

Small-World Synchronized Computing Networks for Scalable Parallel Discrete-Event Simulations

Hasan Guclu¹, György Korniss¹, Zoltán Toroczkai², and Mark A. Novotny³

¹ Department of Physics, Applied Physics, and Astronomy, Rensselaer Polytechnic Institute, 110 8th Street, Troy, NY 12180, USA

² Theoretical Division and Center for Nonlinear Studies, Los Alamos National Laboratory, MS B258 Los Alamos, NM 87545, USA

³ Department of Physics and Astronomy and ERC Center for Computational Sciences, Mississippi State University, P.O. Box 5167, Mississippi State, MS 39762, USA

Abstract. We study the scalability of parallel discrete-event simulations for arbitrary short-range interacting systems with asynchronous dynamics. When the synchronization topology mimics that of the short-range interacting underlying system, the virtual time horizon (corresponding to the progress of the processing elements) exhibits Kardar-Parisi-Zhang-like kinetic roughening. Although the virtual times, on average, progress at a nonzero rate, their statistical spread diverges with the number of processing elements, hindering efficient data collection. We show that when the synchronization topology is extended to include quenched random communication links between the processing elements, they make a close-to-uniform progress with a nonzero rate, without global synchronization. We discuss in detail a coarse-grained description for the small-world synchronized virtual time horizon and compare the findings to those obtained by “simulating the simulations” based on the exact algorithmic rules.

1 Introduction

Synchronization is a fundamental problem in natural or artificial coupled multi-component systems [1]. To achieve it in an autonomous fashion can be a particularly challenging task from a system design viewpoint. In this chapter we discuss such a problem in the context of scalable Parallel Discrete-Event Simulations (PDES) [2–4]. Examples of PDES applications include dynamic channel allocation in cell phone communication network [4,5], models of the spread of diseases [6], battle-field simulations [7], and dynamic phenomena in highly anisotropic magnetic systems [8–10]. In these examples the discrete events are call arrivals, infections, troop movements, and changes of the orientation of the local magnetic moments, respectively. We focus on the basic algorithm suitable for simulating large spatially extended systems with short-range interactions and asynchronous dynamics [11,12].

In discrete-event simulations, the instantaneous local updates (discrete-events) occur in continuous time. The algorithm must faithfully and reproducibly keep track of the asynchrony of the local updates in the system’s configuration. For example standard random-sequential Monte Carlo simulations naturally produce Poisson asynchrony. In fact, such continuous-time simulations (e.g., sin-

gle spin-flip Glauber dynamics) were long believed to be inherently serial until Lubachevsky's illuminating work [11,12] on the parallelization of these simulations *without* altering the underlying dynamics. The essence of the problem is to algorithmically parallelize "physically" non-parallel dynamics of the underlying system. This requires some kind of synchronization to ensure causality. The two basic ingredients of PDES are the set of local simulated times (or virtual times [13]) and a synchronization scheme. First, a scalable PDES scheme must ensure that the average progress rate of the simulation approaches a nonzero constant in the long-time limit as the number of Processing Elements (PEs) N goes to infinity. Second, the "width" of the simulated time horizon (the spread of the progress of the individual PEs) should be bounded as N goes to infinity [14]. The second requirement is crucial for the measurement phase of the simulation to be scalable: a large width of the virtual time horizon hinders scalable data management. Temporarily storing a large amount of data on each PE (being accumulated for "on-the-fly" measurements) is limited by available memory while frequent global synchronizations can get costly for large N . Thus, one aims to devise a scheme where the PEs make a nonzero and close-to-uniform progress without global synchronization. In such a scheme, the PEs autonomously learn the global state of the system (without receiving explicit global messages) and adjust their progress rate accordingly.

As the number of PEs available on parallel architectures increases to hundreds of thousands [15], or grid-computing networks proliferate the internet [16, 17] fundamental questions of the scalability of the underlying algorithms must be addressed. The center of our interest here is to understand the effects of the "microscopic dynamics" (corresponding to the algorithmic synchronization rules) and the effects of the underlying communication network among the PEs on the evolution and the morphological properties of the virtual time horizon. We achieve this by looking at the parallel simulation itself as a complex interacting system. A similar approach was also successful to establish connection [18] between rollback-based (or optimistic) schemes [13] and self-organized criticality [19]. Our main finding is that extending the basic conservative synchronization rules [11,12] to a small-world-like [20] communication topology among the PEs results in both a finite width of the time horizon and a nonzero progress rate of the simulation [21]. Performing additional synchronizational steps through the random links at a very small rate can only reduce the average progress rate infinitesimally while the width is reduced from infinity (in the limit of an infinite number of PEs) to some manageable finite value.

2 The Basic Conservative Scheme

The basic notion of discrete-event simulations is that time is continuous and the discrete events occur instantaneously. Between events, the state (configuration) of the system remains unchanged. If the events occur at random instants of time, the dynamics can be referred to as asynchronous. In *conservative* PDES schemes [22], only those PEs that are guaranteed not to violate causality are

allowed to process their events and increment their local time. The rest of the PEs must “idle”. For simplicity we consider an arbitrary but one-dimensional underlying system (the “physical” system to be simulated) with nearest-neighbor interactions in which discrete events (update attempts in the local configuration) exhibit Poisson asynchrony. Further, we focus on the one site-per-PE scenario where each PE has its own local simulated time $h_i(t)$, constituting the virtual time horizon $\{h_i(t)\}_{i=1}^N$. Here t is the number of parallel steps executed by all PEs (proportional to the wall-clock time) and N is the number of PEs. By construction, $h_i(t)$ is the progress of PE i after parallel step t . In the following, we will use the terms “height”, “simulated time”, or “virtual time” interchangeably, since we refer to the same observable.

According to the basic conservative synchronization scheme, first introduced by Lubachevsky, [11,12], at each parallel step t , only those PEs for which the local simulated time is not greater than the local simulated times of their virtual 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 one in simulated time units [stu].) Thus, for the one-site-per-PE, one-dimensional regular virtual topology, if $h_i(t) \leq \min\{h_{i-1}(t), h_{i+1}(t)\}$, 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 [11,12]. (More general PDES schemes, where events to be processed by a PE are initiated (or generated) by the same PE (such as the basic conservative scheme above), are also referred to as self-initiating discrete-event schemes [23,24].) In the original algorithm, the virtual communication topology between PEs mimics the interaction topology of the underlying system [11,12,25]. When “simulating the simulations” based on the above simple “microscopic” rules for the evolution of the time horizon, we implemented periodic boundary conditions, i.e., the PEs are placed on a ring. In analyzing the performance of the above scheme, it is enormously helpful that the progress of the simulation itself is decoupled from the possibly complex behavior of the underlying system. This is contrary to optimistic approaches, where the evolution of the underlying system and the progress of the PDES simulation are strongly entangled [18], making scalability analysis a much more difficult task.

