Statistical Properties of the Simulated Time Horizon in Conservative Parallel Discrete-Event Simulations

G. Korniss,[†] M.A. Novotny,[‡] A.K. Kolakowska,[‡] and H. Guclu[†]

[†] Department of Physics, Applied Physics, and Astronomy
Rensselaer Polytechnic Institute
110 8th Street, Troy, NY 12180-3590

[‡] Department of Physics and Astronomy
P.O. Drawer 5167
Mississippi State, MS 39762-5167

korniss@rpi.edu, man40@ra.msstate.edu, alicek@erc.msstate.edu, gucluh@rpi.edu

ABSTRACT

We investigate the universal characteristics of the simulated time horizon of the basic conservative parallel algorithm when implemented on *regular lattices*. This technique [1, 2] is generically applicable to various physical, biological, or chemical systems where the underlying dynamics is asynchronous. Employing direct simulations, and using standard tools and the concept of *dynamic scaling* from non-equilibrium surface/interface physics, we identify the universality class of the time horizon and determine its implications for the asymptotic *scalability* of the basic conservative scheme. Our main finding is that while the simulation converges to an asymptotic nonzero rate of progress, the statistical width of the time horizon diverges with the number of PEs in a power law fashion. This is in contrast with the findings of Ref. [3]. This information can be very useful, e.g., we utilize it to understand optimizing the size of a moving "time window" to enforce memory constraints.

Keywords

conservative parallel discrete-event simulation, scalability, non-equilibrium surface growth, Monte Carlo, stochastic processes

1. INTRODUCTION

Massively parallel simulations of complex systems with asynchronous dynamics, often referred to as parallel discrete-event simulations (PDES), are far from new for computer scientists. On the other hand, despite the fact that PDES has a long history as far as scalability and applications are concerned [4, 5], very few of the PDES techniques have filtered through to the physics community. Even the simplest random-site update Monte Carlo schemes [6], where update attempts correspond to Poisson arrivals, were long believed to be inherently serial (at least in the physics community). Since a large number of physicists are engaged in the modeling and large-scale simulation of complex systems, ranging from magnetization dynamics and epidemic models to market models and internet traffic (which are perfect candidates for PDES) it is important to learn the algorithms and schemes, which have long been considered standard among computer scientists working on simulations. In this regard, Lubachevsky's work [1, 2] was rather illuminating, by illustrating how to apply efficiently the conservative scheme [7, 8] to the Ising model on a lattice with Glauber dynamics [9]. In this basic model for ferromagnets, the discrete events are the possible spin-flips, and the actual flipping probability depends on the energetics of the nearest-neighbor magnetic interaction [6]. The nearest-neighbor interaction implies that in order to ensure causality, processing elements (PE) need to exchange their local simulated (or "virtual") times only with "neighboring" processing elements in the virtual topology.

Systems which can be modeled on a lattice (regular grid) with *short-ranged* interactions constitute a large class among physics, biology, and engineering applications. For these systems the conservative scheme can be highly efficient. For example, it was implemented for modeling magnetization switching [10] and the dynamic phase transition in highly anisotropic thin-film ferromagnets [11]. Communication times between PEs and possible idling due to the conservative synchronization protocol can be greatly suppressed by each PE carrying a large block of sites (spins) [1, 2], yielding encouraging efficiencies and utilizations (fraction of non-idling PEs) [10]. Since there is a finite number of processing elements on any architecture, one can typically be satisfied with these "experimental" observations.

There is an obvious question, however, lurking in the background for all massively parallel schemes: how does the utilization behave in the asymptotic limit where the number of PEs goes to infinity? While this question for all practical purposes might seem academic, it truly lies at the heart of any parallel-discrete event simulation scheme. From a "statistical physics" standpoint, one would like to know, e.g., what the leading behavior and/or corrections are of the uti-

SAC 2002, Madrid, Spain

lization as the "infinite" number of PE limit is *approached* through large but finite number of PEs. The answers to these questions can be very useful to model and understand the performance of the algorithm in the limit of large number of PEs. In statistical physics, in the context of interacting systems with many degrees of freedom, these behaviors are referred to as "finite-size effects". Obtaining quantitative answers for these types of questions are far from trivial, and innocent looking assumptions [3] in order to obtain an analytically computable scalability model can yield misleading predictions.

These are exactly the types of scalability questions where statistical physics may have a lot to offer and contribute. We believe that the tools and frameworks of modern statistical physics, in particular those of non-equilibrium interface/surface growth [12, 13, 14], can be very helpful in analyzing and understanding the asymptotic scalability properties of parallel discrete-event simulation schemes. To this end, one must look at the simulation scheme itself as an interacting system of individual PEs where the synchronization rules among the PEs constitute the effective interaction. The most important variable to track is the set of local simulated (or virtual) times. The evolution of this simulated time horizon, in particular, its average progression rate and statistical spread, will determine the scalability properties of the corresponding PDES scheme.

