

Update statistics in conservative parallel-discrete-event simulations of asynchronous systems

A. Kolakowska* and M. A. Novotny†

*Department of Physics and Astronomy and the ERC Center for Computational Sciences, P.O. Box 5167,
Mississippi State, Mississippi 39762-5167, USA*

Per Arne Rikvold‡

*School of Computational Science and Information Technology, Center for Materials Research and Technology,
and Department of Physics, Florida State University, Tallahassee, Florida 32306-4120, USA*

(Received 7 June 2003; published 14 October 2003)

We model the performance of an ideal closed chain of L processing elements that work in parallel in an asynchronous manner. Their state updates follow a generic conservative algorithm. The conservative update rule determines the growth of a virtual time surface. The physics of this growth is reflected in the utilization (the fraction of working processors) and in the interface width. We show that it is possible to make an explicit connection between the utilization and the microscopic structure of the virtual time interface. We exploit this connection to derive the theoretical probability distribution of updates in the system within an approximate model. It follows that the theoretical lower bound for the computational speedup is $s = (L + 1)/4$ for $L \geq 4$. Our approach uses simple statistics to count distinct surface-configuration classes consistent with the model growth rule. It enables one to compute analytically microscopic properties of an interface, which are unavailable by continuum methods.

DOI: 10.1103/PhysRevE.68.046705

PACS number(s): 05.10.-a, 89.20.-a, 02.50.Fz

I. INTRODUCTION

In discrete-event simulations a physical system with stochastic dynamics is modeled on a lattice of discrete points and changes of its state are viewed as discrete events in time. Physical processes interact with each other at various points in simulation time. The stochastic nature of these interactions makes it difficult to utilize a parallel computing environment to the fullest extent because *a priori* there is no global clock to synchronize physical processes. Examples of such complex systems with underlying asynchronous dynamics come from a wide range of fields, such as activated processes in chemistry, contact processes in epidemiology and ecology models, population dynamics, finance markets, and communication networks and internet traffic, to mention a few. In physics an important example is an interacting spin system, where stochastic processes can be simulated with a dynamic Monte Carlo approach. Until recently, a common belief in the physics community was that even the simplest random-site update Monte Carlo schemes [1] were inherently serial. A popular parallelization technique for these systems is the so-called “trivial parallelization,” in which each processor carries a copy of the full system. An obvious limitation of this technique is imposed by the memory requirement, which may exceed available resources for a large-scale simulation. In nontrivial parallelization, a system is spatially partitioned into subsystems, and each subsystem is placed on a different processor. In other words, in this way *physical processes* and physical interactions between subsystems are mapped to *logical processes* and logical dependences between process-

ing elements. Each logical process manages the state of the assigned physical subsystem and progresses in its own local virtual time (LVT). The asynchronous nature of the physical dynamics implies an asynchronous system of logical processes where discrete events are not synchronized by a global clock. Logical processes execute concurrently and exchange time-stamped messages to perform state updates of the entire physical system being simulated. A sufficient condition for preserving causality in simulations (the so-called local causality constraint) requires that each logical process works out the received messages from other logical processes in nondecreasing time-stamp order [2,3].

Parallel-discrete-event simulations (PDES) are classified in two broad categories: conservative PDES and optimistic PDES. In conservative PDES, originally studied by Chandy and Misra [2,4] and introduced by Lubachevsky in the study of dynamic Ising spin systems [5,6], an algorithm does not allow a logical process to advance its LVT (i.e., to proceed with computations) until it is certain that no causality violation can occur. In the conservative update scenario a logical process may have to be blocked and it may have to wait to ensure that no message with a lower time stamp is received later. Recent physics applications of conservative PDES in modeling magnetization switching [7], ballistic particle deposition [8], and a dynamic phase transition in highly anisotropic thin-film ferromagnets [9,10] suggest that the conservative algorithm should be very efficient in simulating the dynamics of complex systems with short-range interactions. In optimistic PDES [11–15], originated by Jefferson’s time warp algorithm [11], an algorithm allows a logical process to advance its LVT regardless of the possibility of a causality error that may happen in the case of receiving a message with a lower time stamp than the local clock. The optimistic scenario detects causality errors and provides a recovery procedure from the violation of the local causality constraint by

*Electronic address: alicjak@bellsouth.net

†Electronic address: man40@ra.msstate.edu

‡Electronic address: rikvold@csit.fsu.edu

rolling back the events that have been processed prematurely. Although there are no general performance studies to date that would provide an unbiased comparison of the two groups of algorithms, a common perception is that an optimistic PDES should outperform a conservative PDES. However, in the context of physics applications to Ising spin systems, recent numerical studies by Sloot *et al.* [16] demonstrate that near the Ising critical temperature, where long-range correlations occur in the physical spin system being modeled, the computational complexity of an optimistic PDES and the physical complexity of the modeled system are entangled, leading to a nonlinear increase of the roll-back length and a sudden deterioration of the run-time behavior when the number of computing processors is increased.

There are several aspects of PDES algorithms that should be considered in systematic efficiency studies. Some important aspects are the synchronization procedures, the utilization of the parallel environment as measured by the fraction of working processors, memory requirements, interprocessor communications handling, scalability as measured by evaluating the performance when the number of computing processors becomes large, and the speedup as measured by comparing the performance with sequential DES. In routinely performed studies to date, the efficiency is investigated in a heuristic fashion by testing the performance of a selected application in a chosen PDES environment (i.e., in a parallel simulator). Recently, Korniss *et al.* [17] introduced a novel and powerful approach in which a PDES algorithm can be studied in an abstract way by extracting key features of the algorithm, simulating its performance, and applying the methods of nonequilibrium surface growth [18] to evaluate its theoretical efficiency. In the Korniss *et al.* approach, the main concept is the simulated time horizon (STH), defined as the collection of LVTs of all logical processes. The growth rule of this virtual time surface is defined by the communication rule among logical processes (i.e., by their communication topology, which in turn is defined by the underlying dynamics of the physical system being simulated) and by the way in which the algorithm handles the advances in LVTs. In this picture, the utilization of the parallel environment is evaluated as the mean density of local update sites of the growing time interface, and the width of the interface at saturation provides a measure of desynchronization that is directly related to the memory requirements [19]. Scalability properties of a PDES algorithm can be assessed from these performance simulation studies [17,19,20].

In the study of the STH generated by a conservative PDES [17], it has been determined that in the worst-case conservative scenario for a closed spin chain, when each processing element (PE) carries only one spin site (i.e., each logical process simply corresponds to the flipping of one spin) and communicates only with its nearest neighbors, the time evolution of the STH on coarse-grained scales is governed by the Kardar-Parisi-Zhang stochastic equation [21]. This proves, by universality arguments, that the simulation phase of conservative PDES is asymptotically scalable, which guarantees a nonzero utilization even for an infinite number of PEs. Using the same argument, it has been shown that the STH becomes infinitely rough in the limit of an

infinite number of PEs, which suggests possible difficulties with data management. Thus, the measurement phase of conservative PDES is not asymptotically scalable [19]. Recent simulation studies [20] show that conservative PDES can be made fully scalable when the algorithm is supplemented with either a moving time window constraint [22,23] or additional scale-free communication patterns between PEs [24].

From the physics point of view, the virtual time surface of the generic conservative PDES, with its morphology and dynamics, can be viewed as a surface growing through deposition of random time increments in accordance with a growth rule defined by a generic conservative PDES update rule. The physics of this growth is reflected in the utilization (the fraction of nonidling PEs) that corresponds to the mean number of deposition events on the surface. In the case of a closed spin chain this is equivalent to the mean density of local minima in the interface. It should be possible, at least for steady-state simulations, to make an explicit connection between the utilization and the microscopic structure of the interface. Such a connection would enable rigorous studies of the update statistics and a closed theoretical formula for the utilization. The coarse-grained methods previously applied to this problem [17] provide a proof of asymptotic scaling properties in the limit of a large number of PEs. Because of their continuum nature they can neither give a detailed microscopic description of the interface nor is it certain if their results are valid for statistically feasible moderate to large numbers of PEs. On the other hand, the mean utilization strictly depends on the microscopic structure of the STH. In this paper we explore the connection between the STH interface morphology on the microscale and the update statistics by addressing the above questions. Recently, similar connections have been established between the interface microstructure and its mobility for Ising and solid-on-solid models with various dynamics [25–28].