To understand the scalability and performance of the basic conservative scheme we study two basic observables: the average utilization $\langle u \rangle$ (the fraction of non-idling PEs), which directly corresponds to the average rate of progress of the simulation, and the average width of the virtual time horizon, which probes the complexity of data management during the simulation. On a *regular* one-dimensional lattice the utilization is the density of local minima

$$\langle u \rangle = \langle \Theta(h_{i-1} - h_i) \Theta(h_{i+1} - h_i) \rangle = \langle \Theta(-\phi_{i-1}) \Theta(\phi_i) \rangle, \quad (1)$$

where $\phi_i \equiv h_{i+1} - h_i$ is the local slope, $\Theta(\dots)$ is the Heaviside step function, and $\langle \dots \rangle$ denotes an ensemble average over the stochastic, exponentially distributed local simulated time increments. For a system of identical PEs (implying translational invariance), the above quantity is independent of i . The width,

characterizing the spread of the time horizon, is defined as

$$\langle w^2 \rangle = \left\langle \frac{1}{N} \sum_{i=1}^N (h_i - \bar{h})^2 \right\rangle \tag{2}$$

where $\bar{h} = (1/N) \sum_{i=1}^N h_i$ is the mean-height.

Here we use a coarse-grained description for the virtual time horizon to perform the scalability analysis [25,26]. It was shown [25] that the virtual time horizon exhibits Kardar-Parisi-Zhang (KPZ)-like [27] kinetic roughening [28] and the steady-state behavior in one dimension is governed by the Edwards-Wilkinson (EW) Hamiltonian [29]. The evolution of the simulated time horizon is effectively governed by the Langevin equation

$$\partial_t h_i(t) = \nabla^2 h_i - \lambda(\nabla h_i)^2 + \dots + \eta_i(t) , \tag{3}$$

where $\eta_i(t)$ is a delta correlated Gaussian noise $\langle \eta_i(t) \eta_j(t') \rangle = 2D \delta_{ij} \delta(t - t')$, and ∇ and ∇^2 are the discrete gradient and discrete Laplacian operators on a regular lattice, respectively. The \dots in (3) stands for infinitely many irrelevant terms in the long-time, large- N limit. Being primarily interested in the steady-state properties of the algorithm, we consider the equal-time height-height correlations, or alternatively, its Fourier transform, the corresponding structure factor $S^{(h)}(k, t)$, defined through

$$S^{(h)}(k, t) N \delta_{k, -k'} \equiv \langle \tilde{h}_k(t) \tilde{h}_{k'}(t) \rangle . \tag{4}$$

Here $\tilde{h}_k = \sum_{j=1}^N e^{-ikj} h_j$ is the Fourier transform of the virtual times with the wave number $k = (2\pi n)/N$, $n = 0, 1, 2, \dots, N-1$. In the long time limit in one dimension (EW stationary state), one has [26]

$$S^{(h)}(k) \equiv \lim_{t \rightarrow \infty} S^h(k, t) = \frac{D}{2[1 - \cos(k)]} . \tag{5}$$

The structure factor essentially contains all the “physics” needed to describe the scaling behavior of the time horizon. Figure 1(a) shows the measured structure factor, obtained by simulating the PDES simulation itself, based on the exact rules for the evolution of the local times. It confirms the $\sim 1/k^2$ coarse-grained prediction for small k values. Using the steady-state structure factor, one can express the width as

$$\langle w^2 \rangle = \frac{1}{N} \sum_{k \neq 0} S^{(h)}(k) . \tag{6}$$

The above summation can be carried out for the structure factor given by (5), yielding

$$\langle w^2 \rangle_N \simeq \frac{D}{12} N \sim N^{2\alpha} , \tag{7}$$




Fig. 1. (a) Steady-state height-height structure factor for various system sizes for the regular one-dimensional lattice, one-site-per-PE basic conservative PDES time horizon. The solid straight line indicates the theoretical $\sim 1/k^2$ behavior, (5), for small k values. (b) Steady-state width and extreme-height fluctuations for the same scheme as a function of the number of PEs. The dashed straight line corresponds to the exact KPZ (EW in one dimension) roughening (7).

(corresponding to a roughness exponent $\alpha=1/2$) in the limit of large N . Figure 1(b) shows the measured width, asymptotically approaching the above scaling form. For later calculations, we will also need the slope-slope steady-state structure factor

$$S^{(\phi)}(k) = 2[1 - \cos(k)]S^{(h)}(k) = D \quad (8)$$

and the corresponding correlation function

$$C^{(\phi)}(l) = \langle \phi_i \phi_{i+l} \rangle = \frac{1}{N} \sum_{k \neq 0} e^{ikl} S^{(\phi)}(k) \quad (9)$$

to study the density of local minima. From (8) and (9) it trivially follows that $C^{(\phi)}(l) = D\delta_{l,0}$ (i.e., the local slopes become independent) in the infinite system-size limit. Then the probability that two neighboring local slopes form a local minima is $1/4$. Hence, the density of local minima and the utilization $\langle u \rangle$ [see (1)] approaches $1/4$. (The steady state is governed by the EW Hamiltonian where the local slopes are independent.)

For more general two-point functions (but still within the coarse-grained Gaussian picture (5), we utilize a simple relationship between the density of local minima and the slope-slope correlation function [26]

$$\langle \Theta(-\phi_{i-1})\Theta(\phi_i) \rangle = \frac{1}{2\pi} \arccos \left(\frac{C^{(\phi)}(1)}{C^{(\phi)}(0)} \right). \quad (10)$$

The above formula can be used, e.g., to extract finite-size corrections to the utilization [26]. From (8) and (9), for a *finite* system, one finds that $C^{(\phi)}(l) =$

$D(\delta_{l,0} - 1/N)$ and from (1) and (10), $\langle u \rangle \simeq 1/4 + 1/(2\pi N)$. Clearly, the specific value $1/4$ in the thermodynamic limit and the prefactor of the $1/N$ finite-size corrections will differ from those of the actual PDES evolution with its specific “microscopic dynamics”. The density of local minima, however, must remain *nonzero* and it displays universal finite-size effects [25,26,30–34],

$$\langle u \rangle_N \simeq \langle u \rangle_\infty + \mathcal{O}\left(\frac{1}{N}\right), \quad \langle u \rangle_\infty \neq 0, \quad (11)$$

based on the universality class (EW in one dimension) the virtual time horizon belongs to. Thus, the average progress rate of the simulation approaches a nonzero constant in the asymptotic long-time, large- N limit. For example, for the one-site-per PE basic conservative PDES scheme $\langle u \rangle_\infty \simeq 0.2464$ [25,26], due to non-universal short-range correlations between the local slopes [35].

The average width of the virtual time horizon, however, diverges as $N \rightarrow \infty$ [see (7)], making the measurement phase of the PDES scheme (data collection) *not* scalable [30]. Since the effect of very large fluctuations in the progress of the individual PEs is also important (after all, delays will be caused by state-saving difficulties on the individual nodes, where extreme events occur), we investigated the properties of the extremal-height fluctuations. We considered the average of the largest height fluctuations above the mean $\Delta_{\max} \equiv h_{\max} - \bar{h}$. The average or typical extreme-height fluctuations in the basic conservative PDES scheme exhibit the same scaling behavior as the width itself, $\langle \Delta_{\max}^2 \rangle \sim N$ [Fig. 1(b)]. This is not particularly surprising in that the extreme fluctuations emerge through the dominating collective long-wavelength modes of the “critical” surface. This finding was also observed [36] for other surface growth models belonging to KPZ universality class.