We illustrate the potential of this approach by tackling fundamental scalability issues for the conservative scheme. We show how the evolution of the simulated time horizon for these schemes fits into the general picture of non-equilibrium surface/interface growth. Then we demonstrate that the evolution of the time horizon exhibits kinetic roughening, observable in numerous other artificial and natural growth processes, including molecular beam epitaxy, electrochemical deposition, fluid flow in porous media, and growth of bacterial colonies [12, 13].

We must note that similar approaches, namely finding analogies between the evolution of the time horizon and that of known *physical* systems proved to be rather helpful for optimistic schemes as well. There is some evidence [15, 16, 17] that the time horizon in rollback-based schemes can exhibit self-organized criticality and power-law spatio-temporal correlations, which can be crucial to extract the scalability properties. Also, analytic models for the event horizon for optimistic schemes were studied by others [18, 19, 20, 21, 22].

2. BASIC CONSERVATIVE SCHEME ON REGULAR LATTICES

We consider a d-dimensional hypercubic regular lattice topology where the underlying physical system has only nearest-neighbor interactions (e.g., Ising model with Glauber spin-flip dynamics). In this paper we consider the case of simple Poisson asynchrony. Update attempts at each site are independent Poisson processes with the same rate (thus the random simulated time increments between two successive update attempts are exponentially distributed) and are also *independent of the state of the underlying physical system*. The consequence of the latter is that the evolution of the simulated time horizon completely *decouples* from the behavior and evolution of the underlying physical system. This not only makes the evolution of the time horizon accessible to direct simulations and the scalability analysis easier, but it is indeed applicable to a large number of physics models and other applications.

For simplicity, we discuss first the "worst-case" scenario, when each PE carries one site (e.g., one spin). In this basic conservative scheme [1, 2], each PE generates its own local simulated time for the next update attempt. The set of local simulated times for N PEs, $\{\tau_i(t)\}_{i=1}^{N}$, constitutes the simulated time horizon. Here t is the discrete number of parallel steps simultaneously performed on each PE directly related to real/wall-clock time, or if the architecture operates in an asynchronous execution mode, t is simply the continuous real time. On a regular d-dimensional hypercubic lattice $N{=}L^d$, where L is the linear size of the lattice. In physics applications one typically specifies the initial configuration (i.e., at $\tau{=}0$) of the underlying physical system. This translates to $\tau_i(0){=}0$ for every site for the initial condition of the parallel simulation. Then at each parallel update, *only* those PEs for which the local simulated time is *not greater* than the local simulated times of their nearest neighbors, can increment their local time by an exponentially distributed random amount, $\eta_i(t)$. Without loss of generality we take iid exponential variables with mean one, $\langle\eta_i(t)\rangle{=}1$. The other PEs must idle. Due to the continuous nature of the random simulated times, for $t > 0$ the probability of equal time updates for any two sites is of measure zero. The comparison of nearest neighbor simulated times and idling if necessary enforces causality. Also, at worst, the PE with the global minimum simulated time can make progress, so the algorithm is free from deadlock. For this basic conservative scheme, the theoretical efficiency or utilization (ignoring communication overheads) is simply the fraction of non-idling PEs. This corresponds to the density of local minima of the simulated stochastic time horizon which determines the average progress rate of the simulation. Another important aspect of the simulation is the tightness of the distribution of the local simulated times. This property can have serious effects on the "measurement part" of the algorithm, e.g. when one attempts to collect and compute simple statistics for the full underlying physical system "on the fly". Therefore, one must determine the *statistical spread* (width) of the time horizon as was pointed out in Ref. [3].

To obtain an analytically tractable scalability model, Greenberg et.al [3] introduced the K-random model. Here at each update attempt, PEs compare their local simulated times to the local simulated times of K *randomly* chosen PEs (rechosen at every update attempt). They showed that in the $t{\to}\infty$, $N{\to}\infty$ limit the average rate of progress of the simulation converges to a *non-zero* constant, $1/(K+1)$. Further, they also showed that the evolution of the time horizon converges to a traveling wave solution described by a *finite width* of the distribution of the local times. Finally, they suggested that the qualitative properties of the K-random model are universal and hold for regular lattice models as well. As we will illustrate, the underlying connection topology has crucial effects on the "universal" behavior of the evolution of the time horizon and their conjecture [3] for the width for *regular lattices* does not hold.

