Extreme Fluctuations in Small-World-Coupled Autonomous Systems with Relaxational Dynamics

H. Guclu and G. Korniss

Department of Physics, Applied Physics, and Astronomy, Rensselaer Polytechnic Institute, 110 8^{th} Street, Troy, New York 12180–3590, USA

ABSTRACT

Synchronization is a fundamental problem in natural and artificial coupled multi-component systems. We investigate to what extent small-world couplings (extending the original local relaxational dynamics through the random links) lead to the suppression of extreme fluctuations in such systems. We use the framework of non-equilibrium surface growth to study and characterize the degree of synchronization in the system. In the absence of the random links, the surface in the steady state is "rough" (strongly de-synchronized state) and the average and the extreme height fluctuations diverge in the same power-law fashion with the system size (number of nodes). With small-world links present, the average size of the fluctuations becomes finite (synchronized state) and the extreme heights diverge only logarithmically in the large system-size limit. This latter property ensures synchronization in a practical sense in coupled multi-component autonomous systems with short-tailed noise and effective relaxation through the links. The statistics of the extreme heights is governed by the Fisher-Tippett-Gumbel distribution. We illustrate our findings through an actual synchronization problem in parallel discrete-event simulations.

Keywords: extreme-value statistics, synchronization, small-world networks, scalable computing

1. INTRODUCTION

Many of our important technological, information, and infrastructure systems form a complex network[1–6] with a large number of components. The network consists of nodes (components of the system) and links connecting the nodes. The links facilitate some kind of effective interaction/dynamics between the nodes. Examples (with the processes inducing the interaction between the nodes) include high-performance scalable parallel or grid-computing networks (synchronization protocols for massive parallelization),[6] load-balancing schemes (relocating jobs among processors),[7] the Internet (protocols for sending/receiving packets),[8] or the electric power grid (generating/transmitting power between generators and buses).[5] Many of theses systems are autonomous (by design or historical evolution), i.e., they lack a central regulator. Thus, fluctuations in the "load" in the respective network (data/state savings or task allocation in parallel simulations, traffic in the Internet, or voltage/phase in the electric grid) are determined by the collective result of the individual decisions of many interacting "agents" (nodes). As the number of processors on parallel architectures increases to hundreds of thousands,[9] grid-computing networks proliferate over the Internet,[10,11] or the electric power-grid covers, e.g., the North-American continent,[5] fundamental questions on the corresponding dynamical processes on the respective underlying networks must be addressed.

Typically, large fluctuations in the above networks are to be avoided (e.g., for scalability or stability reasons). In the absence of global intervention or control, this can be a difficult task. Motivated by a recent example[6] for small-world (SW)[12] synchronized autonomous systems in the context of scalable parallel computing, we investigate the steady-state properties of the extreme fluctuations in SW-coupled interacting systems with relaxational dynamics. Since the introduction of SW networks[12] it has been well established that such networks can facilitate autonomous synchronization.[13,14] In addition to the average "load" in the network, knowing the typical size

Further author information: (Send correspondence to G.K)
H.G.: E-mail: gucluh@rpi.edu
G.K.: E-mail: korniss@rpi.edu

Noise in Complex Systems and Stochastic Dynamics II, edited by Zoltán Gingl,
José M. Sancho, Lutz Schimansky-Geier, Janos Kertesz, Proc. of SPIE Vol. 5471
(SPIE, Bellingham, WA, 2004) · 0277-786X/04/$15 · doi: 10.1117/12.546056

and the distribution of the extreme fluctuations[15–17] is of great importance from a system-design viewpoint, since failures and delays are triggered by extreme events occurring on an individual node.

Relationship between extremal statistics and universal fluctuations in correlated systems has been studied intensively.[18–29] The focus of these studies was to find connections between the probability distribution of global observables (such as the width in surface growth problems[30]) and known universal extreme-value limit distributions. Here we discuss to what extent SW couplings (extending the original dynamics through the random links) lead to the suppression of the extreme fluctuations. We illustrate our findings on an actual synchronization problem in scalable parallel computing.[6] In Sec. 2 we briefly review the well-known extreme-value limit distribution for short-tail distributed random variables. In Sec. 3 we employ recent results[31] for the Edwards-Wilkinson model[32] on SW networks to obtain the scaling behavior of the extreme fluctuations and their distribution. In Sec. 4 we apply these results to study the extreme load fluctuations in SW-synchronized parallel discrete-event simulation (PDES) schemes,[33, 34] applicable to high performance parallel architectures and large-scale grid-computing networks. We finish the paper with a brief summary and outlook in Sec. 5.

2. A BRIEF REVIEW OF THE EXTREME-VALUE LIMIT DISTRIBUTION FOR INDEPENDENT SHORT-TAILED VARIABLES

Here we consider the case when the individual complementary cumulative distribution $P_>(x)$ (the probability that the individual stochastic variable is greater than x) decays faster than any power law, i.e., exhibits an exponential-like tail in the asymptotic large-x limit. (Note that in this case the corresponding probability density function displays the same exponential-like asymptotic tail behavior.) We will assume $P_>(x) \simeq e^{-cx^\delta}$ for large x values, where c and δ are constants. Then the cumulative distribution $P_<^{\max}(x)$ for the largest of the N events (the probability that the maximum value is less than x) can be approximated as[29, 35, 36]

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

where one typically assumes that the dominant contribution to the statistics of the extremes comes from the tail of the individual distribution $P_>(x)$. With the exponential-like tail in the individual distribution, this yields

$$P_<^{\max}(x) \simeq e^{-e^{-cx^\delta + \ln(N)}} . \tag{2}$$

The extreme-value limit theorem states that there exists a sequence of scaled variables $\tilde{x} = (x - a_N)/b_N$, such that in the limit of $N \to \infty$, the extreme-value probability distribution for $\tilde{x}$ asymptotically approaches the Fisher-Tippett-Gumbel (FTG) distribution[15, 16]:

$$\tilde{P}_<^{\max}(\tilde{x}) \simeq e^{-e^{-\tilde{x}}} , \tag{3}$$

with mean $\langle\tilde{x}\rangle=\gamma$ (γ=0.577... being the Euler constant) and variance $\sigma_{\tilde{x}}^2=\langle\tilde{x}^2\rangle-\langle\tilde{x}\rangle^2=\pi^2/6$. From Eq. (2), one can deduce[36, 37] that to leading order the scaling coefficients are $a_N = \left[\frac{\ln(N)}{c}\right]^{1/\delta}$ and $b_N = (\delta c)^{-1} \left[\frac{\ln(N)}{c}\right]^{(1/\delta)-1}$. The average value of the largest of the N original variables then scales as

$$\langle x_{\max}\rangle = a_N + b_N\gamma \simeq \left[\frac{\ln(N)}{c}\right]^{1/\delta} \tag{4}$$