Finally, we note that, in an attempt to construct an analytically tractable model for PDES, Greenberg et al. [14] 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 $t \rightarrow \infty$, $N \rightarrow \infty$ 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 have shown above, their latter conjecture for the width does not hold, thus, the basic conservative PDES scheme for *regular lattices* cannot be equivalently described by K -random model (at least not below the critical dimension of the KPZ universality class [28,37]). Nevertheless, their “annealed” random connection model is highly inspiring in that the underlying connection topology can have crucial effects on the universal behavior of the evolution of the virtual time horizon, and in turn, on the synchronizability of PDES schemes.

3 The Small-World Synchronized Conservative PDES Scheme

3.1 Motivation and Properties for the Synchronization Network

The divergent width and extreme-height fluctuations (with increasing N), discussed in the previous section, are the result of the divergent lateral correlation length $\xi^{(h)}$ of the virtual time surface, which reaches the system size N in the steady state [28,30]. To de-correlate the simulated time horizon, first, we modify the virtual communication topology of the PEs. The resulting communication network *must include* the original short-range (nearest-neighbor) connections to faithfully simulate the dynamics of the underlying system. In the modified network, the connectivity of the nodes (the number of neighbors) should remain *non-extensive* (i.e., only a finite number of virtual neighbors per node is allowed). This is in accordance with our desire to design a PDES scheme where no global “intervention” or synchronization is employed (PEs can only have $\mathcal{O}(1)$ communication exchanges per step). It is clear that the added synchronization links (or at least some of those) have to be long range. (Only short range links would not change the universality class and the scaling properties of the width of the time horizon). Also, fluctuations in the individual connectivity should be avoided for load balancing purposes, i.e., requiring the same number of added links (e.g., one) for each node is a reasonable constraint.

One may wonder how the collective behavior of the PDES scheme would change if each node was connected to the one located at the “maximum” possible distance away from it ($N/2$ on a ring) [Fig. 2(a)] [38]. Consider a linear coarse-grained Langevin equation with Gaussian noise where the effective strength of the added long-range links is Σ ,

$$\partial_t h_i(t) = (h_{i+1} + h_{i-1} - 2h_i) - \Sigma(h_i - h_{i+N/2}) + \eta_i(t) , \tag{12}$$

with periodic boundary conditions. After elementary calculations one obtains for the width

$$\langle w^2 \rangle = \frac{1}{N} \sum_{k \neq 0} S^{(h)}(k) = \frac{1}{N} \sum_{k \neq 0} \frac{D}{2[1 - \cos(k)] + 2\Sigma[1 - \cos(kN/2)]} , \tag{13}$$

where $k = (2\pi n)/N$, $n = 0, 1, 2, \dots, N-1$ as before (and N is even for simplicity). Separating the terms with even and odd n values above, we find

$$\begin{aligned} \langle w^2 \rangle &= \frac{1}{N} \sum_{n=\text{odd}} \frac{D}{2[1 - \cos(2\pi n/N)] + 4\Sigma} \\ &+ \frac{1}{N} \sum_{n=\text{even}} \frac{D}{2[1 - \cos(2\pi n/N)]} . \end{aligned} \tag{14}$$

The first sum yields a finite N independent value in the $N \rightarrow \infty$ limit. The second sum, on the other hand, is identical to the width of the EW model on




Fig. 2. Schematic diagrams for the PDES synchronization networks. (a) Maximal-distance connected network as described in the text. (b) Small-world network where each PE has exactly one quenched random neighbor.

a regular network of size $N/2$. Thus, in the large N limit the width for the “maximal-distance” connected network [Fig. 2(a)] diverges as $\langle w^2 \rangle_N \simeq DN/24$. Indeed, one can realize, that such regularly patterned long-range links make the network equivalent to a $2 \times (N/2)$ quasi one-dimensional system with only nearest-neighbor interactions and helical boundary conditions. The above extreme case suggests, that the purely maximum-range synchronization cannot work either.

We then choose the extra synchronization links in such a way that they cover all lengthscales with equal weight [21]. With the one extra link per PE constraint, we employ *quenched* random bidirectional links, i.e., each PE is connected to exactly one other PE, as illustrated on Fig. 2(b). That is, pairs of sites selected at random, and once they are linked they cannot be selected again. The resulting network resembles a (constrained) small-world-like network [20]. It differs from both the original (“rewiring”) [20,39] and the “soft” version [40,41] of the Small-World (SW) network (where an Erdős-Rényi random graph is thrown on top of a regular lattice). Our construction too, however, exhibits a well balanced coexistence among short- and long-range links (random links are placed on the top of a regular substrate), and we will refer to it as a SW network in what follows. When explicit distinction is needed among the above versions of the SW networks, we will refer to our construction as the “hard” version of the SW network. This terminology is motivated by the eigenvalue spectrum of the Laplacian on the different variations of the SW networks [42,43], discussed in more detail in [43,44] and in the chapter by Hastings and Kozma in this book.

As one can expect, the average path length $\langle l \rangle_N$ (the average minimum number of links connecting two randomly chosen nodes) for our synchronization network scales *logarithmically* with the system size N [Fig 3], i.e., like most other random networks [45], it too exhibits the “small-world” character (or low-degree of separation).

We now describe the modified algorithmic steps for the SW connected PEs [21]. In the modified conservative PDES scheme, at every parallel step each PE with probability p compares its local simulated time with its *full* virtual




Fig. 3. Average shortest path as a function of the logarithm of the number of nodes (PEs) for our small-world synchronization network [Fig. 2(b)]. The straight line represents the slope of the asymptotic large N behavior of the average shortest path $\langle l \rangle_N \simeq 1.42 \ln(N)$.

neighborhood and can only advance if it is a neighborhood minimum, i.e., if $h_i(t) \leq \min\{h_{i-1}(t), h_{i+1}(t), h_{r(i)}(t)\}$, where $r(i)$ is the random connection of PE i . With probability $(1 - p)$ each PE follows the original scheme, i.e., the PE then can advance if $h_i(t) \leq \min\{h_{i-1}(t), h_{i+1}(t)\}$. 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.

3.2 Coarse-Grained Equation of Motion for the Small-World-Coupled Conservative PDES Scheme

We now obtain a coarse-grained description for the evolution of the virtual time horizon. The occasional checking of the virtual time (at every $1/p$ parallel steps on average) through the random links introduces an effective strength \tilde{p} for these links. Note that this is a dynamic “averaging” process, controlled by the parameter p , the probability of checking the random neighbor as well. The only properties we assume about $\tilde{p}(p)$ is that it is a monotonically increasing function of p and is only zero when $p=0$. The effective Langevin equation then becomes

$$\partial_t h_i(t) = (h_{i+1} + h_{i-1} - 2h_i) - \sum_{j=1}^N J_{ij}(h_i - h_j) + \dots + \eta_i(t), \quad (15)$$

where $\eta_i(t)$ is delta-correlated Gaussian noise as in (3) and J_{ij} is proportional to the symmetric adjacency matrix of the random part of the network with exactly one non-zero element (being equal to \tilde{p}) in each row and column. The former property implies that $\sum_l J_{il} = \tilde{p}$ for all i , which is related to our construction that there are no fluctuations in the individual connectivity. The \dots in (15) stand for *all* non-linear terms (involving non-linear interactions through the random links