3. NON-EQUILIBRIUM SURFACE GROWTH AND KINETIC ROUGHENING

The conservative synchronization protocol together with

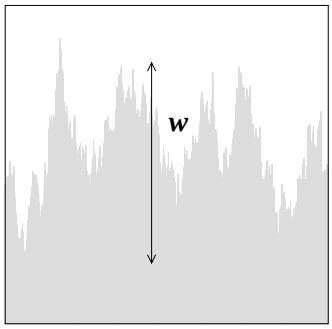

Figure 1: Snapshot configuration of the actual simulated time horizon for the one-dimensional one site per PE topology with L=10,000 PEs. The time horizon propagates "upwards" in the figure. w indicates the typical spread of the time horizon.

the communication topology of the PEs fully specify the "microscopic" dynamics of the growth process associated with the evolution of the time horizon. A snapshot of the evolving and fluctuating time horizon is shown in Fig. 1. In order to identify the universal characteristics of this time "surface", one must define the observables, which carry essential features of this genuinely non-equilibrium growth process. In the basic conservative scheme (one site per PE) the average progress rate (utilization) is the average density of local minima of the evolving and fluctuating time horizon, $\langle u(t)\rangle_N$. The other important quantity is the average width (spread) of the time horizon

$$\left\langle w^2(t)\right\rangle_N = \left\langle \frac{1}{N}\sum_{i=1}^{N}\left[\tau_i(t)-\bar{\tau}(t)\right]^2 \right\rangle , \qquad (1)$$

where $\bar{\tau}(t)=(1/N)\sum_{i=1}^{N}\tau_i(t)$, $N=L^d$. Here $\langle\ldots\rangle$ denotes an ensemble average, i.e. an average over many independent simulations. In the case of translational invariance (e.g. a ring or a torus) the width $\langle w^2(t)\rangle$ is just $\langle[\tau_i(t)-\bar{\tau}(t)]^2\rangle$ independent of site i. The behavior of this quantity alone typically captures and identifies the universality class of the non-equilibrium growth process [12, 13, 14].

The conservative synchronization protocol introduces (although short-range) "interactions" between PEs. As a result, the *correlation length* between local simulated times begins to grow as the simulation evolves. In principle, one can extract this lateral correlation length from the equal-time "height-height" correlation function $\langle\tau_i(t)\tau_j(t)\rangle$, or from its Fourier transform, the structure factor [12, 13, 14].

To characterize this growth process, we can exploit the analogy which exists between the evolution of the time horizon and kinetic roughening in general non-equilibrium surfaces. Then we will build on the wealth of results obtained for various non-equilibrium surfaces in the last two decades [12, 13, 14]. In these systems, in general, the lateral correlation length between sites grows as a power law, $\xi\sim t^{1/z}$, where z is the *dynamic exponent*. Having a finite system, however, the correlation length cannot grow beyond the linear system size, L. Thus, one can formally extract the *crossover time* $t_\times\sim L^z$. For very early times, mostly microscopic details of the dynamics influence the width. For intermediate times, where $t\ll t_\times$ still, the width typically grows as a power law $\langle w^2(t)\rangle_L\sim t^{2\beta}$, where β is called the *growth exponent*. For late times, $t\gg t_\times$, the width *saturates* for any finite system size. In this regime the surface reaches a steady-state evolution and the fluctuations about the mean are stationary. The saturation or steady-state value of the width, however, scales as a power law with the system size, $\langle w^2(\infty)\rangle_L\sim L^{2\alpha}$, where α is the *roughness exponent*. To summarize the temporal and system-size dependence of the width (except for the very early times), we have

$$\left\langle w^2(t)\right\rangle_L \sim \begin{cases} t^{2\beta} & \text{if } t\ll t_\times \\ L^{2\alpha} & \text{if } t\gg t_\times \end{cases} . \qquad (2)$$

The time horizon through its progress exhibits exactly the above scaling behavior [Fig. 2(a)]. This type of temporal and system-size scaling is consistent with the dynamic scaling hypothesis [12, 23] and can be expressed through the Family-Vicsek scaling relation

$$\left\langle w^2(t)\right\rangle_L = L^{2\alpha} f(t/L^z) \qquad (3)$$

together with the important $\alpha=\beta z$ scaling law. Note that the scaling function $f(x)$ depends on t and the linear system size L only through the specific combination t/L^z, reflecting the importance of the crossover time $t_\times$. For small values of its argument $f(x)$ behaves as a power law, while for large arguments it approaches a constant

$$f(x) \sim \begin{cases} x^{2\beta} & \text{if } x\ll 1 \\ \text{const.} & \text{if } x\gg 1 \end{cases} . \qquad (4)$$

The existence of the above scaling function implies that if one plots the rescaled variables $\langle w^2(t)\rangle_L/L^{2\alpha}$ vs t/L^z, then curves for different system sizes collapse for intermediate and late times. We confirmed this data collapse for the simulated time horizon [Fig. 2(b)].

