\left(\frac{\partial\hat{\tau}}{\partial x}\right)^2 + \eta(x,t) \quad (3)$$

where $\hat{\tau}(x,t)$ is the coarse-grained surface-height fluctuation measured from the mean, and the temporal and spatial derivatives are the coarse-grained interpretations of finite differences. Similarly, $\eta(x,t)$ is a coarse-

Fig. 1. Snapshots of the virtual time horizons in the steady state for the one-site-per-PE conservative PDES scheme with $N_{PE} = 10,000$ after $t = 10^6$ parallel algorithmic steps for (**A**) the original algorithm ($p = 0.00$) and (**B**) the modified one ($p = 0.10$) with quenched random connections. The vertical scale is the same in (A) and (B).


Fig. 2. Steady-state structure factor of the virtual time horizon for $p = 0.10$. In addition to ensemble averages over 100 realizations of the random links (solid symbols), single realizations (the same open symbols) are also shown. The inset shows a magnified view of $1/S(k)$ versus k^2 for small k for the largest system. The straight line is the best linear fit used to determine the correlation length.

grained noise, delta-correlated in space and time. The coefficients γ and λ carry the details of the coarse-graining procedure. In particular, $\gamma(p)$ is nonzero for all nonzero values of p and $\gamma(p) \rightarrow 0$ as $p \rightarrow 0$ (i.e., in this limit we recover the KPZ equation for the original scheme). The morphological properties of the surface (the virtual time horizon) governed by Eq. 3 are fundamentally different from the one-dimensional KPZ surface for all $p \neq 0$. The algorithmic rules extended to include the quenched random connections introduce a relaxational term (the first term in Eq. 3) in the evolution of the virtual time horizon. This term converts the system into a mean-field growth model, whose time horizon is macroscopically smooth (i.e., has a finite width).

To support the coarse-graining arguments, we simulated the exact stochastic growth process defined by our modified conservative algorithmic rules. Snapshots of the progress of the simulation (i.e., the virtual time horizon) are shown in Fig. 1 for the original ($p = 0.00$) and modified ($p = 0.10$) cases. The structure of the virtual time horizon can be best understood from the steady-state structure factor $S(k) = \langle\tilde{\tau}_k\tilde{\tau}_{-k}\rangle/N_{PE}$, where $\tilde{\tau}_k$ is the spatial Fourier transform of the surface fluctuations (24). The presence of the relaxational (first) term in the stochastic growth equation, Eq. 3, implies that $\lim_{k\rightarrow 0} S(k) < \infty$, that is, there are no large-amplitude, long-wavelength modes in the surface. Consequently, the width $\langle w^2\rangle_{N_{PE}} = (1/N_{PE})\Sigma_{k\neq 0}S(k)$ is also finite.

If we consider only the linear terms in Eq. 3, we obtain

$$S(k) \propto \frac{1}{\gamma + k^2} \quad (4)$$

In this approximation, the lateral correlation length ξ of the surface fluctuations is $1/\gamma^{1/2}$; that is, it is finite for all $p \neq 0$. From Eq. 4 it also follows that the local slopes of the virtual time horizon remain short-range correlated and the utilization $\langle u(\infty)\rangle_{N_{PE}}$ (average progress rate) is a nonzero constant in the limit $N_{PE} \rightarrow \infty$, just like in the original conservative scheme. The behavior of the structure factor $S(k)$ indicates (Fig. 2) that it approaches a finite value as $k \rightarrow 0$. Further, the linearized version of the theory seems to work well in the small-k regime (Fig. 2, inset), yielding a finite correlation length ($\xi \approx 26$ for $p = 0.10$) for a long chain of PEs.