(up to $\mathcal{O}(\frac{1}{\ln(N)})$ correction) in the asymptotic large-N limit. When comparing with experimental or simulation data, instead of Eq. (3), it is often convenient to use the form of the FTG distribution which is scaled to zero mean and unit variance, yielding

$$\tilde{P}_<^{\max}(y) = e^{-e^{-(ay+\gamma)}} , \tag{5}$$

where $a = \pi/\sqrt{6}$ and γ is the Euler constant. In particular, the corresponding FTG density then becomes

$$\tilde{p}^{\max}(y) = ae^{-(ay+\gamma)-e^{-(ay+\gamma)}} . \tag{6}$$

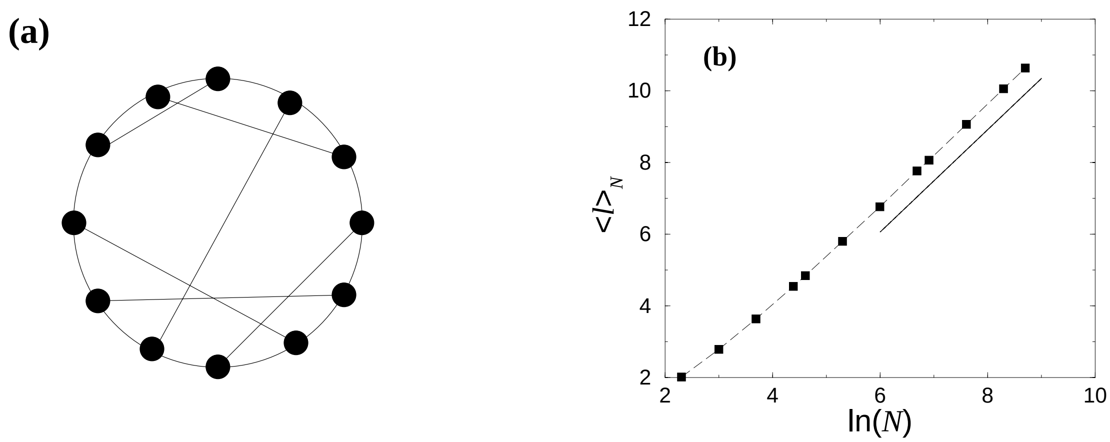

Figure 1. (a) A small-world network where random links are added to the ring, such that each node has exactly one random link. (b) Average shortest path as a function of the logarithm of the number of nodes for our small-world synchronization network shown in (a). The straight line represents the slope of the asymptotic large N behavior of the average shortest path $\langle l \rangle_N \simeq 1.42 \ln(N)$.

3. EXTREME FLUCTUATIONS IN THE SMALL-WORLD-COUPLED EDWARDS-WILKINSON MODEL

We consider the simplest stochastic model with linear relaxation on a SW network,

$$\partial_t h_i(t) = -(2h_i - h_{i+1} - h_{i-1}) - \sum_{j=1}^{N} J_{ij}(h_i - h_j) + \eta_i(t) \ , \tag{7}$$

where $h_i(t)$ is the local height or field variable at node i at time t and $\eta_i(t)$ is a delta-correlated short-tailed (e.g., Gaussian) noise. The symmetric matrix J_{ij} (with matrix elements being equal to 0 or p) represents the quenched random links of strength p on top of a one dimensional regular lattice. In the construction of the SW network presented here, each node has exactly one random neighbor [Fig. 1(a)]. That is, pairs of nodes are selected at random, and once they are linked, they cannot be selected again. This construction is motivated by our application[6, 38] to scalable PDES schemes (see Sec. 4), where fluctuations in the individual degree of the nodes are to be avoided. Our construction of the SW network differs from both the original ("rewiring")[12, 13] and the "soft" version[31, 39–41] of the SW network (where an Erdős-Rényi random graph[42] 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). Further, the average path length $\langle l \rangle_N$ (the average minimum number of links connecting two randomly chosen nodes) scales *logarithmically* with the system size N [Fig. 1(b)], i.e., like most other random networks,[1] it too exhibits the "small-world" property (or low-degree of separation).

Equation (7), the extension of the the Edwards-Wilkinson (EW) model to a SW "substrate", where the strength of the interactions through the random links is p, is a prototypical synchronization problem with "local" relaxation. The width

$$w \equiv \sqrt{\left\langle \frac{1}{N} \sum_{i=1}^{N} (h_i - \bar{h})^2 \right\rangle} \ , \tag{8}$$

borrowing the framework from non-equilibrium surface-growth phenomena, provides a sensitive measure for the average degree of synchronization in coupled multi-component systems.[6, 43] In Eq. (8) $\langle \ldots \rangle$ denotes an ensemble average over the noise in Eq. (7). In addition to the width, we will study the scaling behavior the largest fluctuations (e.g., above the mean) in the steady-state

$$\langle \Delta_{\max} \rangle \equiv \langle h_{\max} - \bar{h} \rangle \ . \tag{9}$$

Equation (7) (and its generalization with a Kardar-Parisi-Zhang (KPZ)-like nonlinearity[44]) is also believed to govern the steady-state progress and scalability properties of a large class of PDES schemes.[6, 43, 45–47] In this context, the local height variables $\{h_i(t)\}_{i=1}^{N}$ correspond to the progress of the individual processors after t parallel steps (Sec. 4). The EW/KPZ-type relaxation at a coarse-grained level originates from the "microscopic" (node-to node) synchronizational rules. In the absence of the random links with purely short-range connections, the corresponding steady-state landscape is rough[30] (de-synchronized state), i.e., it is dominated by large-amplitude long-wavelength fluctuations. The extreme values of the local fluctuations emerge through these long-wavelength modes and the extreme and average fluctuations follow the *same* power-law divergence with the system size[18, 46, 47]

$$\langle \Delta_{\max} \rangle \sim w \sim N^{\alpha} \ , \tag{10}$$

where α is the roughness exponent[30] [Figs. 2(a) and 3(a)]. The diverging width is related to an underlying diverging lengthscale, the lateral correlation length, which reaches the system size N for a finite system. In PDES schemes the average local memory requirement on each node is proportional to the spread of the progress of the individual processors (the width of the landscape of the progress of the simulation). Thus, a diverging width (strongly de-synchronized state) [Fig.2(a)] can seriously hinder scalable data management,[46, 47] motivating the implementation of a SW synchronization network[6] (Sec. 4).