The enormous number of "microscopically" different non-equilibrium growth processes typically fall into a fewer number of (universality) classes. The systems (ranging from biological and chemical systems, surface deposition, or here in this paper, the time horizon), which belong to the same universality class, despite having very different microscopic growth mechanisms, exhibit the same macroscopic characteristics (e.g., for the width $\langle w^2\rangle$), and are described by the same set of exponents α, β (and $z=\alpha/\beta$). Given the microscopic rules for the evolution of the simulated time horizon (that is, the local minima are incremented by an exponentially distributed random number) the most important task is to determine which universality class it belongs to. Recently it was shown by directly simulating the time horizon and also by deriving a coarse-grained stochastic equation of motion for its evolution [24], that the evolution of the time horizon on *regular lattices* belongs to the Kardar-Parisi-Zhang (KPZ) universality class [25]. In the next section we discuss the implication of this finding on the scalability of the conservative scheme.

4. SCALING AND SCALABILITY

We discuss in detail the one dimensional case (N=L) with periodic boundary conditions (i.e. a ring). In this case the KPZ exponents are known *exactly* and fully agree with our simulations. The evolution of the time horizon was also studied in two and three dimensions and the phenomenon of kinetic roughening was observed [26]. In higher dimensions, in general, exact solution is not available and the simulation

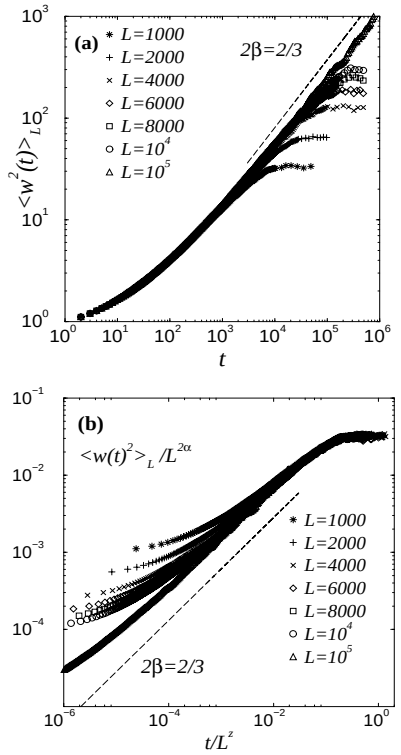

Figure 2: (a) Kinetic roughening of a growing surface illustrated through the actual simulated time horizon for the one-dimensional one site per PE basic conservative scheme. Note the log-log scales, indicating the power-law growth of the width before saturation. The dashed line corresponds to a power law with the exact KPZ exponent 2β=2/3. (b) The same behavior as in (a), using the rescaled variables to demonstrate the dynamic scaling hypothesis, Eq. (3).

of the corresponding surfaces is a challenging task [12]. Currently available best estimates for the roughness exponent (which is the most important for the steady-state performance) are in the range of α=0.20 − 0.40 and α=0.08 − 0.3 in two and three dimensions respectively.

Direct simulation results for the utilization (average rate of progress) for various system sizes are shown in Fig. 3(a). $\langle u(t)\rangle_L$ decreases monotonically with time towards a long-time asymptotic limit well separated from zero. In one dimension α=1/2, β=1/3 (and z=3/2 due to the $\alpha=\beta z$ scaling law). The steady-state KPZ surface in one dimension is essentially a random-walk profile. At coarse-grained length scales the local slopes become independent, yielding a *non-zero* average density of local minima, i.e., a non-zero average rate of progress of the simulation in the $L\to\infty$ limit in the steady-state. It was shown [27] that the finite-size corrections of the density of local minima (and hence the rate of progress) is of the form

$$\langle u(\infty)\rangle_L \simeq \langle u(\infty)\rangle_\infty + \frac{\text{const.}}{L} \ . \qquad (5)$$

This form for the finite-size effects is in full agreement with the simulations [Fig. 3(b)]. Using the above finite-size behavior for the utilization, one can extrapolate to determine its "infinite number of PE" value, yielding $\langle u(\infty)\rangle_\infty$=0.246461(7) for the one-site-per PE nearest-neighbor interaction case. Krug and Meakin obtained [28] universal finite-size effects for the growth rate of generic KPZ-like processes

$$\langle u(\infty)\rangle_L \simeq \langle u(\infty)\rangle_\infty + \frac{\text{const.}}{L^{2(1-\alpha)}} \ , \qquad (6)$$

which can be used to estimate the utilization (average rate of progress) for higher dimensions as well. While the density of local minima is based on "microscopic" measures, whether its asymptotic value vanishes or not, is fully governed by macroscopic characteristics and the corresponding universality class [24, 27]. For the KPZ class this asymptotic value is *non-zero*.