as well). “Phenomenological” results of simulating the simulation (Sect. 3.3) suggest that the dynamic control of the link strength and non-linearities only give rise to a renormalized coupling and a corresponding renormalized mass (in a field theory sense). Thus, the dynamics is effectively governed by EW relaxation in a small world. This motivates the study of the EW model on a SW network, i.e., keeping only the linear terms in (15). That problem is studied in detail in [43,44] and in the chapter by Hastings and Kozma in this book. A disorder-averaged systematic perturbation expansion yields an effective “mass” $\Sigma(\tilde{p}) \sim \tilde{p} + \mathcal{O}(\tilde{p}^{3/2})$ in the asymptotic small- \tilde{p} limit. In our case, when non-linearities are indeed present and the strength of the random links is controlled by the relative frequency p of the synchronization steps through those links, we will only assume that $\Sigma(p)$ is a monotonically increasing function of p and is only zero when $p=0$. In the following, for brevity, $\langle \dots \rangle$ will denote the double average: ensemble average based on the stochastic dynamics [e.g., over the noise in (15)], and disorder average over the random network realizations. The resulting steady-state structure factor (or propagator) for (15) then reads as [44]

$$S^{(h)}(k) = \frac{1}{N} \langle \tilde{h}_k \tilde{h}_{-k} \rangle = \frac{D}{2[1 - \cos(k)] + \Sigma} . \quad (16)$$

The above structure factor contains the essential properties of the SW synchronized PDES scheme at the coarse grained level. In particular, the SW links induce a finite correlation length $\xi^{(h)}$ for the surface fluctuations. In the following we will only discuss the infinite-system small- Σ behavior, when the finite-size effects vanish and the discrete-lattice effects become negligible. In this limit, $\xi^{(h)} \simeq 1/\sqrt{\Sigma}$. Also from (16), for the width of the time horizon one obtains

$$\langle w^2 \rangle = \frac{1}{N} \sum_{k \neq 0} S^{(h)}(k) \simeq \frac{1}{2\sqrt{\Sigma}} . \quad (17)$$

i.e., the width remains finite in the $N \rightarrow \infty$ limit. (Note that $\Sigma(p)$ is only zero when $p=0$.) The implication of this result for the SW synchronized PDES scheme is that the spread of the virtual time horizon will approach a finite value in the limit of infinite number of PEs for any nonzero value of p .

We now discuss some general considerations for the the utilization $\langle u \rangle$ (the average progress rate) for the SW synchronized PDES scheme. From the algorithmic rules it follows that

$$\langle u \rangle = (1 - p) \langle \Theta(-\phi_{i-1}) \Theta(\phi_i) \rangle + p \langle \Theta(-\phi_{i-1}) \Theta(\phi_i) \Theta(h_{r(i)} - h_i) \rangle , \quad (18)$$

where p is the probability to include the random neighbor as well in the synchronization step. Note that the disorder averaging makes the right hand side independent of i . For general p (with the random links present) it is hard to carry out quantitative approximations for the utilization. Since the height fluctuations *become* short-range correlated (16) and the local slopes *remain* short-range correlated [see discussion below, (21)], it is guaranteed that both terms in (18), and subsequently $\langle u \rangle$, *remain non-zero* for any $0 \leq p \leq 1$ [35,46]. Rearranging the terms in (18) one obtains

$$\langle u \rangle = \langle \Theta(-\phi_{i-1})\Theta(\phi_i) \rangle - p \left[\langle \Theta(-\phi_{i-1})\Theta(\phi_i) \rangle - \langle \Theta(-\phi_{i-1})\Theta(\phi_i)\Theta(h_{r(i)} - h_i) \rangle \right]. \quad (19)$$

The first term, $\langle \Theta(-\phi_{i-1})\Theta(\phi_i) \rangle = \langle \Theta(h_{i-1} - h_i)\Theta(h_{i+1} - h_i) \rangle$, is an increasing function of p , as the heights become less correlated [46]. For example, it would be 1/4 for completely independent slopes, and it would be 1/3 for completely independent heights. The actual values differ for the PDES time horizon (the slopes and heights exhibit some short-range correlations), but the above trend remains and $\langle \Theta(-\phi_{i-1})\Theta(\phi_i) \rangle$ saturates rapidly as a function of p [46]. The quantity in [...] in (19) is always positive, bounded from zero, so it will eventually lead to the *decrease* in $\langle u \rangle$ as $\mathcal{O}(p)$, *once* $\langle \Theta(-\phi_{i-1})\Theta(\phi_i) \rangle$ saturates. For very small values of p , however, the leading order correction to $\langle \Theta(-\phi_{i-1})\Theta(\phi_i) \rangle$ may become more dominant than $\mathcal{O}(p)$. In this case, as it is clear from (19), the small- p behavior of $\langle \Theta(-\phi_{i-1})\Theta(\phi_i) \rangle$ alone yields the asymptotic small- p behavior of $\langle u \rangle$.

We now continue to discuss the density of local minima and the utilization for the coarse-grained linear model with Gaussian noise [(15) and (16)], which may capture some of the small- p features of the actual PDES time horizon. We also make the mean-field assumption that the structure factor and correlation functions are self-averaging in the large system-size limit. First, from (16) we find for the slope-slope structure factor

$$S^{(\phi)}(k) = 2[1 - \cos(k)]S^{(h)}(k) = D \left\{ 1 - \frac{\Sigma}{2[1 - \cos(k)] + \Sigma} \right\}, \quad (20)$$

which yields

$$C^{(\phi)}(l) \simeq D \left\{ \delta_{l,0} - \Sigma \frac{e^{-l\sqrt{\Sigma}}}{2\sqrt{\Sigma}} \right\} = D \left\{ \delta_{l,0} - \frac{\sqrt{\Sigma}}{2} e^{-l\sqrt{\Sigma}} \right\} \quad (21)$$

for the slope correlation function in the infinite system-size, small- Σ limit. The above equation shows explicitly, that the local slopes remain short-range correlated for the SW-synchronized time horizon. Using (10) and the above form of the slope correlations, in the small- Σ limit we obtain

$$\langle \Theta(-\phi_{i-1})\Theta(\phi_i) \rangle \simeq \frac{1}{4} + \frac{\sqrt{\Sigma}}{4\pi} - \frac{\Sigma}{8\pi} + \dots \quad (22)$$

This implies that increasing the effective mass increases the density of local minima. This is not surprising, in that increasing Σ reduces the correlation length $\xi^{(h)}$ for the height fluctuations, as discussed above. Using (19) and (22), we obtain for the utilization

$$\langle u \rangle \simeq \frac{1}{4} + \frac{\sqrt{\Sigma(p)}}{4\pi} - \frac{\Sigma(p)}{8\pi} + \dots + \mathcal{O}(p), \quad (23)$$

where we now explicitly indicated the p -dependence of Σ . If $\Sigma(p)$ is known, more precisely, if $\Sigma(p) \sim p^s$ with $s < 2$ for small p values, the above equation

becomes useful to extract the asymptotic small- p behavior of the utilization. It is instructive to consider the mean-field case when the effective strength of the random links [in (15)] scales as p , at least for small p values. Then $\Sigma(p) \sim p$ [44], and to leading order in p , one finds

$$\langle u \rangle \simeq \frac{1}{4} + \frac{\sqrt{p}}{4\pi} + \mathcal{O}(p). \quad (24)$$

The above counterintuitive behavior of increasing $\langle u \rangle$ by actually synchronizing “more” is the result of the gain in $\langle \Theta(-\phi_{i-1})\Theta(\phi_i) \rangle [\mathcal{O}(\sqrt{p})]$ winning over the loss due to the occasional extra random synchronizations $[\mathcal{O}(p)]$, for *asymptotically small* p values. As p is increased, (24) will not be valid anymore; $\langle \Theta(-\phi_{i-1})\Theta(\phi_i) \rangle$ starts to saturate, so the change in $\langle u \rangle$ will be dominated by the $-\mathcal{O}(p)$ factor in (19).