The important feature of the EW model on SW networks is the development of an effective nonzero mass $\Sigma(p)$, corresponding to an actual or pseudo gap in a field theory sense,[31, 41, 48] generated by the quenched-random structure.[31] In turn, both the correlation length $\xi \simeq [\Sigma(p)]^{-1/2}$ and the width $w \simeq (1/\sqrt{2})[\Sigma(p)]^{-1/4}$ approach a finite value (synchronized state) and become self-averaging in the $N\to\infty$ limit.[38] For example, for our above described construction of the SW network,[31] for small p values, $\Sigma(p) \sim p$. Thus, the correlation length becomes *finite* for an arbitrarily small but nonzero strength of the random links (one such link per site). This is the fundamental effect of extending the original dynamics to a SW network: it decouples the fluctuations of the originally correlated system. Then, the extreme-value limit theorems can be applied using the number of independent blocks N/ξ in the system.[29, 36] Further, if the tail of the noise distribution decays in an exponential-like fashion, the individual relative height distribution will also do so,[49] and depends on the combination Δ_i/w, where $\Delta_i = h_i - \bar{h}$ is the relative height measured from the mean at site i. Considering, e.g., the fluctuations above the mean for the individual sites, we will then have $P_{>}(\Delta_i) \simeq \exp[-c(\Delta_i/w)^{\delta}]$, where $P_{>}(\Delta_i)$ denotes the "disorder-averaged" (averaged over network realizations) single-site relative height distribution, which becomes independent of the site i for SW networks. From the above it follows that the cumulative distribution for the extreme-height fluctuations relative to the mean $\Delta_{\max} = h_{\max} - \bar{h}$, if scaled appropriately, will be given by Eq. (3) [or alternatively by Eq. (5)] in the asymptotic large-N limit (such that $N/\xi \gg 1$). Further, from Eq. (4), the average maximum relative height will scale as

$$\langle \Delta_{\max} \rangle \simeq w \left[\frac{\ln(N/\xi)}{c} \right]^{1/\delta} \simeq \frac{w}{c^{1/\delta}} \left[\ln(N) \right]^{1/\delta} \ , \tag{11}$$

where we kept only the leading order term in N. Note, that both w and ξ approach their *finite* asymptotic N-independent values for SW-coupled systems. Also, the same logarithmic scaling with N holds for the largest relative deviations below the mean $\langle \bar{h} - h_{\min} \rangle$ and for the maximum spread $\langle h_{\max} - h_{\min} \rangle$. This is the central point of this paper: in SW synchronized systems with unbounded local variables driven by exponential-like noise distribution (such as Gaussian), the extremal fluctuations increase only *logarithmically* with the number of nodes. This weak divergence, which one can regard as marginal, ensures synchronization for practical purposes in SW coupled multi-component systems with local relaxation in a noisy environment.

4. APPLICATION TO SCALABLE PARALLEL DISCRETE-EVENT SIMULATIONS ON HIGH-PERFORMANCE PARALLEL AND GRID COMPUTING NETWORKS

Developing and implementing massively parallel algorithms is among the most challenging areas in computer science and in computational science and engineering.[50] While there are numerous technological and hardware-related points, e.g., concerning efficient message passing and fast communications between computer nodes, the

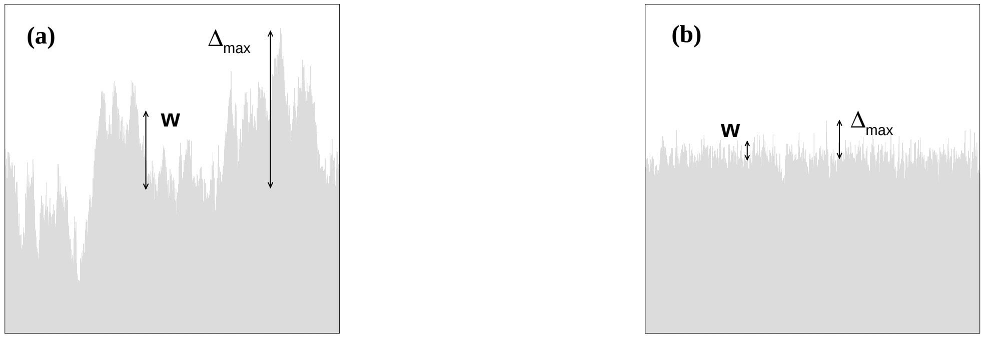

Figure 2. Snapshots of the virtual time horizon for the conservative PDES scheme with N=10^4 processors in the steady state. (a) The processors are connected in a ring-like fashion; (b) the processors are connected by a small-world topology and the additional synchronization through the random link is performed with probability p=0.10 at every parallel step. The vertical scales indicating the progress of the individual processing elements are the same in (a) and (b). The arrows indicate the average (w) and the extreme ($\Delta_{\max}$) fluctuations in virtual time horizon.

theoretical algorithmic challenge is often as important. This is particularly true for cases when the parallel algorithm has to simulate the time evolution of a complex system in which the local changes (discrete events) in the configuration are inherently *asynchronous.* The basic notion of the above discrete-event systems is that time is continuous and the changes in the local configurations occur at random instants of time (hence the asynchrony of the time evolution of the local configuration). Between events, the local configuration remains unchanged. In physics or chemistry these types of simulations are most commonly referred to as dynamic or kinetic Monte Carlo simulations.[51] In computer science they are called discrete-event simulations. PDES schemes[33,34,52] are capable of faithfully simulating such systems in a massively parallel fashion. For very large interacting systems (where trivial or "embarrassing" parallelization is not possible or highly inefficient due to CPU/memory limitations), PDES is the only way to perform parallel simulations *without* changing the original underlying asynchronous dynamics. Examples of PDES applications include dynamic channel allocation in cell phone communication networks,[53,54] models of the spread of diseases,[55] and dynamic phenomena in highly anisotropic magnetic thin films.[56–58] In these examples the discrete events are call arrivals, infections, and changes of the orientation of the local magnetic moments, respectively.