The above findings for the density of local minima, which determines the average rate of progress of the simulation, imply that the "simulation part" of the conservative scheme is scalable. That is, if we run the simulation for long times, the average progress rate approaches a constant. There is some disturbing implication, however, of the kinetic roughening exhibited by the time horizon. In the steady state the width (spread) of the simulated time horizon *diverges* with the number of PEs

$$\langle w^2(\infty)\rangle_L \sim L^{2\alpha} \ . \qquad (7)$$

This scaling behavior for large L is also confirmed by simulations [Fig. 3(c)] and it is contrary to the conclusions of Ref. [3]. This property adds an additional difficulty for collecting statistics (e.g. to perform simple average) "on the fly" through the course of the simulation. The diverging width means that the memory requirement *per PE*, for temporarily storing (buffering) data, diverges as we increase the number of PEs. In this sense we may call the "measurement part" of the bare conservative scheme asymptotically not scalable. Thus, in an actual application, the programer must implement some global synchronization or a moving "window" with respect to the global minimum of the time horizon.

Along these lines of questioning, one should be interested in the *extremal* fluctuations of the time horizon. Namely, what is the typical size of the fluctuations above and below the mean. We have started to study these quantities in detail, $\Delta_{\max}(t)\equiv(\tau_{\max}(t)-\bar\tau(t))$ and $\Delta_{\min}(t)\equiv(\bar\tau(t)-\tau_{\min}(t))$, where $\tau_{\max}(t)$ and $\tau_{\min}(t)$ are the *global* maximum and minimum simulated times among L PEs, respectively. We found that in the steady state

$$\left\langle \Delta^2_{\max}(\infty)\right\rangle_L \sim \left\langle \Delta^2_{\min}(\infty)\right\rangle_L \sim L^{2\alpha} \ , \qquad (8)$$

i.e., they scale the *same way* the width $\langle w^2(\infty)\rangle_L$ does [Fig. 3(c)]. This shouldn't come as a surprise, since this surface is highly correlated, dominated by long-wavelength fluctuations spreading over macroscopic length scales. This finding, again, is consistent with the extremal fluctuations found for general KPZ surfaces [29].

The above universal characteristics hold for the sensible and more efficient many site per PE case and/or when the interaction and the corresponding communication pattern extends beyond nearest-neighbor PEs (but still short-ranged with a finite cutoff). For the many site per PE case , however, the saturation will occur at a later time and the time horizon exhibits a substantially larger width.

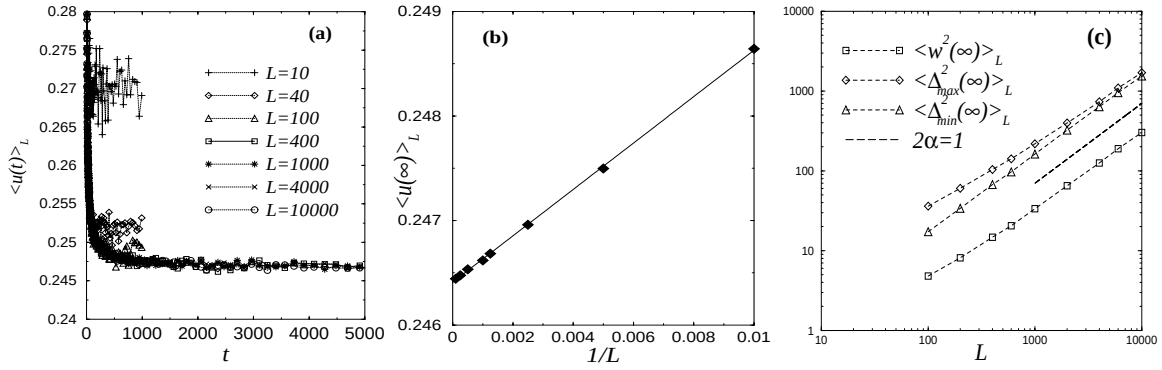

Figure 3: (a) Time dependent utilization for various system sizes. (b) Steady-state utilization (average rate of progress) as function of the inverse system size. (c) Steady-state average width and extremal fluctuations of the time horizon. Note the log-log scales. The dashed line corresponds to a power law with the exact KPZ exponent 2α=1. All three graphs for the one-dimensional one site per PE basic conservative scheme.

5. THE EFFECTS OF A MOVING WINDOW

As seen in the sections above, a statistical analysis of the growing virtual time interface in conservative asynchronous PDES shows that in the steady state, the average utilization $\langle u\rangle_L$ remains finite and $\langle w^2\rangle_L$ diverges as $L\to\infty$. This was demonstrated for one volume element M_V per PE (one site per PE). A higher efficiency can be obtained by increasing M_V. Here the "volume element" is measured in terms of the "nearest-neighbor interaction distance". For simplicity, we continue to consider only the linear chain (with periodic boundary conditions). For M_V=2, each PE will only need to check the neighbor to the left *or* to the right, and can advance its local virtual time if its local virtual time is behind that of the checked neighboring PE's local virtual time. For $M_V > 2$, with probability $(M_V - 2)/M_V$ the PE does not need to check any neighboring PE, and with probability $2/M_V$ it must check τ on one of its neighboring PEs. Because of universality arguments, for all finite M_V one expects that still the average utilization $\langle u\rangle_L$ remains finite and $\langle w^2\rangle_L$ diverges as $L\to\infty$. This is because changing M_V changes only local behavior of the model, but such changes will not affect the asymptotic behavior of the growing surface. Fig. 4 shows $\langle u(t)\rangle$ averaged over 100 trials as a function of t for different L and M_V. Note that for M_V=2 the average utilization is slightly less than 1/2, while for M_V=100 it saturates at about 0.9 at large times, and the differences for L between 10 and 10^3 are within the statistical errors.