If an infinitesimally small p is chosen, the virtual time horizon becomes macroscopically smooth with a finite width. At the same time, the utilization is reduced by only an infinitesimal amount as a result of the occasional extra checking with the random neighbors. This trade-off is substantially favorable for the conservative PDES scheme: By giving up an infinitesimally small fraction of the utilization, the width is reduced from infinity to a finite value, making measurement and data management scalable under the conservative PDES scheme. For example, for $p = 0.10$, $\langle w^2(\infty)\rangle_{N_{PE}}$ is reduced from infinity to
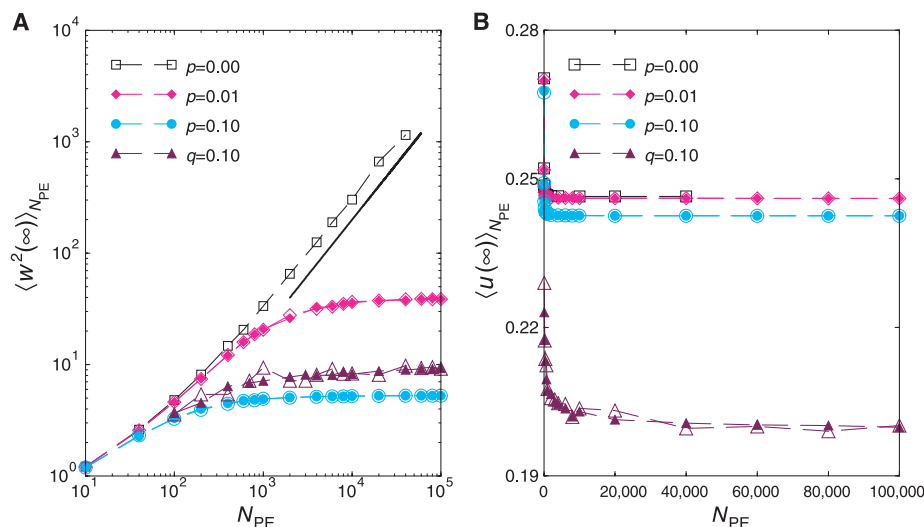
A

B

Fig. 3. (A) Average steady-state width of the virtual time horizon for various values of p for the one-site-per-PE PDES scheme. In addition to ensemble averages over 10 realizations of the random links (solid symbols), a single realization is also shown (the same open symbols). The solid straight line represents the asymptotic one-dimensional KPZ power-law divergence with roughness exponent $\alpha = \frac{1}{2}$ for the $p = 0$ case. Note the log-log scales. The $q = 0.10$ data set corresponds to the case when only 10% of the PEs have random links and those are checked at every step. **(B)** The steady-state utilization (fraction of nonidling PEs) for the same cases as in (A).

about 5.25 (Fig. 3A) while $\langle u(\infty)\rangle_{N_{PE}}$ is reduced by only about 1.6% (Fig. 3B). One can also observe the clear self-averaging property for both global observables, the width and the utilization (Fig. 3). From a broader statistical physics viewpoint, one can ask whether a small, nonzero fraction q of the PEs with random links (checked, e.g., at every step) is sufficient to control the width. Although this choice clearly weakens load balancing and the utilization (Fig. 3B), our results in Fig. 3A and recent work on the closely related XY-model on a small-world network (*25*) suggest that a finite width is achieved. There is growing evidence that systems without inherent frustration exhibit mean-field characteristics when the original short-range interaction topology is modified to a small-world network (*23, 25–27*).

The generalization when random links are added to a higher-dimensional underlying regular lattice is clear: Because the one-dimensional case with random links is governed by the mean-field equation, in higher dimensions it will be even more so [i.e., the critical dimension (*18*) of the model with random links is less than 1]. The generalization for the many-sites-per-PE case also follows from universality arguments: Without the random connections, there is an additional fast-roughening phase for early times when the evolution of the time horizon corresponds to random deposition (*18*). Subsequently, it will cross over to the KPZ growth regime and finally saturate. With many sites per PE, the typically desired and efficient way of implementing PDES (*7, 15*), the saturation value of the width can become extremely large. This underscores the

importance of implementing the additional synchronizations through random links to suppress the roughness of the time horizon.