Some analogy between the evolution of the virtual time horizon with quenched random links added and the sliding state of charge-density waves with “no-passing” rule [47] suggests [48] that the above mean-field coarse-graining argument may break down and the average rate of progress of the SW-synchronized conservative scheme for arbitrary small p is bounded by that of the $p=0$ case.

3.3 Comparison with the Simulated Small-World Synchronized PDES Results

We now turn to discussing the results obtained by simulating the actual PDES scheme, based on the specific update rules for the local simulated times (see the end of Sect. 3.1). The fundamental difference between the original and the SW-synchronized conservative scheme is illustrated on Fig. 4. The snapshots of the virtual times indicate that, indeed, the large-amplitude long-wavelength fluctuations, present in the original time horizon [Fig. 4(a)] are suppressed when the extra synchronizations through the quenched random links are implemented [Fig. 4(b)].

Since all steady-state characteristics are “encoded” in the structure factor of the virtual times, we measured this quantity, and compared it to

$$S^{(h)}(k) \propto \frac{1}{k^2 + \Sigma(p)}, \quad (25)$$

the small k limit of (16). Figure 5(a) shows the disorder-averaged structure factors, as well as individual realizations for various system sizes. As one can observe, finite-size effects become small and $S^{(h)}(k)$ approaches a *finite* value as $k \rightarrow 0$. Thus, there are no large-amplitude, long-wavelength modes in the virtual time horizon. Further, the inset of Fig. 5(a) confirms the “massive” behavior (25) for small k values. It is important to note that for the actual PDES time horizon, the effective mass $\Sigma(p)$ may depend non-trivially on p as a result of the dynamic control of the “link strength” and renormalization by nonlinear effects in the specific “microscopic dynamics”. The form of $S(k)$, however, seems to follow the linear theory, discussed in Sect. 3.2. By plotting $1/S^{(h)}(k)$ vs. k^2




Fig. 4. Snapshots of the virtual time horizon in the steady state (a) for the original regular lattice ($p=0$) and (b) for the small-world synchronized ($p=0.10$) PDEs scheme for $N=10000$. The vertical scale is the same in (a) and (b).

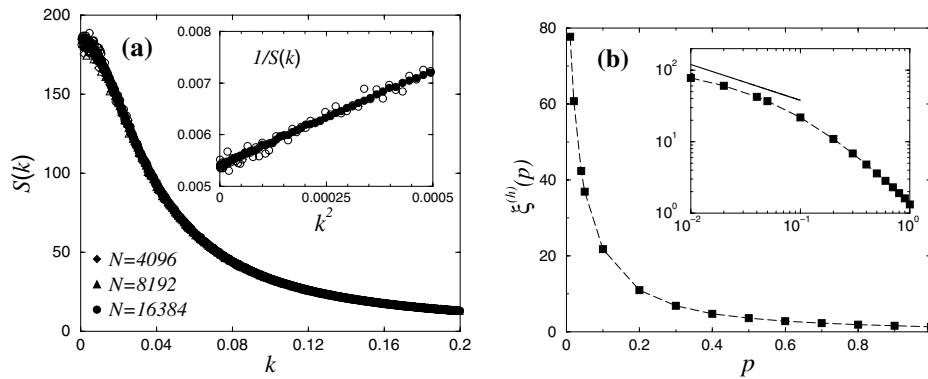


Fig. 5. (a) Steady-state structure factor of the virtual time horizon for $p=0.10$. In addition to ensemble averages over 100 realizations of the random links (filled symbols), single realizations (the same open symbols) are also shown. The inset shows a magnified view of $1/S(k)$ versus k^2 for small k for the largest system. The solid straight line is the best linear fit used to determine the correlation length. (b) Correlation length of the virtual time horizon as a function of the random synchronization frequency p . For reference, the solid straight line corresponds to the power-law divergence with the power $-1/2$.

[inset of Fig. 5(a)] for the measured structure factors, we can “phenomenologically” determine the correlation length of the fluctuations of the virtual times, $\xi^{(h)}(p) \simeq 1/\sqrt{\Sigma(p)}$. The results are shown in Fig. 5(b)]. The inset implies, that the asymptotic small- p scaling regime has not been reached yet at our smallest value $p=0.01$.

The results for the width and for the utilization also agree with the basic prediction of the coarse-grained linear model. Even for small p values, the width saturates, and approaches a system-size independent value for large N

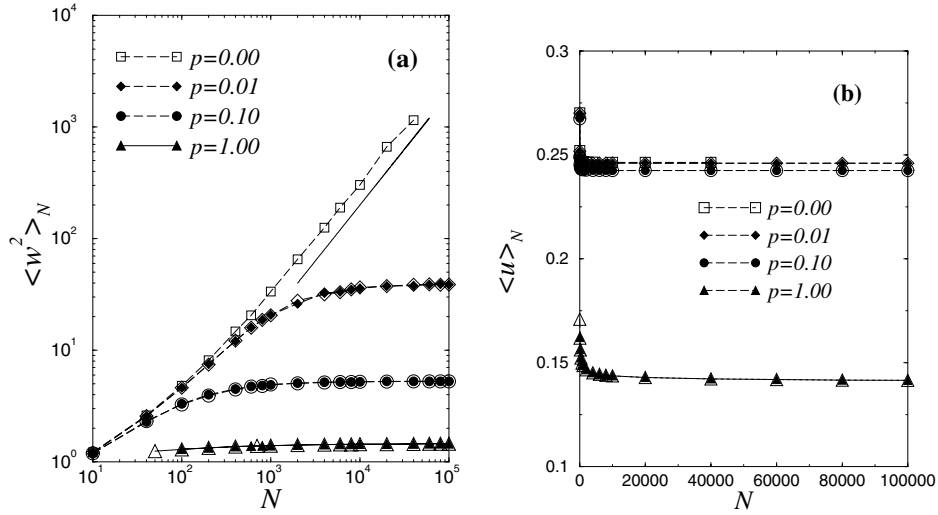


Fig. 6. (a) Average steady-state width of the virtual time horizon as function of the number of PEs for various values of p . In addition to ensemble averages over 10 realizations of the random links (filled symbols), a single realization is also shown (the same open symbols). The solid straight line represents the asymptotic one-dimensional KPZ power-law divergence with roughness exponent $\alpha=1/2$ for the $p=0$ case. Note the log-log scales. (b) The steady-state utilization (fraction of non-idling PEs) for the same cases as in (a).