The convergence of $\langle u(\infty)\rangle_L$ to a finite value as $L\to\infty$, reflects positively on the ability to efficiently implement this type of PDES, in other words the "simulation portion" of the algorithm is scalable. However, the divergence of $\langle w^2(\infty)\rangle_L$ with L means that the "measurement portion" of the PDES algorithm is not scalable. A standard way [1, 2, 4] to try to remedy this problem is to impose a moving window of width Δ. At each step the *global* minimum of the virtual time horizon is calculated, this is called the "global virtual time", which we label as $\tau_{\rm GVT}$. Then the i-th PE can only advance its value τ_i if it could do so from the previous rules and if $\tau_i \le \tau_{\rm GVT} + \Delta$. This prevents the virtual time horizon from roughening. Note that this is at the cost of a *global* calculation to find $\tau_{\rm GVT}$ at each iteration — an expensive

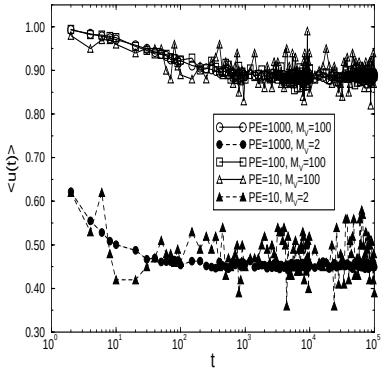

Figure 4: The average utilization $\langle u\rangle$ is shown as a function of t for different numbers of processing elements (PE) and and different M_V.

calculation for large L. Examples of surfaces for different values of Δ for L=1000 are shown in Fig. 5 for two different times. Clearly, using a finite value of Δ keeps the interface width from diverging. This is very useful in actual simulations.

However, the question must now be asked what effect a finite Δ has on the average utilization. Fig. 6 shows for Δ=10 that for small L the average utilization does not change significantly from the case $\Delta=\infty$. However, for L larger than about 10^2 the value of $\langle u\rangle$ decreases. The values in Fig. 6 are averages over 10^3 independent sequences of random numbers and averaged over time steps between 20,000 and 100,000. Error estimates rely on having the second moment of the probability distribution of u approximated reasonably well by the 10^3 independent trials. The error bars in Fig. 6 represent the probable error that was estimated from this sample size. Consequently, the error bars should be regarded as an approximate lower bound on the error. It is seen that $\langle u\rangle$ decreases for Δ=10 fixed, and the utilization seems to decrease in a way that it goes toward zero as $L\to\infty$. Consequently, introducing a finite Δ makes the "measurement phase" of the PDES scalable, but it is not yet clear whether the "simulation phase" of the PDES remains asymptotically

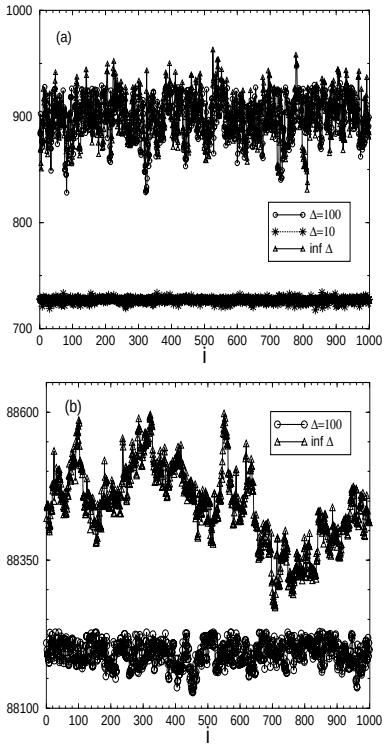

Figure 5: Examples of virtual time horizons for $L = 1000$ and $M_V = 100$ are shown for three values of the moving window width Δ. (a) is for $t = 10^3$, while (b) is for $t = 10^5$. Note that in (b) the values of $\Delta = 10$ is not shown since it is has an average value of $\bar{\tau} = 72100$.

scalable.