Section II outlines the simulation algorithm for modeling the generic conservative PDES of spatially decomposable cellular automata when each PE carries N lattice sites. The steady-state update statistics for $N=1$ is analyzed in Sec. III. Here we derive formulas for the theoretical utilization and the theoretical probability distribution of updates in the system within an approximate model. Our approach uses simple statistics to build and to count distinct surface-configuration classes consistent with a model update rule (or deposition rule). The idea may be generally applied to any surface that grows on a lattice by a known growth rule. When the growth recipe is known, it is possible to construct diagrams of local lattice-site configurations and to translate the rule to dependences among the graphs. Then the event probability on the surface is deduced from the corresponding diagram of possible surface-configuration classes. The performance of conservative algorithms is discussed in Sec. IV, where the results of Sec. III are applied to estimate the theoretical computational speedup for the ideal system of PEs in a ring communication topology. In Sec. V we discuss generalizations of our approach to other growth processes and advantages that follow in terms of practical applications such as the possibility of computing closed-form expressions for quantities that would be unavailable by standard approaches.

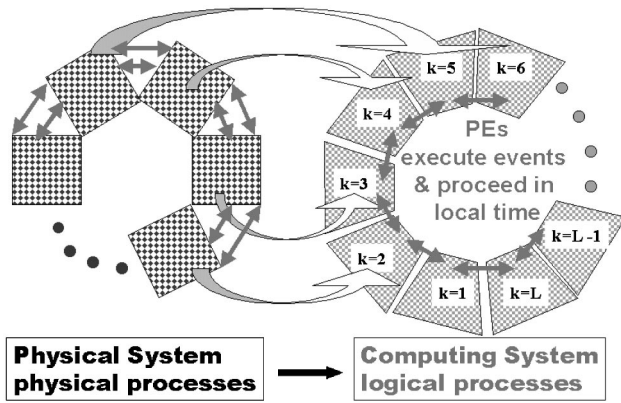


FIG. 1. The mapping of physical processes to logical processes considered in this work. The nearest-neighbor physical interactions (two-sided arrows in the left part) on a lattice with periodic boundary conditions are mapped to the ring communication topology of logical processes (two-sided arrows in the right part). Each PE carries N lattice sites, but communications take place only at border sites. In this study, each PE has at most two effective border sites.

II. MODEL SIMULATIONS OF CONSERVATIVE UPDATE EVENTS

We consider an ideal system of L processors, arranged on a ring (Fig. 1). As an ideal system we understand a system of identical PEs, where communications between PEs take place instantaneously. Each PE carries N lattice sites, N_b of which are border sites and $(N - N_b)$ are interior sites (where all immediate lattice neighbors reside on the same PE). On each PE the simulation algorithm randomly selects one of the N sites. If the selected site is a border site, the PE is required to communicate with its immediate neighbor(s) in an update attempt. If an interior site is selected, the update happens without communication between PEs. For this system, a discrete event means an update attempt. The state of the system does not change between update attempts. Processing elements perform operations concurrently. However, update attempts are not synchronized by a global clock.

An example of the kind of system described above is a large, spatially extended ensemble of spins, arranged on a regular lattice, with a concurrent operation of random Monte Carlo spin-flip attempts. In this picture, the ensemble is spatially decomposed into L subsystems, each of which carries N spin sites. Each subsystem is placed on a PE, and the required communication is the exchange of information about states of the border spins (Fig. 1). In the simplest case of $N=1$, the system is a closed spin chain, and the spin-flip attempt at the k th PE depends on the two nearest-neighbor spins located on the $(k-1)$ th and the $(k+1)$ th PEs. The k th PE is not allowed to update until it receives information from the neighboring PEs. For general N , a sublattice assigned to a PE has N_b border spins. However, for example, in Monte Carlo simulations, at each update attempt only one of the border sites may be randomly selected at a time: either a site from the left border slice or a site from the right border slice. Therefore, considering communications between logical processes, there are only two effective border sites per PE when $N \geq 2$. The case when $N > 1$ and the effective $N_b = 1$ is real-




FIG. 2. The growth and roughening of the STH for $L=100$ and $N=1$: snapshots at t_1 (lower surface) and t_2 (upper surface). Here, $t_1 < t_2 \leq t_x \approx 3700$. Local heights τ_k are in arbitrary units.

ized when on each PE the left and the right border slices coincide. This case is equivalent to a closed spin chain, i.e., to the case of $N=1$.

In generic conservative PDES, to simulate asynchronous dynamics employing L processors, the k th PE generates its own local simulated time τ_k for the next update attempt. The k th local simulated time models the LVT of the k th logical process. Update attempts are simulated as independent Poisson-random processes, in which the k th random time increment η_k (i.e., the random time interval between two successive attempts) is exponentially distributed with unit mean. A processor is allowed to update its local time only if the update is guaranteed not to violate causality. Otherwise, it remains idle. The time step t is the index of the simultaneously performed update attempt. It corresponds to an integer wall-clock time with each PE attempting an update at each value of t . Explicitly, in our model simulations the generic conservative update rule allows the k th PE to update at any time step $(t+1)$ if either of the two conditions is satisfied. First condition: the randomly chosen lattice site is in the interior. Second condition: the randomly chosen lattice site is a border site and either of the following update conditions is satisfied:

$$N=1: \tau_k(t) \leq \min\{\tau_{k-1}(t), \tau_{k+1}(t)\}, \quad (1)$$

$$N \geq 2: \tau_k(t) \leq \tau_r(t), \quad (2)$$

where $r=k-1$ when the left border site is chosen and $r=k+1$ when the right border site is chosen. Following a successful update attempt, the local simulated time is incremented for the next update attempt: $\tau_k(t+1) = \tau_k(t) + \eta_k(t)$. The random time increment $\eta_k(t)$ is computed at each k and t as $\eta_k(t) = -\ln(r)$, where $r \in (0; 1]$ is a uniform deviate. The periodicity condition requires communication between the first and the last PEs in the chain: $\tau_{L+1}(t) = \tau_1(t)$. In simulations we iterate either the update rule (1) or the update rule (2), starting with the initial condition $\tau_k(t=0) = 0$ for all k .

For the set of L processing elements, we define the STH as the set of L local simulated times at time step t . The mean height of the STH is given by the mean virtual time $\langle \tau(t) \rangle_L = 1/L \sum_{k=1}^L \tau_k(t)$. Figure 2 presents the STH gener-

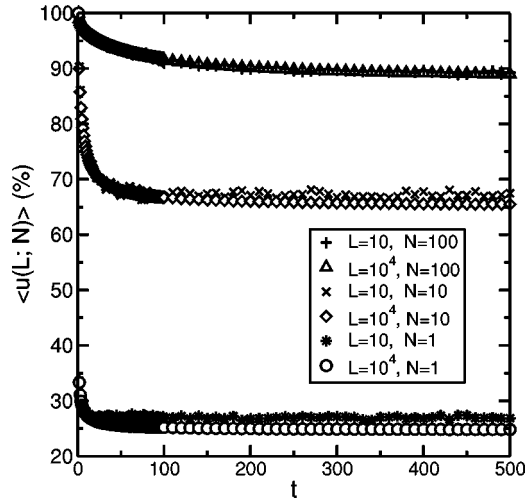


FIG. 3. The time evolution of the utilization $\langle u(t) \rangle$ (averaged over $K=1024$ simulations) for $L=10$ and 10^4 with $N=1, 10, 100$. The result depends most strongly on N .

ated for a closed chain of $L=100$ processors. As the time index advances, the STH grows and roughens. The time evolution of the statistical spread of the interface is characterized by two distinct phases, the growth phase (when $t \ll t_\times$) and the saturation phase (when $t \gg t_\times$), separated by the crossover time t_\times . For a finite L , t_\times marks the transition to the steady state, where the average width of the interface is constant in time and is given by the power law $(NL)^{1/2}$ [17,29].