[Fig. 6(a)] while the utilization remains non-zero [Fig. 6(b)]. For example, for a hypothetically infinite system, for $p=0.01$, $\langle w^2 \rangle$ is reduced from “infinity” (the width for the KPZ surface) to about 40, while the utilization drops from 0.2464 only to about 0.2460. For $p=0.10$, the width is further reduced to about 5, while the utilization is down only to 0.242. One can also observe the clear self-averaging property for both global observables (the width and the utilization), i.e., their values become independent of the realization of the SW network for large enough N .

To extract the asymptotic small- p behavior of the width and the utilization, one would need larger system sizes, longer steady-state PDES time series, and more network realizations to obtain reliable statistics, and to compare all aspects of the linearized coarse-grained model with the actual PDES simulations. In particular, it would be interesting to see, whether the utilization increases initially for sufficiently small p values (the subtle prediction of the coarse-grained linear theory). We have not observed this, but the systems we simulated have not yet reached their asymptotic scaling regime [inset of Fig. 5(b)]. Further, finite-size corrections and error bars may become comparable to this possible asymptotically small effect in $\langle u \rangle$.

3.4 Extremal Fluctuations of the Virtual Time Horizon

In addition to the average value of the fluctuations of the local field variables (such as the height in the context of surface growth models), the typical value of the largest fluctuations can also be of great importance [49–51] in a number of applications. For example, in load balancing networks [52] or state-saving schemes for PDES schemes [30,53], extreme (load or accumulated data) fluctuations on an *individual* node will cause the delays. Thus, in interacting multi-component systems such as the above examples, failures or delays are triggered by extreme-events occurring on the individual components [51].

Relationship between extremum statistics and universal fluctuations in *correlated* systems have been discussed intensively in recent years. [36,54–60]. For the original PDES scheme ($p=0$, regular lattice synchronization) exhibiting a KPZ-like rough (or critical) surface, we illustrated (Sect. 2) that the extremal fluctuations of the time horizon diverge in the same fashion as the width itself [Fig. 1(b)]. We now discuss to what extent SW synchronizations lead to the suppression of the extreme-height fluctuations in the virtual time horizon [53], closely related to the measurement scalability of the conservative PDES scheme.

First, consider N *independent* identically distributed stochastic variables with a complemter cumulative distribution $P_{>}(x)$ (the probability that the individual stochastic variable is greater than x). Then the cumulative distribution $P_{<}^{\max}(x)$ for the largest of the N events (the probability that the extremal value is less than x) can be approximated as [60,61]

$$P_{<}^{\max}(x) = [P_{<}(x)]^N = [1 - P_{>}(x)]^N = e^{N \ln[1 - P_{>}(x)]} \simeq e^{-NP_{>}(x)}, \quad (26)$$

where one typically assumes that the dominant contribution to the statistics of the extremes comes from the tail of the individual distribution $P_{>}(x)$. For example, for exponentially-tailed individual variables, $P_{>}(x) \simeq e^{-cx}$, the above equation yields

$$P_{<}^{\max}(x) \simeq e^{-e^{-cx + \ln(N)}}. \quad (27)$$

Thus, the sequence of scaled variables $\tilde{x} = c(x - \ln(N)/c)$ asymptotically approaches the standard Fisher-Tippett-Gumbel (FTG) distribution [49,50]

$$\tilde{P}_{<}^{\max}(\tilde{x}) \simeq e^{-e^{-\tilde{x}}} \quad (28)$$

with mean $\langle \tilde{x} \rangle = \gamma$ ($\gamma = 0.577 \dots$ being the Euler constant) and variance $\langle \tilde{x}^2 \rangle - \langle \tilde{x} \rangle^2 = \pi^2/6$. It immediately follows that the average value of the largest of the N original random variables then scales as

$$\langle x^{\max} \rangle = \gamma/c + \ln(N)/c \simeq \ln(N)/c \quad (29)$$

for large N values. When comparing with simulation or experimental data, it is often convenient to use the scaled variables $\hat{x} = (x - \langle x \rangle)/\sigma_x$ which for the above case yields the FTG limit distribution with zero mean and unit variance




Fig. 7. Disorder-averaged probability density (histogram) for the individual simulated time fluctuations for various system sizes at $p=0.10$ (log-normal scales). The solid straight line indicates a pure exponential tail.

$$\hat{P}_{<}^{\max}(\hat{x}) \simeq e^{-e^{-(a\hat{x}+\gamma)}}, \quad (30)$$

where $a=\pi/\sqrt{6}$. Note that with *appropriately* chosen scaled variables, the convergence to the FTG distribution holds not only for exponential variables, but also for more general ones with “exponential-like” tails $P_{>}(x) \simeq e^{-cx^\delta}$ (i.e., decaying faster than any power law) [49,50,60,61]. For any $\delta \neq 1$, however, the convergence to (28) or (30) is *extremely* (logarithmically) slow [61].

For the SW synchronized PDES scheme with N PEs we showed that a finite correlation length $\xi^{(h)}(p) \simeq 1/\sqrt{\Sigma(p)}$ effectively decouples the local simulated times. Then, the extreme-value limit theorems can be applied [60,61] using the number of independent blocks $N/\xi^{(h)}$ in the system. Further, we found [62] that the tail of disorder-averaged distribution of the individual relative height fluctuations (independent of the site i) are simple exponentials $P_{>}(h_i - \bar{h}) \simeq \exp[-c(h_i - \bar{h})/w]$ with $w \equiv \sqrt{\langle w^2 \rangle}$. The histogram for the corresponding probability *density* function, $p(h_i - \bar{h})$, is shown in Fig. 7.

From the general extreme-value limit theorems, discussed above, it follows that the scaled extreme-height fluctuations are governed by the FTG distribution (30) (if scaled to zero mean and unit variance). Further, from (29), the average maximum relative height, $\Delta_{\max} = h_{\max} - \bar{h}$, will scale as

$$\langle \Delta_{\max} \rangle \simeq \frac{w}{c} \ln(N/\xi^{(h)}) \simeq \frac{w}{c} \ln(N), \quad (31)$$

where we dropped all N -independent terms. Note that both w and $\xi^{(h)}$ approach their finite N -independent values for any non-zero p , and the only N dependent factor is $\ln(N)$ for large N values.

Agreement between the simulated PDES extremal fluctuations and the above considerations are rather convincing. In Fig. 8(a) we show the scaled histograms (to zero mean and unit variance) for the extreme-height fluctuations together with the probability *density*, corresponding to (30)




Fig. 8. Extreme fluctuations in the the small-world synchronized PDES time horizon for $p=0.10$. (a) Probability density of the scaled extremal fluctuations and comparison with the FTG density equation (32). (b) Average of the extremal fluctuations (also shown is the width for comparison). The solid straight line indicates the logarithmic divergence.

$$\hat{p}(\hat{x}) \simeq a e^{-(a\hat{x}+\gamma)} - e^{-(a\hat{x}+\gamma)}, \quad (32)$$

We note again, that the underlying reason for the fast convergence to the FTG density of the simulated time horizon is that the local relative height distributions exhibit *pure* exponential tails. Also, for the more general distribution $P_{>}(h_i - \bar{h}) \simeq \exp[-c((h_i - \bar{h})/w)^\delta]$, the approach to the FTG limit distribution would be very slow and the corresponding maximum fluctuations would scale as $\sim [\ln(N)]^{1/\delta}$ [53,61,62], as opposed to (31).