The difficulty of parallel discrete-event simulations is that the local changes (updates) in the system are not synchronized by a global clock. The essence of the corresponding PDES schemes, capable of faithfully simulating these systems, is to algorithmically parallelize "physically" non-parallel dynamics of the underlying systems. This requires some kind of synchronization to ensure causality.[33] The two basic ingredients of PDES schemes is a set of local simulated times (or virtual times[59]) and a synchronization scheme. First, a scalable parallel scheme must ensure that the average progress rate of the simulation approaches a non-zero asymptotic value in the long-time limit as the number of processors (or nodes) goes to infinity. Second, the spread of the virtual time horizon (the spread of the progress of the individual processors) should be bounded as the number of processors goes to infinity.[60] The second requirement is crucial for the measurement phase of the simulation to be scalable: a diverging spread of the virtual time horizon (as the number of processors goes to infinity) hinders scalable data management.[46,47] Temporarily storing a large amount of simulated data on each node (proportional to the spread of the virtual time horizon) is limited by available memory, while frequent global synchronizations can get computationally costly for a large number of nodes on certain parallel architectures. In the latter case, one aims to devise a parallel scheme where the processors make a nonzero close-to-uniform progress *without global synchronization.* In such a scheme, the processors autonomously learn the global state of the system (without explicit global messages) and adjust their progress rate accordingly.[6]

PDES algorithms concurrently advance the local simulated time on each processor [or processing element (PE)], without violating causality. In a "conservative" PDES scheme,[61–64] only those PEs that are guaranteed

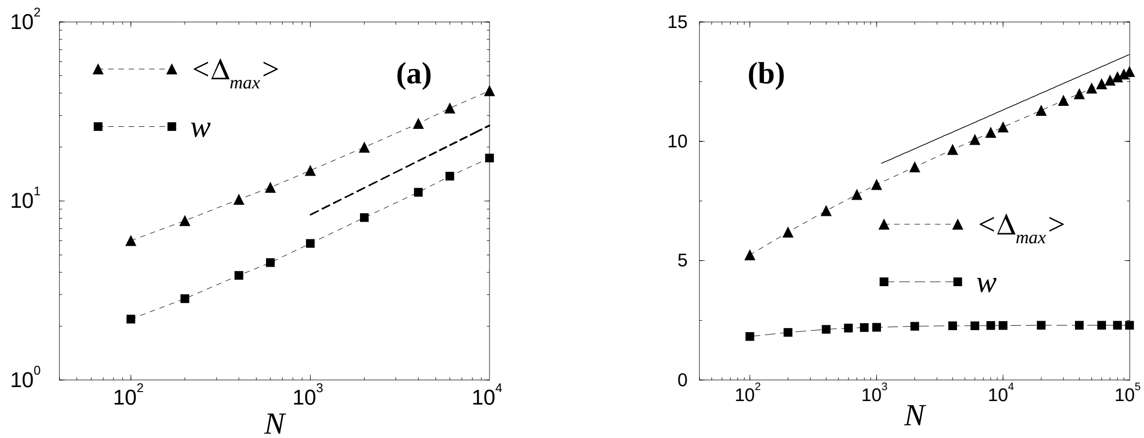

Figure 3. (a) Scaling behavior of the average (w) and the extreme ($\Delta_{\max}$) fluctuations in the virtual time horizon for the conservative PDES scheme in the steady state. The processors are connected in a ring-like fashion (log-log scales). The dashed line represents the theoretical power law with the roughness exponent α=1/2. (b) The same quantities as in (a), but the processors are connected by a small-world topology and the additional synchronization through the random link is performed with probability p=0.10 at every parallel step (log-normal scales). The solid straight line indicates the weak logarithmic increase of the extreme fluctuations with the system size.

not to violate causality are allowed to process their events and increment their local time. The rest of the PEs must "idle." In an "optimistic" approach,[59] the processors do not have to idle, but since causality is not guaranteed at every update, the simulated history on certain processors can become corrupted. This requires a complex "rollback" protocol to correct erroneous computations. Both simulation approaches lead to an evolving and fluctuating time horizon during algorithmic execution. Similar to our earlier results[43] in finding a connection between certain conservative PDES schemes[63, 64] and kinetic roughening in nonequilibrium surfaces,[30, 44, 65] a "complex system" approach was also successful to establish the connection[66, 67] between rollback-based (or optimistic) PDES schemes[59] and self-organized criticality.[68, 69] In what follows, we will focus of the synchronizability of conservative PDES schemes, in particular, the behavior of the width and the largest fluctuations of the virtual time horizon.

Consider an arbitrary one-dimensional system with nearest-neighbor interactions, in which the discrete events (update attempts in the local configuration) exhibit Poisson asynchrony. In the one-site-per-PE scenario, each PE has its own local simulated time, constituting the virtual time horizon $\{h_i(t)\}_{i=1}^{N}$ (essentially, the progress of the individual nodes). Here t is the discrete number of parallel steps executed by all PEs, which is proportional to the wall-clock time and N is the number of PEs. According to the basic conservative synchronization scheme,[63, 64] at each parallel step t, only those PEs for which the local simulated time is not greater then 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.) Thus, denoting the virtual neighborhood of PE i by S_i, if $h_i(t) \leq \min_{j\epsilon S_i}\{h_j(t)\}$, 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.[63, 64] In the original algorithm,[63, 64] the virtual communication topology between PEs mimics the interaction topology of the underlying system. For example, for a one-dimensional system with nearest-neighbor interactions, the virtual neighborhood of PE i, S_i, consists of the left and right neighbor, PE $i{-}1$ and PE $i{+}1$. It was shown[43] that then the virtual time horizon exhibits KPZ-like kinetic roughening and the steady-state behavior in one dimension is governed by the EW Hamiltonian. Thus, both the average (the spread in the progress of the individual PEs) and the extreme fluctuations of the virtual time horizon diverge when $N\to\infty$ [Figs. 2(a) and 3(a)], hindering efficient data collection in the measurement phase of the simulation.[46] To achieve a near-uniform progress of the PEs *without* employing frequent global synchronizations, it was shown[6]