To study the parallel efficiency, we define the utilization $u(t)$ as the fraction of PEs that perform an update at the parallel time step t . The simulated utilization $\langle u(t) \rangle$ is computed as an ensemble average over many independent simulations. The time evolution of the simulated utilization reaches a steady state $\langle u(t) \rangle = \text{const}$ that depends on the system size (Fig. 3): the steady-state utilization grows monotonically with N . Note, for $N=1$, according to the conservative update rule (1), at t the update at the k th PE site does not happen unless its cumulative local simulated time after $(t-1)$ steps is not larger than the cumulative local simulated times at its neighboring PE sites. This means that an update at the k th PE site corresponds to a local minimum of the STH at the k th site. Accordingly, the mean utilization $\langle u(t) \rangle$ represents the mean number of local minima in the STH interface at t , averaged over many independent simulations. In an individual simulation, the utilization $u(t)$ is the density of the local minima in the STH that is generated in this simulation. When $N \geq 2$ the utilization $u(t)$ is the density of updating sites in the interface. It is important to distinguish between $u(t)$ and $\langle u(t) \rangle$ as $u(t)$ is the characteristic of a particular class of the STH configurations, while $\langle u(t) \rangle$ is the average measurement of $u(t)$ taken over all possible configuration classes. In analyzing the steady-state update statistics the steady-state utilization is denoted by u .

III. STEADY-STATE UPDATE STATISTICS FOR $N=1$

The STH of the generic PDES can be identified with a one-dimensional (1D) interface growing on a ring with the

deposition of random time increments η_k in accordance with the deposition (update) rule given by Eq. (1). The physics of this growth is reflected in the utilization. Because the utilization is strictly related to the microscopic structure of the interface, it is possible to make an explicit connection between the utilization and the morphology of the STH and to derive an analytical formula for the theoretical mean utilization as a function of the system size L . In this section we make this connection for $N=1$ when the STH growth has reached the saturation phase (i.e., when $t \gg t_\times$). Our derivation of the update distribution makes the following two simplifying assumptions. First, we neglect correlations between nearest-neighbor local slopes. These depend on the type of deposition, i.e., our derivation is not specific to the distribution from which the deposited η_k are sampled. This simplification is reflected in the assumption of equal statistical weights assigned to the legs of binary transition graphs that represent possible choices of neighboring local sites. Second, we neglect temporal correlations among the groups of the surface-configuration classes. Because of the above two simplifications, our theoretical result for the mean utilization is a mean-field-like approximation to the mean utilization measured in simulations.

A. Theoretical utilization

There are only four groups of elementary local site configurations of the STH that correspond to four mutually exclusive discrete events that take place at the k th PE site at t . These are as follows: “A” denotes an event when the update rule (1) is satisfied from the left and from the right, i.e., when $\tau_{k-1} \geq \tau_k$ and $\tau_k \leq \tau_{k+1}$; “B” denotes an event when the update rule (1) is not satisfied from the right, i.e., when $\tau_{k-1} \geq \tau_k$ and $\tau_k > \tau_{k+1}$; “C” denotes an event when the update rule (1) is not satisfied from the left, i.e., when $\tau_{k-1} < \tau_k$ and $\tau_k \leq \tau_{k+1}$; and “D” denotes an event when the update rule (1) is not satisfied from either side, i.e., when $\tau_{k-1} < \tau_k$ and $\tau_k > \tau_{k+1}$. The corresponding elementary local configurations of the STH at the k th PE site are denoted by A, B, C, and D (Fig. 4). Because of the periodicity condition (i.e., $\tau_{L+1} = \tau_1$), during the steady state not all L sites can have the same elementary site configuration [35]. Therefore, in the set of L sites there must be at least one site with configuration A. Without losing generality, we assign the index $k=1$ to one of the sites that are in the local configuration A and enumerate the other sites accordingly, progressing to the right. Its right neighbor (having index $k=2$) can be only either in configuration C or in configuration D. Similarly, its left neighbor (having index $k=L$) can only be in either B or D. If site $k=2$ is in configuration C, then site $k=3$ can be only either in configuration C or D. If site $k=2$ is in D, then site $k=3$ can be only either in B or A. These choices are presented as transition graphs (binary trees) in Fig. 5. We adopt an approximation in which, during the steady state, the possible choices of transitions from the k th site to the right neighboring $(k+1)$ site are realized on average with equal frequency. Consequently, we assign equal statistical weights to each leg of the transition graph in Fig. 5. Starting from the site $k=1$ and progressing to the right towards $k=L$, with the




FIG. 4. The four groups of elementary local surface configurations of the STH at the k th site. The index k denotes the k th PE in the chain ($N=1$). Each group corresponds to one of the four mutually exclusive discrete events A, B, C, and D at an update attempt. A denotes an event when the update rule is satisfied. B denotes an event when the update rule is not satisfied from the right. C denotes an event when the update rule is not satisfied from the left. D denotes an event when the update rule is not satisfied from the left and the right.

help of elementary transition graphs we can construct all possible configuration equivalency classes of the entire surface generated by the deposition (update) rule (1). These can be categorized into groups (called p groups) based on the number p of the deposition (update) events at t , i.e., the number of local minima in the surface configuration (coded by A) at t . The utilization of the p group is $u(p) = p/L$. The probability distribution $f(p;L)$ of the deposition (update) events is obtained as the quotient of the multiplicity $M(p)$ of the p -group configuration class and the total number M of configuration classes [36].

For example, the binary tree for the construction of possible surface-configuration classes for $L=5$ is shown in Fig. 6. Looking along its branches, starting from the leading A at the fixed $k=1$ position, it is easy to identify a total of eight possible configuration classes of the entire surface: (1) ACCCD; (2) ACCDB; (3) ACDBB; (4) ACDAD; (5) ADBBB; (6) ADBAD; (7) ADACD; and (8) ADADB. Note, according to the surface construction rule, the class representative (4) is not equivalent to the class representative (7). This is because the leading A in configuration (7) has a local maximum as its right neighbor and configuration (4) does not have this property. If the assignment of an index to a site were irrelevant, all configurations that can be obtained under




FIG. 5. Binary branching of possible choices in constructing a surface configuration from the elementary local configurations A, B, C, and D of Fig. 4. (a) The alternatives that must be followed starting with A at $k=1$ and progressing towards $k=L$ to the right. (b) The only possible alternative for a periodic chain closed at $k=L$: the left neighbor of site $k=1$ must have configuration either B or D.




FIG. 6. Binary tree for the construction of all possible configurations of the surface for $L=5$. A number to the left of the configuration symbol denotes the level of branching. A number in parenthesis to the right of the configuration symbol denotes the number of branching levels in a subtree. Notice the recurrent structure: the graph consists of the nested trees A(4), D(3), A(2), and D(1). The dashed lines mark the transition cuts to the lower level trees. A(1) and D(1) denote the one-level branches A and D, respectively, that mark the end of branching. See discussion in Sec. III A and Appendix A.

an even permutation of sites would have fallen into one equivalency class. The surfaces representing configurations (1)–(8) are sketched in Fig. 7. Each surface configuration represents a class of infinitely many topologically equivalent deformations because the deposited random time increment is a real positive number that can take on continuous values in the interval $[0; \infty)$. There are only two p groups. In the first group there are four classes with one letter A: $M(1) = 4$, $f(1;5) = 1/2$, and $u(1) = 1/5$. In the second group there are four classes with two letters A: $M(2) = 4$, $f(2;5) = 1/2$, and $u(2) = 2/5$. Thus, for $L=5$ the mean utilization that is measured during steady-state simulations is $\langle u(L=5;N=1) \rangle = f(1;5)u(1) + f(2;5)u(2) = 3/10$.




FIG. 7. The graphs of possible surface-configuration classes that correspond to the configurations read along the branches from Fig. 6: (1) ACCCD; (2) ACCDB; (3) ACDBB; (4) ACDAD; (5) ADBBB; (6) ADBAD; (7) ADACD; and (8) ADADB. Each graph represents a class of infinitely many topologically equivalent deformations.




FIG. 8. The recurrent structure of the binary tree in constructing the classes of surface configurations for general L . The meaning of the symbols is the same as in Fig. 6. For general L , the highest level tree is $1A(L-1)$ that has 2^{L-2} branches. Each branch represents a class of surface configurations. The branches are categorized in distinct groups. Each group contains configurations with exactly p repetitions of $A(1)$. The smallest p is 1, the largest p is $\lfloor L/2 \rfloor$. The utilization in each group is p/L .