In Fig. 8(b) we show the average of the largest fluctuations above the mean, Δ_{\max} , for the simulated PDES time horizon. The figure confirms that for large enough N (when $\langle w^2 \rangle$ essentially becomes system-size independent) Δ_{\max} increases *logarithmically* with the system size, according to (31) [Fig. 8(b)]. Simulation results for the actual PDES scheme also indicate [62] that the largest deviations below the mean, $\Delta_{\min} = \bar{h} - h_{\min}$, and the maximum separation, $\Delta = h_{\max} - h_{\min}$, scale the same way as Δ_{\max} , i.e., diverge logarithmically with the system size. Note, that similar to the width and the utilization, the extremal height fluctuations are also self-averaging [53,62].

The implication of these findings is that while the width becomes finite for SW-synchronized virtual times, any node can exhibit fluctuations of size $\mathcal{O}(\ln(N))$ in its local simulated time (related to the local memory need). We refer to this property as “marginally” scalable for the measurement phase (due to the weak logarithmic divergence). This property still ensures synchronization in a practical sense for the SW-synchronized PDES scheme with millions of PEs. Note, that this logarithmic system-size dependence of the extreme fluctuations is generic to coupled multi-component system, where some local relaxational dynamics is extended to a SW network [53,62].

4 Summary

Based on a mapping [25] between the evolution of the virtual time horizon for the basic conservative PDES scheme [11,12] and kinetically grown non-equilibrium surfaces [28], we constructed a coarse-grained description for the scalability and performance of such large-scale parallel simulation schemes. These schemes can be applied to large spatially extended systems with short-range interactions and asynchronous dynamics. The one-site-per PE basic PDES was shown to exhibit KPZ-like kinetic roughening. This scheme is scalable in that the average progress rate of the PEs approaches a non-zero value. The spread of the virtual time horizon, however, diverges as the square root of the number of PEs, leading to “de-synchronization” and difficulties in data management.

Universality arguments, and actual PDES simulations suggest [31], that the above characteristics generically hold for any underlying system with short-range interactions for any finite number of sites per PE implementations. Possible idling due to the conservative synchronization rules and actual communication times can be greatly suppressed by each PE carrying a large block of sites [11, 12], yielding encouraging values for the utilization for actual implementations [8]. When the PEs carry many sites, however, the saturation value of the width can become extremely large. More precisely, there is an additional fast-roughening phase for early times when the evolution of the time horizon corresponds to random deposition [31]. Subsequently, it will cross over to the KPZ growth regime and finally saturate. This further motivates the need for some sort of extra synchronizations to suppress the roughness of the time horizon.

Our goal here was to achieve synchronization without any global intervention. We constructed a specific version of the SW network, where each PE was connected to exactly one other randomly chosen PE. The extra synchronization steps through the random links are merely used to control the width. The virtual time horizon for the SW-synchronized PDES scheme becomes “macroscopically” smooth and essentially exhibits mean-field like characteristics. The random links, on top of a regular lattice, generate an effective “mass” for the propagator of the virtual time horizon, corresponding to a nonzero correlation length. The width becomes finite, for an arbitrary small rate of synchronization through the random links, while the utilization remains nonzero, yielding a fully scalable PDES scheme. The former statement is only marginally weakened by observing that the extreme fluctuations in the time horizon can exhibit logarithmically large values as a function of the total number of PEs. The above predictions of the coarse-grained PDES model were confirmed by actually “simulating the simulations”.

The generalization when random links are added to a higher-dimensional underlying regular lattice is clear: since our construction of the SW network (“hard” version) on a one-dimensional regular substrate is already mean-field, in higher dimensions it will be even more so [44] (i.e., the critical dimension of our SW network is less than one). Note that synchronizability on scale-free networks [45,63] was also studied recently. The results indicate that a PDES scheme is marginally scalable if the communication topology between the PEs is a scale-

free network [35]. The implication of this finding is that the internet, which is already exploited for distributed computing for mostly “embarrassingly parallel” problems through existing GRID-based schemes [16,17], may have the potential to accommodate efficient complex system simulations (such as asynchronous PDES) where the nodes frequently have to synchronize with each other.

The above construction of a fully scalable algorithm for simulating large systems with asynchronous dynamics and short-range interactions is an example for the enormous “computational power and synchronizability” [20] that can be achieved by SW couplings. The suppression of critical fluctuations of the virtual time horizon is also closely related to the emergence of mean-field-like phase transitions and phase ordering in *non-frustrated* interacting systems [1,64–70].

Recent theoretical work also supports [43,44,71] that systems without inherent frustration exhibit strict or anomalous mean-field characteristics when the original short-range interaction topology is modified to a SW network.

Acknowledgements

Discussions with G. Györgyi, M.B. Hastings, G. Istrate, B. Kozma, B.D. Lubachevsky, Z. Rácz, and P.A. Rikvold, are gratefully acknowledged. G.K. and M.A.N thank CNLS, Los Alamos National Laboratory for their hospitality during their stays in summer 2002 and 2003 where part of this work was completed. We acknowledge the support of US NSF through Grant No. DMR-0113049 and the support of the Research Corporation Grant No. RI0761. Z.T. is supported by the US DOE under contract W-7405-ENG-36. H.G. was also supported in part by the Los Alamos summer student program in 2002 and 2003, DOE W-7405-ENG-36.

References

1. S.H. Strogatz, *Nature* **410**, 268 (2001).
2. R. Fujimoto, *Commun. ACM* **33**, 30 (1990).
3. D.M. Nicol, R.M. Fujimoto, *Ann. Oper. Res.* **53**, 249 (1994).
4. B.D. Lubachevsky, *Bell Labs Tech. J.* **5** April-June 2000, 134 (2000).
5. A.G. Greenberg et al. in *Proc. 8th Workshop on Parallel and Distributed Simulation (PADS'94), Edinburgh, UK, 1994* (SCS, San Diego, CA, 1994), p. 187.
6. E. Deelman, B.K. Szymanski, T. Caraco, in *Proc. 28th Winter Simulation Conference*, (ACM, New York, 1996), p. 1191.
7. D.M. Nicol, *Proc. 1988 SCS Multiconference*, Simulation Series, SCS, Vol. 19, p. 141 (1988).
8. G. Korniss, M.A. Novotny, P.A. Rikvold, *J. Comput. Phys.* **153**, 488 (1999).
9. G. Korniss, C.J. White, P.A. Rikvold, M.A. Novotny, *Phys. Rev. E* **63**, 016120 (2001).
10. G. Korniss, P.A. Rikvold, M.A. Novotny, *Phys. Rev. E* **66**, 056127 (2002).
11. B.D. Lubachevsky, *Complex Syst.* **1**, 1099 (1987).
12. B.D. Lubachevsky, *J. Comput. Phys.* **75**, 103 (1988).
13. D.R. Jefferson, *ACM Trans. Prog. Lang. and Syst.* **7**, 404 (1985).