6. CONCLUSIONS AND OUTLOOK

We studied the statistical properties of the basic conservative parallel scheme for regular lattice topologies. We found that the evolution of the simulated time horizon belongs to the well known KPZ universality class of non-equilibrium surfaces. This type of growth is characterized by a *non-zero* density of local minima, i.e., implying a non-zero rate of propagation in the infinite PE limit. We also determined the asymptotic finite-size corrections to this constant when the number of PEs is large. Thus, the "simulation" part of the algorithm is scalable. Further, we showed that the spread (width) of the time horizon approaches a finite constant for *finite* number of PEs, but this constant *diverges* in the infinite number of PE limit in a power-law fashion. The same holds for the extremal fluctuations above and below the mean. We may refer to this "macroscopic" roughness of the simulated time horizon as the "measurement part" of the bare conservative scheme not being scalable. That is, there is some extra difficulty associated with measurement taking "on the fly". Intermittent data on each PE has to be stored until all PE reaches the simulated time instant at which some statistics collection, e.g., simple averaging over the full physical application is to be performed. The diverging spread of the time horizon, however, implies diverging storage need for this purpose on every PE. Thus, the programmer has to implement some global synchronization, or windowing technique to limit the spread of the simulated time horizon in order not to exceed the memory constraint. By knowing exactly the finite-size dependence of the spread, one can determine the optimal time between global synchronizations or the optimal window size.

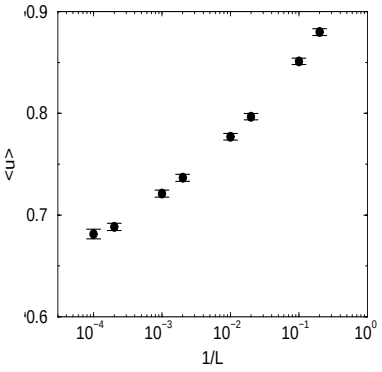

Figure 6: The average steady-state utilization for $\Delta = 10$ is shown as a function of L, averaged over times between 2×10^4 and 10^5.

Our findings are universal in the sense that they hold for any *short-range* "interaction" topology for PEs on regular lattices. Also they are valid in the case when each PE carries a block of sites. The *asymptotic* scaling behavior is again governed by the KPZ exponents, in such a way that for larger and larger blocks, there is a crossover from the almost "random deposition" [12] to KPZ-like growth at a later and later time.

Regarding the width of the time horizon of the bare conservative scheme, we arrived at a very different conclusion than that of Greenberg et.al. [3]. They suggested that convergence to a traveling-wave solution in the $t\to\infty$, $N\to\infty$ limit is universal and applicable for regular lattices as well. In obtaining this result they made the assumption that replacing the "interaction" between nearest-neighbor PEs on a regular grid with the same interaction between *randomly* chosen PEs does not change the universality class of the time horizon. It does. Comparing the local time for each PE to K randomly chosen others essentially turns the model in to a mean-field-like one where the time surface is short-range correlated and has a finite width in the infinite number of PE limit. However, using their idea, and realizing how crucial the communication topology of the PEs, we currently investigate how to turn the original conservative scheme on *regular lattices* into a fully scalable one, where both the "simulation" and the "measurement" parts are scalable, without the need for any global synchronization or windowing technique.

7. ACKNOWLEDGMENTS

We thank P.A. Rikvold, Z. Toroczkai, B.D. Lubachevsky, and Z. Rácz for useful discussions. We acknowledge the support of NSF through Grant Nos. DMR-9871455 and DMR-0113049.

8. REFERENCES

[1] B.D. Lubachevsky. Efficient Parallel Simulations of

Asynchronous Cellular Arrays. *Complex Systems* **1**, 1099–1123 (1987).

[2] B.D. Lubachevsky. Efficient Parallel Simulations of Dynamic Ising Spin Systems. *J. Comput. Phys.* **75**, 103–122 (1988).

[3] A.G. Greenberg, S. Shenker, and A.L. Stolyar. Asynchronous Updates in Large Parallel Systems. *Proc. ACM SIGMETRICS, Performance Eval. Rev.* **24**, 91–103 (1996).

[4] R. Fujimoto. Parallel Discrete Event Simulation. *Commun. of the ACM* **33**, 30–53 (1990).

[5] D.M. Nicol and R.M. Fujimoto. Parallel Simulation Today. *Annals of Operations Research* **53** 249–285 (1994).

[6] K. Binder and D.W. Heermann. *Monte Carlo Simulation in Statistical Physics. an Introduction*, 3rd ed. (Springer, Berlin, 1997).

[7] K.M. Chandy and J. Misra. Distributed Simulation: a Case Study in Design and Verification of Distributed Programs. *IEEE Trans. on Softw. Eng.* **SE-5**, 440–452 (1979).

[8] K.M. Chandy and J. Misra. Asynchronous Distributed Simulation via a Sequence of Parallel Computations. *Commun. ACM* **24**, 198–205 (1981).

[9] R.J. Glauber. Time-Dependent Statistics of the Ising Model. *J. Math. Phys.* **4**, 294 (1963).