For general L , the utilization measured in simulations during the steady state is the mean frequency of the local surface minima, averaged over all admissible surface configurations. It can be obtained from the generally valid formula for the computation of averages:

$$\langle u(L;N) \rangle = \sum_p f(p;L)u(p), \quad (3)$$

where the summation extends over all p groups of the admissible surface-configuration classes, $u(p)$ is the utilization characteristic for each group, and $f(p;L)$ is the frequency of the occurrence of p group during the steady state. To find the theoretical $f(p;L)$, one can exploit the recurrent structure of the corresponding binary tree (Fig. 8) in counting the classes of the surface configurations (branches) that contain the elementary site configuration A at exactly p number of sites, $p = 1, 2, 3, \dots, p_{max} = \lfloor L/2 \rfloor$ ($\lfloor L/2 \rfloor$ denotes the integral part, which is $L/2$ for even L and $(L-1)/2$ for odd L). The details of the derivation are given in Appendix A. The total number of configuration classes is $M = 2^{L-2}$. The number of branches with exactly p occurrences of A is $M(p) = (L-1)! / [(2p-1)!(L-2p)!]$. The frequency of occurrence of the p group is $f(p;L) = M(p)/M$. Thus, the theoretical mean utilization of the steady state is

$$\langle u(L;1) \rangle = \frac{1}{2^{L-2}} \sum_{p=1}^{\lfloor L/2 \rfloor} \binom{L-1}{2p-1} \frac{p}{L} = \begin{cases} 1/2, & L=2 \\ (L+1)/4L, & L \geq 3. \end{cases} \quad (4)$$

The theoretical utilization $\langle u(L;1) \rangle$ is bounded from below by $\langle u(L \rightarrow \infty;1) \rangle = 1/4$ (Fig. 9).

In classifying individual configurations, the underlying principle is provided by the deposition rule given by Eq. (1). Therefore, the local A configuration represents four types of




FIG. 9. The steady-state mean utilization as a function of the system size for $N=1$. The continuous curve represents the analytical result [Eq. (4)]. It converges to $\lim_{L \rightarrow \infty} \langle u(L;1) \rangle = 1/4$ (horizontal line). The circles represent the utilization measured in simulations, with error bars smaller than the symbol size.

update events and the local configuration B (or C) represents two types of no-update events (Fig. 4). The small differences between the simulation results and Eq. (4), clearly observed in Fig. 9, come mainly from neglecting temporal correlations among p groups of surface-configuration classes in our derivation. These correlations are intrinsically present in the computation of averages over time series in simulations but are absent in our model. They depend on the type of deposition, i.e., the probability distribution from which the random time increments η_k are sampled. A possible second source of discrepancies is the assumption of equal statistical weights in the transition graphs (Fig. 5). When the actual weights are only approximately equal, this modifies the frequency $f(p;L)$ of the occurrence of a p group in Eq. (3), so a particular surface-configuration class may occur slightly more (or less) often in simulations than would result from our assumption. Note that this modifies only $f(p;L)$; the utilization $u(p)$ of a p group is not changed. In deriving $f(p;L)$ the underlying assumption implies that any class of the entire surface configurations is equally probable. The factor $1/M = 1/2^{L-2}$ in Eq. (4) has the meaning of this probability (Appendix A).

B. Computation of averages

In simulations, the average steady-state utilization is measured at each t as the arithmetic average over an ensemble of K independent simulations and then averaged over a series of T time steps during the steady state. At each t , this is equivalent to the computation of averages over the surface-configuration classes in accordance with Eq. (3), where $f(p;L)$ is estimated from the steady-state simulation data. Denoting by $G(p;L)$ such an “experimental” frequency, we write explicitly

$$\langle u(t) \rangle_K \approx \frac{1}{K} \sum_{i=1}^K u(i,t) = \sum_{p=1}^{\lfloor L/2 \rfloor} G(p,t)u(p), \quad (5)$$




FIG. 10. The probability distribution for $L=50$: the theoretical $f(p;L)$ (histogram) and $\langle G(p;L) \rangle_T$ measured in simulations (symbols). The error bars represent one standard deviation from the mean of the measured time sequence at saturation [the quantity $\delta G(p)$ that enters Eq. (7)]. The measured frequencies were obtained from an ensemble of $K=2048$ independent simulations as $K(p)/K$, where $K(p)$ is the number of trials that produced the p group of the surface-configuration classes.

where the right-hand side (rhs) follows simply from grouping the summation terms. This is possible because $u(i,t)$ takes on only the values $u(p)$ that characterize the p group of the surface-configuration classes. Having a sequence of measured frequencies $G(p;L)$ over the steady-state time interval, the time average $\langle G(p;L) \rangle_T$ can be computed for each p . After time averaging, Eq. (5) gives

$$\langle u \rangle_{K,T} \approx \sum_{p=1}^{[L/2]} \langle G(p) \rangle_T u(p). \quad (6)$$





FIG. 11. The time sequence of frequencies of the surface configurations characterized by the utilization $u(p)=p/L$. The continuous horizontal lines represent the theoretical $f(p;L)$. Symbols are simulation data $G(p;L)$. The dashed horizontal lines represent time averages $\langle G(p;L) \rangle_T$ over an interval of 1000 t steps, beginning at $t=10^3$. The error bars represent one standard deviation from $\langle G(p;L) \rangle_T$ as in Fig. 10. The data were taken in $K=1024$ simulations: (a) For $L=4$, $\langle G(1;4) \rangle_T=0.7323 \pm 0.0138$ and $\langle G(2;4) \rangle_T=0.2677 \pm 0.0138$ and (b) for $L=11$, $\langle G(1;11) \rangle_T=0.0268 \pm 0.0051$, $\langle G(2;11) \rangle_T=0.2583 \pm 0.0144$, $\langle G(3;11) \rangle_T=0.4759 \pm 0.0157$, $\langle G(4;11) \rangle_T=0.2194 \pm 0.0134$, and $\langle G(5;11) \rangle_T=0.0194 \pm 0.0044$.

The corresponding statistical spread of the measured average utilization $\delta \langle u \rangle$, i.e., the standard deviation of the mean $\langle u(L) \rangle$, can then be determined from the measured standard deviations of $G(p;L)$:

$$\delta \langle u(L) \rangle \approx \sqrt{\sum_p [u(p) \delta G(p)]^2}, \quad (7)$$

where $\delta G(p)$ denotes the empirical standard deviation of the $G(p;L)$ time sequence. At each t , the frequencies $G(p;L)$ are found by directly counting the simulations that produced $u(p)=p/L$ and, subsequently, computing the quotient of this count $K(p)$ and the total number K of simulations in an ensemble. Explicitly, for $p=1, 2, \dots, [L/2]$, the measured frequency is $G(p;L)=K(p)/K$, where $K=\sum_p K(p)$ (Fig. 10).

A typical time sequence of $G(p;L)$, measured in $K=1024$ independent simulations, is shown in Fig. 11. For $L=4$, the theoretical steady-state frequencies, $f(1;4)=3/4$ and $f(2;4)=1/4$, differ slightly from the averages $\langle G(p;L) \rangle_T \pm \delta G(p)$ computed over an interval of $T=1000$ steps, beginning at $t=10^3$. The measured steady-state utilization $\langle u(L;1) \rangle \pm \delta \langle u(L) \rangle$ is $\langle u(4;1) \rangle=0.3169 \pm 0.0077$. Similarly, for $L=11$ the measured frequencies are in close agreement with the theory: $f(1;11)=f(5;11)=5/256$, $f(2;11)=f(4;11)=15/64$, and $f(3;11)=63/128$. The measured steady-state utilization is $\langle u(11;1) \rangle=0.2678 \pm 0.0073$. The theoretical utilizations $\langle u(4;1) \rangle=5/16$ and $\langle u(11;1) \rangle=3/11$ [from Eq. (4)] compare with the utilizations measured in simulations well within the statistical error bars when $K=1024, 2048$; likewise, there is very good agreement for general L . However, when $K=4096$ the statistical spread $\delta \langle u(L) \rangle$ is small enough to see that the results of Eq. (4) lie above the simulation data in Fig. 9.




FIG. 12. The probability distribution $f(p;L)$ of p updates in a closed linear chain of L PEs, each carrying one lattice site and following the conservative update rule. $L=250, 500, 1000,$ and 1500 (from left to right).