14. A.G. Greenberg, S. Shenker, and A.L. Stolyar, *Performance Eval. Rev.* **24**, 91 (1996).
15. *Blue Gene/L project*, partnership between IBM and DoE, announced Nov.9, 2001; 64K processors, expected scale 200 teraflops, a step towards a petaflop scale; expected completion 2005; IBM Research Report, RC22570 (W0209-033) September 10, 2002.
16. See, e.g., www.gridforum.org and setiathome.ssl.berkeley.edu.
17. S. Kirkpatrick, *Science* **299**, 668 (2003).
18. P.M.A. Sloot, B.J. Overeinder, A. Schoneveld, *Comput. Phys. Commun.* **142**, 76 (2001).
19. P. Bak, C. Tang, K. Wiesenfeld, *Phys. Rev. Lett.* **59**, 381 (1987).
20. D.J. Watts and S.H. Strogatz, *Nature* **393**, 440 (1998).
21. G. Korniss, M.A. Novotny, H. Guclu, Z. Toroczkai, and P.A. Rikvold, *Science* **299**, 677 (2003).
22. K.M. Chandy, J. Misra, *Commun. ACM* **24**, 198 (1981).
23. R.E. Felderman and L. Kleinrock, *ACM Trans. Model. Comput. Simul.* **1**, 386 (1991).
24. D.M. Nicol, *ACM Trans. Model. Comput. Simul.* **1**, 24 (1991).
25. G. Korniss, Z. Toroczkai, M.A. Novotny, and P.A. Rikvold, *Phys. Rev. Lett.* **84**, 1351 (2000).
26. Z. Toroczkai, G. Korniss, S. Das Sarma, and R.K.P. Zia, *Phys. Rev. E* **62**, 276 (2000).
27. M. Kardar, G. Parisi, Y.-C. Zhang, *Phys. Rev. Lett.* **56**, 889 (1986).
28. A.-L. Barabási and H.E. Stanley, *Fractal Concepts in Surface Growth* (Cambridge University Press, Cambridge, 1995).
29. S.F. Edwards and D.R. Wilkinson, *Proc. R. Soc. London, Ser A* **381**, 17 (1982).
30. G. Korniss, M.A. Novotny, A.K. Kolakowska, and H. Guclu, SAC 2002, *Proceedings of the 2002 ACM Symposium on Applied Computing*, pp. 132-138, (2002).
31. A. Kolakowska, M. A. Novotny, and G. Korniss, *Phys. Rev. E* **67**, 046703 (2003).
32. A. Kolakowska, M. A. Novotny, and P.A. Rikvold, *Phys. Rev. E* **68**, 046705 (2003).
33. A. Kolakowska and M. A. Novotny, *Phys. Rev. B* **69**, 075407 (2004).
34. J. Krug and P. Meakin, *J. Phys. A* **23**, L987 (1990).
35. Z. Toroczkai, G. Korniss, M. A. Novotny, and H. Guclu, in *Computational Complexity and Statistical Physics*, edited by A. Percus, G. Istrate, and C. Moore, Santa Fe Institute Studies in the Sciences of Complexity Series (Oxford University Press, 2004, in press); arXiv:cond-mat/0304617 (2003).
36. S. Raychaudhuri, M. Cranston, C. Przybyla, and Y. Shapir, *Phys. Rev. Lett.* **87**, 136101 (2001).
37. E. Marinari, A. Pagnani, G. Parisi, and Z. Rácz, *Phys. Rev. E* **65**, 026136 (2002).
38. Z. Rácz, private communications.
39. D.J. Watts, *Small Worlds* (Princeton Univ. Press, Princeton, 1999).
40. M.E.J. Newman, *J. Stat. Phys.* **101**, 819 (2000).
41. M.E.J. Newman and D.J. Watts, *Phys. Lett. A* **263**, 341 (1999).
42. R. Monasson, *Eur. Phys. J. B* **12**, 555 (1999).
43. B. Kozma and G. Korniss, in *Computer Simulation Studies in Condensed Matter Physics XVI*, edited by D.P. Landau, S.P. Lewis, and H.-B. Schüttler, Springer Proceedings in Physics (Springer, Berlin, 2004, in press); arXiv:cond-mat/0305025 (2003).
44. B. Kozma, M.B. Hastings, and G. Korniss, arXiv:cond-mat/0309196: *Phys. Rev. Lett.*, in press (2004).

45. Réka Albert and Albert-László Barabási, *Rev. Mod. Phys.* **74**, 47 (2002).
46. H. Guclu, G. Korniss, M.A. Novotny, and Z. Toroczkai, in preparation.
47. A.A. Middleton, *Phys. Rev. Lett.* **68**, 670 (1992).
48. M.B. Hastings, private communications (2003).
49. R.A. Fisher and L.H.C. Tippett, *Proc. Camb. Philos. Soc.* **24**, 180 (1928)
50. E.J. Gumbel, *Statistics of Extremes* (Columbia University Press, New York, 1958).
51. *Extreme Value Theory and Applications*, edited by J. Galambos, J. Lechner, and E. Simin (Kluwer, Dordrecht, 1994).
52. Y. Rabani, A. Sinclair, and R. Wanka, *Proceedings of the 39th Annual Symposium on Foundations of Computer Science* (IEEE Comput. Soc, Los Alamitos, CA, 1998) pp. 694-703.
53. H. Guclu and G. Korniss, arXiv:cond-mat/0311575 (2003).
54. S.T. Bramwell, P.C.W. Holdsworth, and J.-F. Plinton, *Nature* **396** 552 (1998).
55. S.T. Bramwell *et al.*, *Phys. Rev. Lett.* **84** 3744 (2000).
56. S.T. Bramwell *et al.*, *Phys. Rev. E* **63** 041106 (2001).
57. T. Antal, M. Droz, G. Györgyi, and Z. Rácz, *Phys. Rev. Lett.* **87**, 240601 (2001).
58. S. C. Chapman, G. Rowlands, and N. W. Watkins, *Nonlinear Processes in Geophysics* **9**, 409 (2002); arXiv:cond-mat/0106015.
59. V. Aji and N. Goldenfeld, *Phys. Rev. Lett.* **86**, 1007 (2001).
60. J.-P. Bouchaud and M. Mézard, *J. Phys. A* **30**, 7997 (1997).
61. A. Baldassarri, *Statistics of Persistent Extreme Events*, Ph.D. Thesis, De l'Université Paris XI Orsay (2000).
62. H. Guclu and G. Korniss, in preparation.
63. A.-L. Barabási and R. Albert, *Science* **286** 509 (1999).
64. A. Barrat, M. Weigt, *Eur. Phys. J. B* **13**, 547 (2000).
65. M. Gitterman, *J. Phys. A* **33**, 8373 (2000).
66. B.J. Kim *et al.*, *Phys. Rev. E* **64**, 056135 (2001).
67. H. Hong, B.J. Kim, M.Y. Choi, *Phys. Rev. E* **66**, 018101 (2002).
68. H. Hong, M.Y. Choi, and B.J. Kim, *Phys. Rev. E* **65**, 047104 (2002).
69. C.P. Herrero, *Phys. Rev. E* **65**, 066110 (2002).
70. M.A. Novotny and Shannon M. Wheeler, *Braz. J. Phys., Proceedings for the III Brazilian Meeting on Simulational Physics* (2003); arXiv:cond-mat/0308602 (2003).
71. M.B. Hastings, *Phys. Rev. Lett.* **91**, 098701 (2003).