[10] G. Korniss, M.A. Novotny, and P.A. Rikvold. Parallelization of a Dynamic Monte Carlo Algorithm: a Partially Rejection-Free Conservative Approach. *J. Comput. Phys.* **153**, 488–508 (1999).

[11] G. Korniss, C.J. White, P.A. Rikvold, and M.A. Novotny. "Dynamic Phase Transition, Universality, and Finite-Size Scaling in the Two-Dimensional Kinetic Ising Model in an Oscillating. *Phys. Rev. E* **63**, 016120 (2001).

[12] A.-L. Barabási and H.E. Stanley. *Fractal Concepts in Surface Growth.* (Cambridge University Press, Cambridge, 1995)

[13] Timothy Halpin-Healy and Yi-Cheng Zhang. Kinetic roughening phenomena, stochastic growth, directed polymers and all that. Aspects of multidisciplinary statistical mechanics. *Phys. Rep.* **254** 215–414 (1995).

[14] J. Krug. Origins of scale invariance in growth processes. *Adv. Phys.* **46**, 139-282 (1997).

[15] B.J. Overeinder. Distributed Event-driven Simulation: Scheduling Stategies and Resource Management. Ph.D. thesis, Universiteit van Amsterdam, 2000.

[16] B.J. Overeinder, A. Schoneveld, and P.M.A. Sloot. Self-organized Criticality in Optimistic Simulation of Correlated Systems. *J. Parallel and Distributed Computing*, submitted.

[17] B.J. Overeinder, A. Schoneveld, and P.M.A. Sloot. Spatio-temporal correlations and rollback distributions in optimistic simulations. *Proceedings of the 15th Workshop on Parallel and Distributed Simulation* (IEEE Comput. Soc., Los Alamitos, CA, 2001) p. 145–152.

[18] R.E. Felderman and L. Kleinrock. Bounds and Approximations for Self-Initiating Distributed Simulation without Lookahead. ACM Trans. Model. Comput. Simul. **1**, 386 (1991).

[19] D.M. Nicol. Parallel Self-Initiating Discrete-Event Simulations. ACM Trans. Model. Comput. Simul. **1**, 24 (1991).

[20] J.S. Steinman. Breathing Time Warp. in *Proceedings of the 7th Workshop on Parallel and Distributed Simulation (PADS'93)*, Vol. 23, No. 1, edited by R. Bagrodia and D. Jefferson (Soc. Comput. Simul., San Diego, CA, 1993), p. 109-118.

[21] J.S. Steinman. Discrete-Event Simulation and the Event Horizon. in *Proceedings of the 8th Workshop on Parallel and Distributed Simulation (PADS'94)*, edited by D.K. Arvind, R. Bagrodia, and J.Y.-B. Lin (Soc. Comput. Simul., San Diego, CA, 1994), p. 39-49.

[22] J.S. Steinman. Discrete-Event Simulation and the Event Horizon II: Event List Management. in P*roceedings of the 9th Workshop on Parallel and Distributed Simulation (PADS'96)*, (IEEE Comput. Soc. Press, Los Alamitos, CA, 1996), p. 170-178.

[23] F. Family and T. Vicsek. Scaling of the active zone in the Eden process on percolation networks and the ballistic deposition model. J. Phys. A **18**, L75–81 (1985)

[24] G. Korniss, Z. Toroczkai, M.A. Novotny, and P.A. Rikvold. From Massively Parallel Algorithms and Fluctuating Time Horizons to Non-equilibrium Surface Growth. *Phys. Rev. Lett.* **84**, 1351–1354 (2000).

[25] M. Kardar, G. Parisi, and Y.-C. Zhang. Dynamic Scaling of Growing Interfaces. *Phys. Rev. Lett.* **56**, 889–892 (1986).

[26] G. Korniss, M.A. Novotny, Z. Toroczkai, and P.A. Rikvold. Nonequilibrium Surface Growth and Scalability of Parallel Algorithms for Large Asynchronous Systems. in *Computer Simulation Studies in Condensed Matter Physics XIII*, Springer Proceedings in Physics, Vol. 86, editors D.P. Landau, S.P. Lewis, and H.-B. Schüttler (Springer-Verlag, Berlin, Heidelberg, 2001), p. 183–188.

[27] Z. Toroczkai, G. Korniss, S. Das Sarma, and R.K.P. Zia. Extremal Point Densities of Interface Fluctuations. *Phys. Rev. E* **62**, 276 (2000).

[28] J. Krug and P. Meakin. Universal Finite-Size Effects in the Rate of Growth Processes. *J. Phys. A* **23**, L987 (1990).

[29] S. Raychaudhuri, M. Cranston, C. Przybyla, and Y. Shapir. Maximal Height Scaling of Kinetically Growing Surfaces. *Phys. Rev. Lett.* **87**, 136101 (2001).