The standard deviation of the distribution of $u(p)$ among admissible p groups of the surface-configuration classes can be measured directly in simulations as the square root of the variance $\text{var}(u)$:

$$\begin{aligned} \text{var}(u) &\approx \langle u^2 \rangle_{K,T} - \langle u \rangle_{K,T}^2 \\ &= \sum_{p=1}^{[L/2]} \langle G(p;L) \rangle_T u(p)^2 - \left(\sum_{p=1}^{[L/2]} \langle G(p;L) \rangle_T u(p) \right)^2, \end{aligned} \quad (8)$$

where the rhs comes from grouping terms in the summations. The statistical spread in $u(p)$ is the largest for $L=4$ and decreases when L increases. Equation (8) gives the measured average variance of the probability distribution of updates in the system, scaled by L^2 .

C. Probability distribution of updates

The derived theoretical probability distribution of updates in a closed linear chain of L processors, each carrying $N=1$ lattice sites and following the conservative update rule, is

$$f(p;L) = \frac{1}{2^{L-2}} \binom{L-1}{2p-1}, \quad (9)$$

where p is the number of updates at the t th update attempt in the steady-state simulation. In an equivalent interpretation, Eq. (9) gives the probability of generating a surface with p local minima at saturation when the surface growth obeys rule (1). In other words, in the latter interpretation $f(p;L)$ is the probability distribution of the deposition events on the surface. Equation (9), derived in Appendix A, was already used in calculating $\langle u(L;1) \rangle$ in Eq. (4). Figure 12 presents $f(p;L)$ for various system sizes. It can be seen that the measured utilization is $\langle u \rangle = \langle p \rangle / L$, where $\langle p \rangle$ is the mean of

$f(p;L)$. The variance σ^2 of $f(p;L)$ (and, thus the standard deviation of u) can be obtained in the usual way (Appendix B) as

$$\sigma^2 = \sum_m (m - \langle p \rangle)^2 f(m;L) = \frac{L-1}{16} \quad (10)$$

for $L \geq 4$ and $\sigma^2 = 0$ for $L=2,3$, where

$$\langle p^2 \rangle = \sum_m m^2 f(m;L) = \frac{L(L+3)}{16} \quad (11)$$

for $L \geq 4$ and $\langle p^2 \rangle = 1$ for $L=2,3$, and

$$\langle p \rangle = \sum_m m f(m;L) = \frac{L+1}{4} \quad (12)$$

for $L \geq 3$ and $\langle p \rangle = 1$ for $L=2$. Also, using Eq. (12), it can be derived that $\langle 1/p \rangle = 1$ for $L=2$, and for $L \geq 3$,

$$\begin{aligned} \left\langle \frac{1}{p} \right\rangle &= \sum_m \frac{1}{m} f(m;L) \\ &= \frac{8}{L(L+1)} \left(\sum_k (2k-1) f(k;L+2) - \frac{L+1}{2^L} \right) \\ &= \frac{4}{L} \left(1 - \frac{1}{2^{L-1}} \right) \end{aligned} \quad (13)$$

and for $L \geq 4$,

$$\langle p^3 \rangle = \frac{L^3 + 6L^2 + 3L - 2}{64}, \quad (14)$$

$$\langle p^4 \rangle = \frac{L(L^3 + 10L^2 + 15L - 10)}{256}. \quad (15)$$

The form of Eq. (9) permits exact computations of all moments of $f(p;L)$. The corresponding recursion formula is given in Appendix B. In this way we find that the skewness of $f(p;L)$ is zero, which means that $f(p;L)$ is symmetric about $\langle p \rangle$. The computed kurtosis (the fourth central moment) is strictly positive for $L \geq 4$, which means that $f(p;L)$ is more pointed than a Gaussian.

The theoretical standard deviation of the utilization distribution among various p groups, i.e., the statistical spread of $u(p)$ in an ensemble of independent simulations, can be computed from Eq. (10):

$$\sigma(u) = \frac{\sqrt{L-1}}{4L}. \quad (16)$$

In an equivalent interpretation, Eq. (16) gives the distribution of $u(p)$ among various p groups of surface-configuration classes.

The statistical distribution of the updates in the system of L processors can also be estimated directly from the simulation data, without any presupposed underlying model. It is sufficient to notice that the update at a PE site happens when

the STH has a local minimum at that site and that the 1D surface of L sites may have no more than $[L/2]$ local minima. Denoting by $K(p)$ the number of simulations that produced surfaces with exactly p local minima in the sequence of K independent simulations [p and $K(p)$ are directly counted in simulations], at any t the experimental distribution is $G(p;L) = K(p)/K$. Its time average in the steady state is $\langle G(p;L) \rangle_T$. The variance of $\langle G(p;L) \rangle_T$ can be obtained directly from the simulation data as described in Sec. III B [by the rhs of Eq. (8) multiplied by L^2]. In the infinite K limit $\langle G(p;L) \rangle_T$ converges to the exact steady-state distribution $g(p;L)$. At any t the average of any observable Q that depends on the number of local minima can be evaluated as

$$\begin{aligned} \langle Q \rangle_K &\approx \frac{1}{K} \sum_{r=1}^K Q(r) \\ &= \frac{1}{K} \sum_{p=1}^{[L/2]} K(p) \left(\frac{1}{K(p)} \sum_{r=1}^{K(p)} Q(r) \right) \\ &= \sum_{p=1}^{[L/2]} G(p;L) \langle Q \rangle_{K(p)}, \end{aligned} \quad (17)$$

where $\langle \cdot \rangle_{K(p)}$ is the mean over the measured configuration classes in the p group. The exact mean in the steady state is

$$\langle Q \rangle = \sum_{p=1}^{[L/2]} g(p;L) Q(p,L), \quad (18)$$

where $Q(p,L)$ is the value typical for the p group.

The question that naturally arises at this point is how close the theoretical distribution $f(p;L)$ given by Eq. (9) represents the exact distribution $g(p;L)$ that enters Eq. (18). The results for $L=50$ obtained with $K=2048$, presented as a histogram in Fig. 10, show that $f(p;L)$ mimics the overall shape of the experimental $\langle G(p;L) \rangle_T$ very well. Increasing K to 4096 improves the error bars of $\langle G(p;L) \rangle_T$ (in particular, for the extreme values of p); yet, as K gets larger the overall shape of $\langle G(p;L) \rangle_T$ remains unchanged and the difference $\langle G(p;L) \rangle_T - f(p;L)$ does not entirely vanish. This difference is most pronounced when L is small and is largest for $L=4$. Since for $L=4$ there are only two values of p , it is easy to estimate the exact $g(p;L)$ to a high degree of confidence. For $K=1024$, $\langle G(1;4) \rangle_T = 0.732\,26 \pm 0.013\,78$ and $\langle G(2;4) \rangle_T = 0.267\,73 \pm 0.013\,78$. For $K=5120$, $\langle G(1;4) \rangle_T = 0.731\,86 \pm 0.006\,11$ and $\langle G(2;4) \rangle_T = 0.268\,13 \pm 0.006\,11$. As K increases, $\langle G(p;L) \rangle_T$ converges to $g(1;4) \approx 0.732$ and $g(2;4) \approx 0.268$, while the theoretical estimate is $f(1;4) = 0.75$ and $f(2;4) = 0.25$. The small difference $g(p;L) - f(p;L)$ for the worst case of $L=4$ indicates that Eq. (9) is a close approximation to the time-averaged exact distribution $g(p;L)$. We believe the primary reason for this difference is the lack of temporal correlations among various p groups in our computational model, as was pointed out in Sec. III A. Regardless of this simplification, the theoretical standard deviations $\sigma(u)$ given by Eq. (16) agree with $\sigma(u)$ computed directly in simulations via Eq. (8).

IV. PERFORMANCE OF A CONSERVATIVE ALGORITHM

The computational speedup s of a parallel algorithm is defined as the ratio of the time required to perform a computation in serial processing on one PE to the time the same computation takes in parallel processing on L processors. It is easy to derive from the above definition that for an ideal system of processors the computational speedup is the product of the number L of PEs in the system and the utilization $\langle u(L;N) \rangle$ of the parallel processing environment: $s = L \langle u(L;N) \rangle$. In other words, the speedup is measured by the average number of PEs that work concurrently between two successive update attempts.

We observe that for ideal PEs the speedup as a function $F(L)$ must be such that the equation $F(L) = s$ has a unique solution, where s is a fixed positive number. This requirement follows naturally from the logical argument that distributing the computations over L ideal processors gives a unique speedup, i.e., two ideal systems having sizes L_1 and L_2 , respectively, may not give the same s . This means that $F(L)$ must be a monotonically increasing function of L .