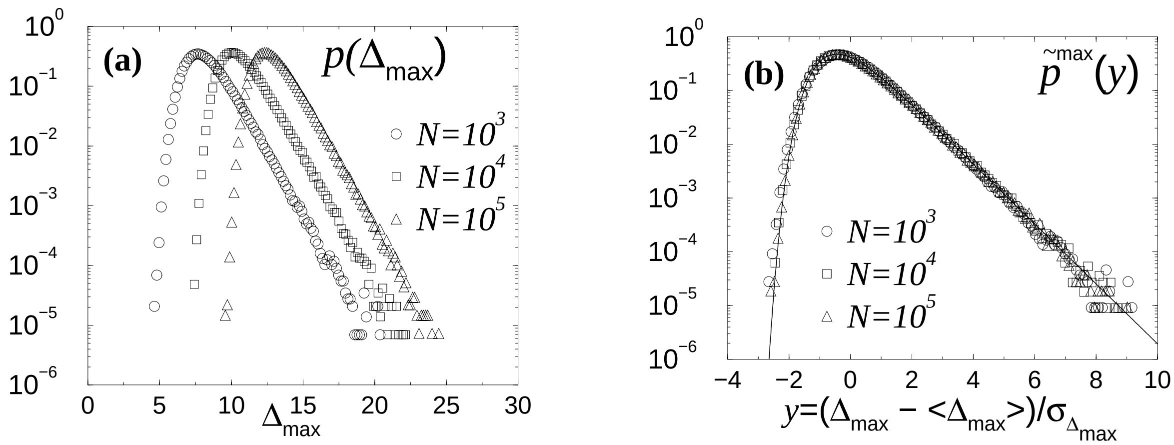

Figure 4. (a) Disorder-averaged probability densities for the extreme-height fluctuations for the SW-synchronized conservative PDES time horizons with p=0.10 for three system sizes indicated in the figure. Note the log-normal scales. (b) The same as (a) but the probability densities are scaled to zero mean and unit variance. The solid curve corresponds to the similarly scaled FTG density Eq. (6) for comparison.

that including randomly chosen PEs (*in addition* to the nearest neighbors) in the virtual neighborhood, results in a finite average width [Figs. 2(b) and 3(b)]. Here we demonstrate that SW synchronization in the above PDES scheme results in logarithmically increasing extreme fluctuations in the simulated time horizon, governed by the FTG distribution.

In the SW synchronized version of the conservative PDES scheme each PE has exactly one random neighbor (in addition to the nearest neighbors) and the local simulated time of the random neighbor is checked only with probability p at every simulation step. Thus, the effective "strength" of the random links is controlled by the relative frequency p of the basic synchronizational steps through those links. 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.[6]

To study the extreme fluctuations of the SW-synchronized virtual time-horizon, we "simulated the simulations", i.e., the evolution of the local simulated times based on the above exact algorithmic rules.[70] By constructing histograms for Δ_i, we observed that the tail of the disorder-averaged individual relative-height distribution decays exponentially (δ=1). Then, we constructed histograms for the extreme-height fluctuations Fig. 4(a). The scaled histograms, together with the similarly scaled FTG density Eq. (6), are shown in Fig. 4(b). We also observed that the distribution of the extreme values becomes *self-averaging*, i.e., independent of the network realization. Figure 3(b) shows that for sufficiently large N (such that w essentially becomes system-size independent) the average (or typical) size of the extreme-height fluctuations diverge *logarithmically*, according to Eq. (11) with δ=1. We also found that the largest relative deviations below the mean $\langle\bar{h}-h_{\min}\rangle$, and the maximum spread $\langle h_{\max}-h_{\min}\rangle$ follow the same scaling with the system size N. Note, that for our specific system (PDES time horizon), the "microscopic" dynamics is inherently non-linear, but the effects of the non-linearities only give rise to a renormalized mass $\Sigma(p)$ (leaving $\Sigma(p)>0$ for all $p>0$).[6] Thus, the dynamics is effectively governed by relaxation in a small world, yielding a finite correlation length and, consequently, the slow logarithmic increase of the extreme fluctuations with the system size [Eq. (11)]. Also, for the PDES time horizon, the local height distribution is asymmetric with respect to the mean, but the average size of the height fluctuations is, of course, finite for both above and below the mean. This specific characteristic simply yields different prefactors for the extreme fluctuations [Eq. (11)] above and below the mean, leaving the logarithmic scaling with N unchanged.

5. SUMMARY AND OUTLOOK

We considered the extreme-height fluctuations in a prototypical model with local relaxation, unbounded local variables, and short-tailed noise. We argued, that when the interaction topology is extended to include random links in a SW fashion, the statistics of the extremes is governed by the FTG distribution. This finding directly addresses synchronizability in generic SW-coupled systems where relaxation through the links is the relevant node-to-node process and effectively governs the dynamics. We illustrated our results on an actual synchronizational problem in the context of scalable parallel simulations. Analogous questions for heavy-tailed noise distribution and different types of networks have relevance to various transport and transmission phenomena in natural and artificial networks[71] and to the corresponding discrete-event systems with fat-tail (non-Poisson) asynchrony. For example, in Internet or WWW traffic, in part, as a result of universal "heavy-tailed" file-size distributions,[72,73] service times exhibit power-law distributions.[74–76] In turn, when simulating these systems, the corresponding PDES should use power-law tail distributed local simulated time increments. This will correspond to a surface-growth problem where, the "substrate" is a complex network, and the noise is power-law distributed. Heavy-tailed noise typically generates similarly tailed local field variables through the collective dynamics. Then, the largest fluctuations can still diverge as a power law with the system size (governed by the Fréchet distribution[16,17]), motivating further research for the properties of extreme fluctuations and synchronizability in complex networks.