References and Notes
1. R. Fujimoto, *Commun. ACM* **33**, 30 (1990).
2. D. M. Nicol, R. M. Fujimoto, *Ann. Oper. Res.* **53**, 249 (1994).
3. B. D. Lubachevsky, *Bell Labs Tech. J.* **5** (April–June), 134 (2000).
4. A. G. Greenberg *et al.*, in *Proceedings of the 8th Workshop on Parallel and Distributed Simulation (PADS'94), Edinburgh, UK, 1994* (Society for Computer Simulation, San Diego, CA, 1994), pp. 187–194.
5. E. Deelman, B. K. Szymanski, T. Caraco, in *Proceedings of the 28th Winter Simulation Conference* (Association for Computing Machinery, New York, 1996), pp. 1191–1198.
6. D. M. Nicol, in *Proceedings of the SCS Multiconference on Distributed Simulation* (Society for Computer Simulation, San Diego, CA, 1988), vol. 19, pp. 141–146.
7. G. Korniss, M. A. Novotny, P. A. Rikvold, *J. Comput. Phys.* **153**, 488 (1999).
8. G. Korniss, C. J. White, P. A. Rikvold, M. A. Novotny, *Phys. Rev. E* **63**, 016120 (2001).
9. P. M. A. Sloot, B. J. Overeinder, A. Schoneveld, *Comput. Phys. Commun.* **142**, 76 (2001).
10. D. R. Jefferson, *ACM Trans. Prog. Lang. Syst.* **7**, 404 (1985).
11. P. Bak, C. Tang, K. Wiesenfeld, *Phys. Rev. Lett.* **59**, 381 (1987).
12. A. G. Greenberg, S. Shenker, A. L. Stolyar, *Performance Eval. Rev.* **24**, 91 (1996).
13. K. M. Chandy, J. Misra, *Commun. ACM* **24**, 198 (1981).
14. B. D. Lubachevsky, *Complex Syst.* **1**, 1099 (1987).
15. ———, *J. Comput. Phys.* **75**, 103 (1988).
16. G. Korniss, Z. Toroczkai, M. A. Novotny, P. A. Rikvold, *Phys. Rev. Lett.* **84**, 1351 (2000).
17. G. Korniss, M. A. Novotny, Z. Toroczkai, P. A. Rikvold, in *Computer Simulation Studies in Condensed Matter Physics XIII*, D. P. Landau, S. P. Lewis, H.-B. Schüttler, Eds., vol. 86 of *Springer Proceedings in Physics* (Springer-Verlag, Berlin, 2001), pp. 183–188.
18. A.-L. Barabási, H. E. Stanley, *Fractal Concepts in Surface Growth* (Cambridge Univ. Press, Cambridge, 1995).
19. T. Halpin-Healy, Y.-C. Zhang, *Phys. Rep.* **254**, 215 (1995).
20. M. Kardar, G. Parisi, Y.-C. Zhang, *Phys. Rev. Lett.* **56**, 889 (1986).
21. Z. Toroczkai, G. Korniss, S. Das Sarma, R. K. P. Zia, *Phys. Rev. E* **62**, 276 (2000).
22. D. J. Watts, S. H. Strogatz, *Nature* **393**, 440 (1998).
23. M. E. J. Newman, *J. Stat. Phys.* **101**, 819 (2000).
24. Here we define the Fourier transform of the surface fluctuations as
$$\tilde{\tau}_k = \sum_{j=1}^{N_{PE}} [\exp(-ikj)](\tau_j - \bar{\tau})$$
where $k = 2\pi n/N_{PE}$ and $n = 0, 1, 2, \ldots, N_{PE} - 1$.
25. B. J. Kim *et al.*, *Phys. Rev. E* **64**, 056135 (2001).
26. A. Barrat, M. Weigt, *Eur. Phys. J. B* **13**, 547 (2000).
27. M. Gitterman, *J. Phys. A* **33**, 8373 (2000).
28. We thank G. Istrate and Z. Rácz for discussions. Supported by NSF grants DMR-0113049 and DMR-9981815, Research Corporation grant RI0761, and U.S. Department of Energy grant W-7405-ENG-36 (Z.T.).

A Gate-Controlled Bidirectional Spin Filter Using Quantum Coherence

J. A. Folk,[1,2]* R. M. Potok,[1] C. M. Marcus,[1]* V. Umansky[3]

We demonstrate a quantum coherent electron spin filter by directly measuring the spin polarization of emitted current. The spin filter consists of an open quantum dot in an in-plane magnetic field; the in-plane field gives the two spin directions different Fermi wavelengths resulting in spin-dependent quantum interference of transport through the device. The gate voltage is used to select the preferentially transmitted spin, thus setting the polarity of the filter. This provides a fully electrical method for the creation and detection of spin-polarized currents. Polarizations of emitted current as high as 70% for both spin directions (either aligned or anti-aligned with the external field) are observed.

The ability to controllably generate and detect electron spin in mesoscopic systems is important for the field of spintronics and is also a powerful tool for investigating basic properties of spin in coherent electronic systems. The development of tunable spin filters