In our model for $N=1$ the average number of PEs that work in parallel, i.e., the speedup, is the mean number of local minima in the STH during steady-state simulations. In Sec. III C this number is computed explicitly as the first moment of a theoretical distribution given by Eq. (9). In this way, translating Eq. (12) to the language of applications, the theoretical speedup that can be obtained in an ideal system of PEs, performing conservative PDES in the ring communication topology with $N=1$, is given by

$$s = \begin{cases} 1, & L=2,3 \\ (L+1)/4, & L \geq 4. \end{cases} \quad (19)$$

In what follows we discuss the consequences of Eq. (19) [or its alternative, Eq. (4)] from the point of view of applications.

One of the consequences is that the theoretical upper bound for $\langle u(L;1) \rangle$ is $1/2$. This corresponds to the situation when only one of the two PEs is working at a time while the other one is idle. In the picture when the simulations represent operations performed by a parallel algorithm, when $L=2$ or $L=3$ the parallelization within the conservative update scheme does not give an advantage in terms of the computation time because the processors work alternately. For an ideal system of PEs, where communications between PEs take place instantaneously, such an operation will not produce speedup. For a real system of PEs, the communication overhead will produce an actual slowdown, i.e., the parallel execution time will be longer than the sequential execution time on one PE. Between the update attempts during the steady state the average number of PEs working in parallel is $L \langle u(L;1) \rangle = (L+1)/4$. This means, when $L=4$ or 5 the actual number of working PEs is still either 1 or 2; and when $L=6$ or 7 the actual number of working PEs is 1, 2, or 3 at a time and on average this actual number is still either 1 or 2. This will produce a small speedup for an ideal system of PEs, but for a real system of PEs this speedup may be negligible or not present at all. A noticeable advantage in terms

of speedup should be expected when the average number of working PEs is $(L+1)/4 > 2$, which gives $L \geq 8$. For a real system of PEs the best speedup will not be larger than the average speedup for an ideal system, i.e., not larger than s given by Eq. (19).

The parallel utilization efficiency $\langle u(L;N) \rangle$ and the speedup, as measured by the average number $L\langle u(L;N) \rangle$ of PEs working in parallel, depend on the number N of lattice sites per PE, as well as on the number N_b of border lattice sites per PE, and on the communication topology among the PEs. Our earlier large-scale simulations [20] show that the worst-case scenario of $N=1$, studied in this work, can be greatly improved when N is increased while retaining the ring communication topology with the effective $N_b=2$ (Fig. 3). On the other hand, the case of $N_b=2$ seems to have limited applications and should rather be considered as an intermediate model towards the study of more realistic conservative PDES where the effective N_b may be arbitrary. From this perspective the case of $N_b=2$ is really the best-case scenario since the utilization declines with increased number of communications between PEs, i.e., when the effective N_b increases. In the ring communication topology, it can be expected that the actual speedup in the ideal system of PEs should be larger than it is in the worst-case scenario with $N=1$, but smaller than it is in the best-case scenario with large N and $N_b=2$. The actual speedup in the real system of PEs should not be larger than a theoretical upper bound for the ideal system. Deriving a theoretical estimate for this upper bound requires first finding a closed expression for $\langle u(L;2) \rangle$, a problem that is still awaiting solution.

V. DISCUSSION

We showed in Sec. III that for steady-state simulations for a closed chain it is possible to explicitly correlate deposition statistics with the surface morphology on a microscopic scale. In our approach we used simple laws of statistics to build distinct configuration classes of the virtual time horizon. For one particular rule of surface growth, we constructed binary trees from which we could read the surface equivalency classes, serving our purpose of counting a particular type of configurations, relevant in deriving an approximate analytical expression for the measured utilization, i.e., the measured frequency of the time deposition. In the case solved in this work, the summation process was technically easy once the symmetries in the binary graphs became apparent (Fig. 8). In principle, our method may be applied to any surface that grows on a lattice by a known growth rule, and it can be generalized to any measurable quantity Q . When the growth recipe is available it should be possible to construct diagrams of elementary site configurations and to translate the growth rule to dependences among the graphs. An observable Q can then be computed as the average over all available groups of classes of the surface configurations: $Q = \sum_p f(p)q(p)$, where $f(p)$ is the probability distribution and $q(p)$ is the value of Q characteristic for each group. In the example given in this work, the statistical weights assigned to each leg of an elementary transition diagram (Fig. 5) were taken to be equal. However, a different deposition

rule may require a different choice of elementary diagrams and a different choice of statistical weights. The present application of the method to the 1D deposition problem on the STH surface is a promising example that could be generalized to a variety of other growth processes.

In a common approach one finds the universal properties of growing interfaces from a stochastic growth equation that is solved in a coarse-grained approximation at large scales. One powerful technique is the renormalization group approach [21,30–32]. The coarse-grained solution of the stochastic dynamics provides asymptotic scaling properties in the limit of large system sizes. Because of its continuum nature this method is not capable of giving a detailed microscopic description of interfaces such as the probability distribution of events on the growing surface. When applied to the model of conservative PDES studied in this paper, the continuum method neither gives any estimate of the utilization of the parallel environment and the speedup nor does it give the scaling behavior for the utilization in the limit of infinite system size. In earlier work, Korniss *et al.* [17] used a coarse-grained method to determine that, in the steady state for $N=1$, the conservative STH is governed by the Edwards-Wilkinson Hamiltonian [33], which implies a nonzero utilization in the infinite L limit, i.e., the asymptotic scalability of a generic conservative PDES. This finding explained the observed tendencies in the time evolution of the large-scale simulation data for the utilization, which clearly showed that the steady-state mean utilization settles down at a nonzero value, slightly lower than 1/4. The simplified analysis of the microscopic structure of the conservative STH at saturation, presented in this work, enabled us to derive the analytical formula for the utilization, Eq. (4), in the approximation where correlations among various surface-configuration classes are absent during the steady state. Equation (4) provides the explicit scaling relation for the utilization, which shows directly the asymptotic scalability of conservative PDES. The limiting value $\lim_{L \rightarrow \infty} \langle u(L;1) \rangle = 1/4$ coincides with the estimate in Ref. [17]. The actual simulated values for the mean utilization $\langle u(L;1) \rangle$ fall below the analytical curve (Fig. 9). This small, but statistically significant, difference is evidence for small spatial and temporal correlations inherently present in the simulation data during the steady state. Theoretical treatment of these correlations, which would provide a correction to the theoretical distribution derived in this work, remains to be explored.

The closed form of the event distribution during steady-state simulations [Eq. (9)] enables one to compute analytically the mean of any observable that depends on the number of local minima in the STH surface [Eq. (18)]. In this way we derived the explicit expressions for $\langle p^n \rangle$ [Eqs. (10)–(15)] that are valid for all values of L within the adopted model. When $N=1$, the measured mean number of local minima $\langle p \rangle$ directly translates to the utilization and to the speedup of conservative PDES. This approach has an advantage over any of the common continuum methods because it gives not only the exact scaling relations in the limit of large system sizes but also enables one to compute analytic quantities valid for any system size. While the asymptotic properties of

a PDES (such as the scalability as L increases) are of theoretical interest, the explicit formulas for the performance evaluation of a PDES for finite L and N are of practical value in algorithm design.

When the number of lattice sites per PE is $N \geq 3$, each PE carries two kinds of sites, interior sites and border sites. This generates two groups of mutually exclusive update events: updates when an interior site is randomly chosen and updates when a border site is randomly chosen. These groups of events are mutually exclusive because during the t th update attempt the drawing of a lattice site is performed only once per PE. Finding the mean fraction of PEs that made an update while a border site was selected requires solving the case of $N=2$. This case is different from the case of $N=1$ because the update rule changes. Now, at each t we first randomly select a border site on the k th PE, and when the neighboring PE of the selected site has its local simulated time larger than τ_k the k th PE makes the update. This “one-sided” rule changes the deposition pattern (the STH growth rule) because now the k th PE site does not need to be in elementary configuration A to make an update. A new rule implies a new definition of p -group configuration classes, which is now characterized by p updates (for $N=1$ it is defined by p local minima). Also, the way in which new frequencies $f(p;L)$ are computed must incorporate a random choice of border sites. We leave the questions of deriving update distributions for $N \geq 2$ open for future investigations.