From a broader statistical physics viewpoint, the lines of investigations we pursue contribute not only to scalability and synchronizability, but also to general studies of collective phenomena on SW,[12,13,39–41,77–84] and on scale-free[3,85–92] networks. In particular, there is growing evidence that systems without inherent frustration exhibit (strict or anomalous)[31,84] mean-field-like behavior when the original short-range interaction topology is modified to a SW network.[31,77–84]

ACKNOWLEDGMENTS

We thank Z. Rácz, Z. Toroczkai, and M.A. Novotny for comments and discussions. G.K. thanks CNLS LANL for their hospitality during Summer 2003. This research was supported by NSF Grant No. DMR-0113049 and the Research Corp. Grant No. RI0761. H.G. was also supported in part by the LANL summer student program in 2003 through US DOE Grant No. W-7405-ENG-36.

REFERENCES

1. R. Albert and A.-L. Barabási, "Statistical mechanics of complex networks", *Rev. Mod. Phys.* **74**, 47–97, 2002.
2. S.N. Dorogovtsev and J.F.F. Mendes, "Evolution of networks", *Adv. Phys.* **51**, 1079–1187, 2002.
3. A.-L. Barabási and R. Albert, "Emergence of Scaling in Random Networks", Science **286**, 509–512, 1999.
4. M. Faloutsos, P. Faloutsos,and C. Faloutsos, "On Power-Law Relationships of the Internet Topology", Proc. ACM SIGCOMM, Comput. Commun. Rev. **29**, 251–262, 1999.
5. E.J. Lerner "What's Wrong with the Electric Grid?", *The Industrial Physicist*, Oct./Nov. 2003, pp. 8-13.
6. G. Korniss, M.A. Novotny, H. Guclu, Z. Toroczkai, and P.A. Rikvold, "Suppressing Roughness of Virtual Times in Parallel Discrete-Event Simulations", *Science* **299**, 677–679, 2003.
7. Y. Rabani, A. Sinclair, R. Wanka, "Local Divergence of Markov Chains and the Analysis of Iterative Load-Balancing Schemes", in *Proceedings of the 39th Annual Symposium on Foundations of Computer Science* (IEEE Comput. Soc., Los Alamitos, CA, 1998) pp. 694–703.
8. L.L. Peterson and B.S. Davie, *Computer Networks, A Systems Approach*, 2nd edn (Morgan Kaufmann, San Francisco, 2000).
9. *Blue Gene/L project*, partnership between IBM and DoE, announced Nov.9, 2001; 64K processors, expected scale 200teraflops, a step towards a petaflop scale; expected completion 2005; IBM Research Report, RC22570 (W0209-033) September 10, 2002.
10. See, e.g., `www.gridforum.org` and `setiathome.ssl.berkeley.edu`.
11. S. Kirkpatrick, "Rough Times Ahead", *Science* **299**, 668–669, 2003.

12. D.J. Watts and S.H. Strogatz, "Collective Dynamics of 'Small-World' Networks", *Nature* **393**, 440–442, 1998.
13. D.J. Watts, *Small Worlds* (Princeton Univ. Press, Princeton, 1999).
14. S.H. Strogatz, "Exploring Complex Networks", *Nature* **410**, 268–276, 2001.
15. R.A. Fisher and L.H.C. Tippett, "The Frequency Distribution of the Largest or Smallest Member of a Sample", *Proc. Camb. Philos. Soc.* **24**, 180–191, 1928.
16. E.J. Gumbel, *Statistics of Extremes* (Columbia University Press, New York, 1958).
17. *Extreme Value Theory and Applications*, edited by J. Galambos, J. Lechner, and E. Simin (Kluwer, Dordrecht, 1994).
18. S. Raychaudhuri, M. Cranston, C. Przybyla, and Y. Shapir, "Maximal Height Scaling of Kinetically Growing Surfaces", *Phys. Rev. Lett.* **87**, 136101 [4 pages], 2001.
19. S.T. Bramwell, P.C.W. Holdsworth, and J.-F. Plinton, "Universality of Rare Fluctuations in Turbulance and Critical Phenomena", *Nature* **396** 552–554, 1998.
20. S.T. Bramwell, K. Christensen, J.-Y. Fortin, P.C.W. Holdsworth, H.J. Jensen, S. Lise, J.M. López, M. Nicodemi, J.-F. Pinton, and M. Sellitto, "Universal Fluctuations in Correlated System", *Phys. Rev. Lett.* **84**, 3744–3747, 2000.
21. N.W. Watkins, S.C. Chapman, and G. Rowlands, Comment on "Universal Fluctuations in Correlated Systems", *Phys. Rev. Lett.* **89**, 208901, 2002.
22. S.T. Bramwell, K. Christensen, J.-Y. Fortin, P.C.W. Holdsworth, H.J. Jensen, S. Lise, J.M. López, M. Nicodemi, J.-F. Pinton, and M. Sellitto, Reply to Comment on "Universal Fluctuations in Correlated Systems", *Phys. Rev. Lett.* **89**, 208902, 2002.
23. S.T. Bramwell, J.-Y. Fortin, P.C.W. Holdsworth, S. Peysson, J.-F. Pinton, B. Portelli, and M. Sellitto, "Magnetic fluctuations in the classical XY model: The origin of an exponential tail in complex systems", *Phys. Rev.* E **63**, 041106, 2001.
24. V. Aji and N. Goldenfeld, "Fluctuations in finite Critical and Turbulent Systems", *Phys. Rev. Lett.* **86**, 1007–1010, 2001.
25. T. Antal, M. Droz, G. Györgyi, and Z. Rácz, "$1/f$ Noise and Extreme-Value Statistics", *Phys. Rev. Lett.* **87**, 240601 [4 pages], 2001.
26. K. Dahlstedt and H.J. Jensen, "Universal fluctuations and extreme value statistics", *J. Phys. A* **34** 11193–11200, 2001.
27. S. C. Chapman, G. Rowlands, and N. W. Watkins, "Extremum statistics: A framework for data analysis", *Nonlinear Processes in Geophysics* **9**, 409–418, 2002; arXiv:cond-mat/0106015.
28. G. Györgyi, P.C.W. Holdsworth, B. Portelli, and Z. Rácz, "Statistics of extremal intensities for Gaussian interfaces", *Phys. Rev.* E **68**, 056116 [14 pages], 2003.
29. J.-P. Bouchaud and M. Mézard, "Universality Classes for Extreme-Value Statistics", *J. Phys. A* **30**, 7997–8015, 1997.
30. A.-L. Barabási and H.E. Stanley *Fractal Concepts in Surface Growth*, (Cambridge Univ. Press, Cambridge, 1995).
31. B. Kozma, M.B. Hastings, and G. Korniss, "Roughness Scaling for Edwards-Wilkinson Relaxation in Small-World Networks", *Phys. Rev. Lett.* **92**, 108701 [4 pages], 2004.
32. S.F. Edwards and D.R. Wilkinson, "The Surface Statistics of a Granular Aggregate", *Proc. R. Soc.* London, Ser. A **381**, 17–31, 1982.
33. R. Fujimoto, "Parallel Discrete Event Simulation", *Commun. of the ACM* **33**, 30–53, 1990.
34. D.M. Nicol and R.M. Fujimoto, "Parallel Simulation Today", *Ann. Oper. Res.* **53**, 249–285, 1994.
35. J.-P. Bouchaud and M. Potters, *Theory of Financial Risk*, (Cambridge Univ. Press, Cambridge, 2000)
36. A. Baldassarri, *Statistics of Persistent Extreme Events*, Ph.D. Thesis, De l´Université Paris XI Orsay (2000); `http://axtnt3.phys.uniroma1.it/~andreab/these.html`.
37. Note that for $\delta\neq 1$, while the convergence to Eq. (2) is fast, the convergence for the appropriately scaled variable to the universal FTG distribution Eq. (3) is *extremely* slow.[15, 21, 27, 36]