VI. SUMMARY

We simulated the performance of an ideal closed chain of PEs that work in parallel in an asynchronous manner, with updates following a generic conservative algorithm. The conservative update rule can be seen as the mechanism that determines the growth of the virtual time surface of a conservative process. The physics of this growth is reflected in the utilization and in the interface width.

We showed that it is possible to make an explicit connection between the steady-state utilization and the microscopic structure of the virtual time interface at saturation. We exploited this connection to derive an analytical formula for the probability distribution of the update events in the system within an approximate model. Then, having the model probability distribution, we computed explicit expressions for the mean utilization and the computational speedup as functions of the system size. Our result states that the speedup for the ideal closed chain of PEs, each carrying one lattice site, grows linearly with the system size as $s=(L+1)/4$ for $L \geq 4$, and $s=1$ for $L=2,3$.

Finally, we observe that our approach could be applied to a variety of other growth processes. In this sense, the present 1D application to the update problem in conservative PDES is a promising example. The main advantage of the approach is that it enables one to compute analytically quantities that otherwise can be only estimated qualitatively in an asymptotic fashion by continuum methods.

ACKNOWLEDGMENTS

The authors thank G. Korniss for discussions. P.A.R. appreciates the hospitality of the MSU Department of Physics

and Astronomy and the ERC Center for Computational Sciences. This work was supported by NSF Grant Nos. DMR-0113049 and DMR-0120310, and by the Department of Physics and Astronomy at MSU and the ERC Center for Computational Sciences at MSU. This research used resources of the National Energy Research Scientific Computing Center, which was supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC03-76SF00098.

APPENDIX A: DERIVATION OF $f(p;L)$

To find the theoretical frequency $f(p;L)=M(p)/M$, one should compute the number $M(p)$ of surface-configuration classes that contain the elementary configuration A (a local minimum of the STH) at exactly p sites. Both $M(p)$ and M (the total number of configuration classes) can be easily found by simply counting the branches of the binary trees presented in Fig. 8.

For a given L , we start with the highest level tree $1A(L-1)$ (the left tree in Fig. 8). This tree has $(L-1)$ branching levels. The $1A(L-1)$ tree branches to $2C(L-2)$ and $2D(L-2)$ (the right tree in Fig. 8), which have $(L-2)$ branching levels, etc. The branching ends up at $(L-1)A(1)$, $(L-1)B(1)$, $(L-1)C(1)$, and $(L-1)D(1)$ that do not branch. Therefore the total number of branches is $M=(1/2)\prod_{n=1}^{L-1}2=2^{L-2}$.

The factor $1/M$ is the probability that any class of the entire surface configuration appears in an individual simulation at t . Assuming that each leg of the binary branch in Fig. 8 carries the statistical weight of $1/2$, this probability is $1/2^n$, where $n=L-2$ is the highest binary branching level.

To find $M(p)$, notice (Figs. 8 and 6) that in each A tree there is at least one sub-branch that contains no A (along the C branch), and in each D tree there is exactly one sub-branch that contains no A (along the B branch). Let $n(p;kX(r))$ denote the number of branches that contain exactly p number of A's in the sub-branch $kX(r)$. Here, “X” stands either for A or D. In this notation the “exactly one branch with no A in the D tree” means

$$n(0;kD(r))=1. \quad (\text{A1})$$

The k th level A tree has $n(1;kA(r))$ branches that contain exactly one A:

$$n(1;kA(r))=1+\sum_{s=1}^{r-1}n(0;(k+r-s)D(s)). \quad (\text{A2})$$

The meaning of Eq. (A2) can be clarified by the example of the tree in Fig. 6. The first branch cut (the left dashed line) separates the C sub-branch from the remainder. The C sub-branch (to the left of the first cut) has only one (leading) A. Therefore, there is 1 on the rhs of Eq. (A2). The remainder (to the right of the first cut) is the sum of the three D trees: $4D(1)$, $3D(2)$, and $2D(3)$. Therefore, the sum in Eq. (A2) has three terms. Since the remainder has already one A in the

leading position, the sum must contain only the subtrees that contain $p=0$ number of A's. In a similar fashion we obtain for the D tree:

$$n(1;kD(r)) = \sum_{s=1}^{r-1} n(1;(k+r-s)A(s)), \quad (\text{A3})$$

where the summation extends over all A trees because the left sub-branch of the D tree contains no A's (Figs. 6 and 8).

To find $n(2;kA(r))$, one must first sum over all D trees (to the right of the first cut in Fig. 6 or along the left graph in Fig. 8) and then sum over all A trees (to the right of the second cut in Fig. 6 or along the right graph in Fig. 8):

$$\begin{aligned} n(2;kA(r)) &= \sum_{s=1}^{r-1} n(1;(k+r-s)D(s)) \\ &= \sum_{s=1}^{r-1} \sum_{l=1}^{s-1} n(1;(k+r-l)A(l)) \\ &= \sum_{s=1}^{r-1} \sum_{l=1}^{s-1} \left(1 + \sum_{m=1}^{l-1} n(0;(k+r-m)D(m)) \right). \end{aligned} \quad (\text{A4})$$

In the last step of Eq. (A4) we used Eq. (A2).

In summary, in counting the branches with exactly p number of A's we utilize the recurrent structure of the binary trees presented in Fig. 8. We perform the branch cuts that mark the transition from a higher level tree to a lower level tree and iterate along the branch-cut lines. The iteration is completed when the summation term has a form given by Eq. (A1). In this way, continuing from Eq. (A2), we obtain

$$\begin{aligned} M(1) &= n(1;1A(L-1)) \\ &= 1 + \sum_{s=1}^{L-2} n(0;(L-s)D(s)) \\ &= 1 + \sum_{s=1}^{L-2} 1 = L-1, \end{aligned} \quad (\text{A5})$$

where we used Eq. (A1). Similarly, continuing from Eq. (A4), after substituting Eq. (A1), gives

$$\begin{aligned} M(2) &= n(2;1A(L-1)) \\ &= \sum_{s=1}^{L-2} \sum_{l=1}^{s-1} \left(1 + \sum_{m=1}^{l-1} 1 \right) \\ &= \sum_{k=1}^{L-3} \frac{k(k+1)}{2} = \sum_{k=0}^{L-4} \binom{k+2}{2} \\ &= \binom{L-1}{3}. \end{aligned} \quad (\text{A6})$$

For $p=3$ the iteration leads to:

$$\begin{aligned} M(3) &= n(3;1A(L-1)) \\ &= \sum_{s=1}^{L-2} n(2;(L-s)D(s)) \\ &= \sum_{s=1}^{L-2} \sum_{l=1}^{s-1} n(2;(L-l)A(l)) \\ &= \sum_{s=1}^{L-2} \sum_{l=1}^{s-1} \sum_{r=1}^{l-1} \sum_{k=1}^{r-1} \left(1 + \sum_{m=1}^{k-1} 1 \right) \\ &= \sum_{k=1}^{L-5} \frac{k(k+1)(k+2)(k+3)}{2 \times 3 \times 4} \\ &= \sum_{k=0}^{L-6} \binom{k+4}{4} = \binom{L-1}{5}. \end{aligned} \quad (\text{A7})$$

where we used Eq. (A4) and Eq. (A1). For general $p = 1, 2, \dots, [L/2]$, we obtain

$$\begin{aligned} M(p) &= n(p;1A(L-1)) \\ &= \sum_{k_1=1}^{L-2} \dots \sum_{k_{2p-2}=1}^{k_{2p-3}-1} \left(1 + \sum_{k=1}^{k_{2p-2}-1} 1 \right) \\ &= \sum_{k=0}^{L-2p} \binom{k+2p-2}{2p-2} = \binom{L-1}{2p-1}. \end{aligned} \quad (\text{A8})$$

APPENDIX B: MOMENTS OF $f(p;L)$

For the purpose of the computation of moments we introduce the following notation:

$$\langle p^n \rangle_L = \sum_{m=1}^{[L/2]} m^n f(m;L), \quad (\text{B1})$$

where $n=0,1,2,\dots$. By manipulating the factorials in $f(m;L)$ and shifting the summation index on the rhs of Eq. (B1), it is straightforward to show that

$$\langle p^n \rangle_L = \frac{8}{L(L+1)} \sum_{k=1}^{[(L+2)/2]} (k-1)^{n+1} (2k-1) f(k;L+2). \quad (\text{B2})$$

With the substitution $L \rightarrow L-2$, Eq. (B2) becomes a recursion relation for $\langle p^n \rangle_L$:

$$\langle p^n \rangle_{L-2} \frac{(L-2)(L-1)}{8} = \langle (p-1)^{n+1} (2p-1) \rangle_L. \quad (\text{B3})$$

Using the binomial formula in Eq. (B3), the recursion relation can be explicitly written out as

$$\begin{aligned} \langle p^{n+2} \rangle_L &= \langle p^n \rangle_{L-2} \frac{(L-2)(L-1)}{4^2} + \frac{\langle p^{n+1} \rangle_L}{2} \\ &+ \sum_{i=0}^n \binom{n+1}{i} (-1)^{n+i+1} \left(\langle p^{i+1} \rangle_L - \frac{\langle p^i \rangle_L}{2} \right). \end{aligned} \quad (\text{B4})$$

To obtain $\langle p^n \rangle_L$ for arbitrary n , we iterate Eq. (B4), starting with the initial $n=0$ and using the identities

$$\langle p^0 \rangle_L = 1, \quad (\text{B5})$$

$$\langle p^1 \rangle_L = \frac{L+1}{4}. \quad (\text{B6})$$

Equation (B5) expresses the normalization condition. Equation (B6) follows from Eq. (B1) after simple algebra and from Eq. (B5).

For $n=0$, Eq. (B4) gives

$$\langle p^2 \rangle_L = \frac{L(L+3)}{4^2}. \quad (\text{B7})$$

For $n=1$, Eq. (B4) gives

$$\langle p^3 \rangle_L = \frac{(L-2)(L-1)}{4^2} \langle p^1 \rangle_{L-2} + \frac{5\langle p^2 \rangle_L}{2} - 2\langle p^1 \rangle_L - \frac{\langle p^0 \rangle_L}{2}. \quad (\text{B8})$$

Substituting Eqs. (B5)–(B7) leads to

$$\langle p^3 \rangle_L = \frac{L^3 + 6L^2 + 3L - 2}{4^3}. \quad (\text{B9})$$

In a similar fashion, for $n=2$, Eqs. (B4)–(B7) and Eq. (B9) give

$$\langle p^4 \rangle_L = \frac{L(L^3 + 10L^2 + 15L - 10)}{4^4}. \quad (\text{B10})$$

The variance σ^2 , the skewness, and the kurtosis of $f(k;L)$ can be computed in the standard way [34]. The variance is given in Eq. (10). For the skewness we obtain $\text{skew}(f)=0$. The kurtosis is a positive function of L . Explicitly, for $L \geq 4$,

$$\text{kurt}(f) = 2 \frac{3L^3 + 6L^2 - 10L + 1}{(L-1)^2}. \quad (\text{B11})$$

Equation (13) is derived in a similar way as Eq. (B3), by simple algebra and by shifting the summation index. An arbitrary power $\langle p^{-n} \rangle$ can be obtained by deriving the corresponding recursion relation, following the lines outlined above for $n > 0$.

-
- [1] K. Binder and D.W. Heermann, *Monte Carlo Simulation in Statistical Physics. An Introduction*, 3rd ed. (Springer, Berlin, 1997).
- [2] K.M. Chandy and J. Misra, *IEEE Trans. Software Eng.* **5**, 440 (1979).
- [3] R. Fujimoto, *Commun. ACM* **33**, 30 (1990).
- [4] J. Misra, *ACM Comput. Surv.* **18**, 39 (1986).
- [5] B.D. Lubachevsky, *Complex Syst.* **1**, 1099 (1987).
- [6] B.D. Lubachevsky, *J. Comput. Phys.* **75**, 103 (1988).
- [7] G. Korniss, M.A. Novotny, and P.A. Rikvold, *Comput. Phys.* **153**, 488 (1999).
- [8] B.D. Lubachevsky, V. Privman, and S.C. Roy, *J. Comput. Phys.* **126**, 152 (1996).
- [9] G. Korniss, C.J. White, P.A. Rikvold, and M.A. Novotny, *Phys. Rev. E* **63**, 016120 (2001).
- [10] G. Korniss, P.A. Rikvold, and M.A. Novotny, *Phys. Rev. E* **66**, 056127 (2002).
- [11] D.A. Jefferson, *ACM Trans. Programming Languages Syst.* **7**, 404 (1985).
- [12] Phillip M. Dickens and Paul F. Reynolds, Jr., in *Proceedings of the SCS Multiconference on Distributed Simulation, San Diego*, edited by D. Nicol and R. Fujimoto, Simulation Series Vol. 22, pp. 161-164.
- [13] A. Prakash and R. Subramanian, in *Proceedings of the Sixth Parallel and Distributed Simulation Workshop, 1992 SCS Western Multiconference* (IEEE Press, New York, 1992), p. 85.
- [14] J.S. Steinmann, in *Proceedings of the Seventh Workshop on Parallel and Distributed Simulation*, edited by R. Bagrodia and D. Jefferson (IEEE Computer Society Press, Los Alamitos, CA, 1993), p. 109.
- [15] A. Ferscha and G. Chiola, in *Proceedings of the 27th Annual Simulation Symposium, LaJolla, 1994* (IEEE Computer Society Press, Los Alamitos, CA, 1994).
- [16] P.M.A. Sloot, B.J. Overeinder, and A. Schoneveld, *Comput. Phys. Commun.* **142**, 76 (2001).
- [17] G. Korniss, Z. Toroczka, M.A. Novotny, and P.A. Rikvold, *Phys. Rev. Lett.* **84**, 1351 (2000).
- [18] A.-L. Barabási and H.E. Stanley, *Fractal Concepts in Surface Growth* (Cambridge University Press, Cambridge, 1995).
- [19] G. Korniss, M.A. Novotny, A.K. Kolakowska, and H. Guclu, in *Proceedings of the 2002 ACM Symposium on Applied Computing* (ACM, Inc., 2002), p. 132.
- [20] A. Kolakowska, M.A. Novotny, and G. Korniss, *Phys. Rev. E* **67**, 046703 (2003).
- [21] M. Kardar, G. Parisi, and Y.-C. Zhang, *Phys. Rev. Lett.* **56**, 889 (1986).
- [22] B.D. Lubachevsky, in *Proceedings of the SCS Multiconference on Distributed Simulation*, edited by B. Unger and D. Jefferson (SCS, San Diego, 1988), Vol. 19, p. 183.
- [23] D.M. Nicol, *ACM Trans. Model. Comput. Simul.* **1**, 24 (1991).
- [24] G. Korniss, M.A. Novotny, H. Guclu, Z. Toroczka, and P.A. Rikvold, *Science* **299**, 677 (2003).

- [25] P.A. Rikvold and M. Kolesik, *J. Stat. Phys.* **100**, 377 (2000).
- [26] P.A. Rikvold and M. Kolesik, *J. Phys. A* **35**, L117 (2002).
- [27] P.A. Rikvold and M. Kolesik, *Phys. Rev. E* **66**, 066116 (2002).
- [28] P.A. Rikvold and M. Kolesik, *Phys. Rev. E* **67**, 066113 (2003).
- [29] A. Kolakowska, M. A. Novotny, G. Korniss, and P. Verma (unpublished).
- [30] D. Forster, D. Nelson, and M. Stephen, *Phys. Rev. A* **16**, 732 (1977).
- [31] E. Medina, T. Hwa, M. Kardar, and Y.-C. Zhang, *Phys. Rev. A* **39**, 3053 (1989).
- [32] C. Castellano, M. Marsili, and L. Pietronero, *Phys. Rev. Lett.* **80**, 3527 (1998).
- [33] S.F. Edwards and D.R. Wilkinson, *Proc. R. Soc. London, Ser. A* **381**, 17 (1982).
- [34] W.H. Press, S.A. Teukolsky, W.T. Vetterling, and B.P. Flannery, *Numerical Recipes in Fortran 77* (Cambridge University Press, Cambridge, 1992).
- [35] Strictly speaking, for $t > 0$ the set of events when all sites are in the same elementary configuration A is of measure zero. In the absence of the periodicity condition, all sites could be in either B or C .
- [36] For convenience we drop the parameter t in the notation when the analysis concerns the steady state.