38. H. Guclu, G. Korniss, Z. Toroczkai, and M.A. Novotny, "Small-World Synchronized Computing Networks for Scalable Parallel Discrete-Event Simulations", to appear as a Chapter in Complex Networks, edited by E. Ben-Naim, H. Frauenfelder, and Z. Toroczkai, Lecture Notes in Physics (Springer, 2004).
39. M.E.J. Newman, "Models of a Small World", *J. Stat. Phys.* **101**, 819–841, 2000.
40. M.E.J. Newman and D.J. Watts, "Renormalization Group analysis of the Small-World Network Model", *Phys. Lett.* A **263**, 341–346, 1999.
41. R. Monasson, "Diffusion, Localization and Dispersion Relations on "Small-World" Lattices", *Eur. Phys. J.* B **12**, 555-567 (1999).
42. P. Erdős and A. Rényi, "On the Evolution of Random Graphs", *Publ. Math. Inst. Hung. Acad. Sci.* **5**, 17–61 (1960).
43. 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.
44. M. Kardar, G. Parisi, and Y.-C. Zhang, "Dynamic Scaling of Growing Interfaces", *Phys. Rev. Lett.* **56**, 889–892, 1986.
45. Toroczkai, G. Korniss, S. Das Sarma, and R.K.P. Zia, "Extremal Point Densities of Interface Fluctuations", *Phys. Rev.* E **62**, 276–294, 2000.
46. G. Korniss, M.A. Novotny, A.K. Kolakowska, and H. Guclu, "Statistical Properties of the Simulated Time Horizon in Conservative Parallel Discrete-Event Simulations", SAC 2002, Proceedings of the 2002 ACM Symposium on Applied Computing, pp. 132–138, 2002.
47. G. Korniss, M.A. Novotny, P.A. Rikvold, H. Guclu, and Z. Toroczkai, "Going through Rough Times: from Non-Equilibrium Surface Growth to Algorithmic Scalability", in *Materials Research Society Symposium Proceedings Series*, Vol. 700, Fall Meeting, Boston, 2001, pp. 297-308.
48. B. Kozma and G. Korniss, "Stochastic Growth in a Small World", 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).
49. The exponent δ for the tail of the local relative height distribution may differ from that of the noise as a result of the collective (possibly non-linear) dynamics, but the exponential-like feature does not change.
50. K. Hwang, *Advanced Computer Architecture: Parallelism, Scalability, and Programmability* (McGraw–Hill, New York, 1993).
51. K. Binder and D.W. Heermann, *Monte Carlo Simulation in Statistical Physics : An Introduction* (Springer, 1998).
52. B.D. Lubachevsky, "Fast simulation of Multicomponent Dynamic Systems", *Bell Labs Tech. J.* **5** April-June 2000, 134–156, 2000.
53. A.G. Greenberg, B.D. Lubachevsky, D.M. Nicol, and P.E. Wright, "Efficient Massively Parallel Simulation of Dynamic Channel Assignment Schemes for Wireless Cellular Communications", in *Proc. 8th Workshop on Parallel and Distributed Simulation (PADS'94), Edinburgh, UK, 1994* (SCS, San Diego, CA, 1994), pp. 187–194.
54. S.C. Borst, S.A. Grandhi, C.L. Kahn, K. Kumaran, B.D. Lubachevsky, and D.M. Sand, "Wireless Simulation and Self-Organizing Spectrum Management", *Bell Labs Tech. J.* **2**, 81–98, 1997.
55. E. Deelman, B.K. Szymanski, T. Caraco, "Simulating Lyme Disease Using Parallel Discrete-Event Simulations", in *Proc. 28th Winter Simulation Conference*, (ACM, New York, 1996), pp. 1191–1198.
56. 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.
57. G. Korniss, C.J. White, P.A. Rikvold, M.A. Novotny, "Dynamic Phase Transition, Universality, and Finite-size Scaling in the Two-dimensional Kinetic Ising Model in an Oscillating Field", *Phys. Rev.* E **63**, 016120 [15 pages], 2001.
58. G. Korniss, P.A. Rikvold, and M.A. Novotny, "Absence of First-order Transition and Tricritical Point in the Dynamic Phase Diagram of a Spatially Extended Bistable System in an Oscillating Field", *Phys. Rev.* E **66**, 056127 [12 pages], 2002.
59. D.R. Jefferson, "Virtual Time", *Assoc. Comput. Mach. Trans. Programming Languages and Systems* **7**, 404–425, 1985.

60. A.G. Greenberg, S. Shenker, and A.L. Stolyar, "Asynchronous Updates in Large Parallel Systems", *Performance Eval. Rev.* **24**, 91–103, 1996.
61. 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.
62. K.M. Chandy, J. Misra, "Asynchronous Distributed Simulation via a Sequence of Parallel Computations", *Commun. ACM* **24**, 198–205, 1981.
63. B.D. Lubachevsky, "Efficient Parallel Simulations of Asynchronous Cellular Arrays", *Complex Systems* **1**, 1099–1123, 1987.
64. B.D. Lubachevsky, "Efficient Parallel Simulations of Dynamic Ising Spin Systems", *J. Comput. Phys.* **75**, 103–122, 1988.
65. M. Plischke, Z. Rácz, and D. Liu, "Time-reversal invariance and universality of two-dimensional growth models", *Phys. Rev.* B **35**, 3485–3495, 1987.
66. B.J. Overeinder, "Distributed Event-driven Simulation: Scheduling Strategies and Resource Management", Ph.D. thesis, Universiteit van Amsterdam, 2000.
67. P.M.A. Sloot, B.J. Overeinder, A. Schoneveld, "Self-organized Criticality in Simulated Correlated Systems", *Comput. Phys. Commun.* **142**, 76–81, 2001.
68. P. Bak, C. Tang, K. Wiesenfeld, "Self-Organized Criticality: An Explanation of $1/f$ Noise", *Phys. Rev. Lett.* **59**, 381–384, 1987.
69. P. Bak, C. Tang, and Wiesenfeld, "Self-Organized Criticality", *Phys. Rev.* A **38**, 364–374, 1988.
70. H. Guclu and G. Korniss, "Extreme Fluctuations in Small-Worlds with Relaxational Dynamics", arXiv:cond-mat/0311575, 2003.
71. M. Argollo de Menezes and A.-L. Barabási "Fluctuations in Network Dynamics" *Phys. Rev. Lett.* **92**, 028701 [4 pages], 2004.
72. M.E. Crovella and A. Bestavros, "Self-Similarity in World Wide Web Traffic: Evidence and Possible Causes", *IEEE/ACM Trans. on Networking* **5**, 835–846, 1997.
73. M.E. Crovella, M.S. Taqqu, and A. Bestavros, "Heavy-Tailed Probability distributions in the World Wide Web", in *A Practical Guide To Heavy-Tails: Statistical Techniques and Applications*, edited by R.J. Adler, R.E. Feldman, and M.S. Taqqu (Birkhäuser, Boston, 1998), pp 3-25.
74. W.E. Leland, M.S. Taqqu, W. Willinger, and D.V. Wilson, "On the Self-Similar Nature of Ethernet Traffic", *IEEE/ACM Trans. on Networking* **2**, 1–15, 1994.
75. I. Csabai, "$1/f$ noise in Computer Network Traffic", *J. Phys.* A **27**, L417–421, 1994.
76. V. Paxson and S. Floyd, "Wide Area Traffic: The Failure of Poisson Modeling", *IEEE/ACM Trans. on Networking* **3**, 226–244, 1995.
77. A. Barrat and M. Weigt, "On the properties of small-world network models", *Eur. Phys.* J. B **13**, 547–560, 2000.
78. M. Gitterman, "Small-world phenomena in physics: the Ising model", *J. Phys.* A **33**, 8373–8381, 2000.
79. B.J. Kim, H. Hong, P. Holme, G.S. Jeon, P. Minnhagen, and M.Y. Choi, "XY model in small-world networks", *Phys. Rev.* E **64**, 056135 [5 pages], 2001.
80. Comment on "Ising model on a small-world network", H. Hong, B.J. Kim, and M.Y. Choi, *Phys. Rev.* E **66**, 018101 [2 pages], 2002.
81. H. Hong, M.Y. Choi, and B.J. Kim, "Phase ordering on small-world networks with nearest-neighbor edges", *Phys. Rev.* E **65**, 047104 [4 pages, 2002.
82. C.P. Herrero, "Ising model in small-world networks", *Phys. Rev.* E **65**, 066110 [6 pages], 2002.
83. D. Jeong, H. Hong, B.J. Kim, and M.Y. Choi, "Phase transition in the Ising model on a small-world network with distance-dependent interactions", *Phys. Rev.* E **68**, 027101 [4 pages], 2003.
84. M.B. Hastings, "Mean-Field and Anomalous Behavior on a Small-World Network", *Phys. Rev. Lett.* **91**, 098701 [4 pages], 2003.
85. A.V. Goltsev, S.N. Dorogovtsev, and J.F.F. Mendes, "Critical phenomena in networks", *Phys. Rev.* E **67**, 026123 [5 pages], 2003.
86. M. Leone, A. Vázquez, A. Vespignani, and R. Zecchina, "Ferromagnetic ordering in graphs with arbitrary degree distribution", *Eur. Phys. J.* B **28**, 191-197, 2002.

87. S.N. Dorogovtsev, A.V. Goltsev, and J.F.F. Mendes, “Ising model on networks with an arbitrary distribution of connections”, *Phys. Rev.* E **66**, 016104 [5 pages], 2002.
88. G. Bianconi, “Mean-field solution of the Ising model on a Barabási-Albert network”, *Phys. Lett.* A **303**, 166–169, 2002.
89. A. Aleksiejuk, J. Holyst, and D. Stauffer, “Ferromagnetic Phase Transition in Barabási-Albert Networks”, *Physica* A **310**, 260–266, 2002.
90. F. Igloi and L. Turban “First- and second-order phase transitions in scale-free networks”, *Phys. Rev.* E **66**, 036140 [4 pages], 2002.
91. S.N. Dorogovtsev, A.V. Goltsev, and J.F.F. Mendes, “Potts model on complex networks”, arXiv:cond-mat/0310693, 2003.
92. Z. Toroczkai, G. Korniss, M. A. Novotny, and H. Guclu, “Virtual Time Horizon Control via Communication Network Design”, 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.