Self-organized criticality in simulated correlated systems

Sloot, P.M.A.; Overeinder, B.J.; Schoneveld, A.

Link to publication

Self-organized criticality in simulated correlated systems

P.M.A. Sloot *, B.J. Overeinder, A. Schoneveld

University of Amsterdam, Faculty of Science, Section Computational Science, Kruislaan 403, 1098 SJ Amsterdam, The Netherlands

Abstract

In this paper we study the influence of spatio-temporal correlations on the dynamic runtime behavior of the optimistic parallel Time Warp simulation method. By means of Ising spin simulation, we show experimentally that the probability distribution of the number of rolled back events behaves as a power-law distribution over a large range of sub-critical Ising temperatures and decays exponentially for super-critical Ising temperatures. The experimental results indicate that for critical Ising temperatures, where long-range correlations occur, the computational complexity of Time Warp and physical complexity of the Ising spin model are entangled and contribute both to the runtime behavior in a nonlinear way. © 2001 Elsevier Science B.V. All rights reserved.

1. Introduction

The increasing interest in the application of discrete event simulation to solve problems from natural sciences, is motivated by the ability of the discrete event simulation to capture the asynchronous behavior that is characteristic to a large class of natural phenomena. We are specifically interested in the use of parallel discrete event simulation (PDES) methods, in particular the Time Warp method [1], to study dynamic complex systems modeled by asynchronous cellular automata (ACA) [2,3]. Asynchronous cellular automata differ from regular synchronous cellular automata (CA) in their update scheme. The danger of simply imposing an update scheme has been considered in an article by Huberman and Glance [4].

The dynamic behavior of many natural systems is strongly imposed by the statistical (spatio-temporal)

correlation between the individual components of the system. In particular, critical phenomena in Ising spin systems (i.e. magnetization in ferro-metals) or individual-based population dynamics are explained by spatio-temporal correlations [5]. The fundamental question addressed in this paper is, to which extent correlations in the simulation model influence the runtime behavior, or dynamics, of the PDES method. To study this phenomenon, both correlated and uncorrelated system dynamics are simulated, and their effect on causality errors in the Time Warp method is examined.

We conjecture that the Time Warp dynamics can be characterized as a so-called self-organized critical system [6]. Self-organized critical behavior is found in many complex systems, for example, in sand-pile and earthquake dynamics. Similar to the Gutenberg–Richter histogram of earthquake magnitudes, the histogram of rollback lengths is very close to a power law. The analog suggests that large rollbacks are re-

---

* Corresponding author.
 E-mail address: sloot@wins.uva.nl (P.M.A. Sloot).

lated in a scale-invariant fashion to smaller rollbacks, i.e. large rollbacks and small rollbacks are basically different manifestations of the same phenomenon.

2. The APSIS optimistic simulation environment

The Amsterdam Parallel Simulation System (APSIS) environment is a research vehicle for both optimistic simulation protocol design and evaluation, and parallel simulation development of dynamic complex systems [7]. Specifically, requirements for computational science applications are taken into consideration to assess the potential of PDES methods to solve problems originating from, e.g., physics, chemistry, or biology. These requirements put special constraints on the design of the simulation environment and necessitate extensions to the basic Time Warp method.

In optimistic simulation, the parallel simulation processes execute events and proceed in local simulated time as long as they have any input at all. A consequence of the optimistic execution of events is that the local clock or Local Virtual Time (LVT) of a process may get ahead of its neighbors' LVTs, and it may receive an event message from a neighbor with a timestamp smaller than its LVT, that is, in the past of the simulation process. If we allow causality errors happen, we must provide a mechanism to recover from these errors in order to guarantee a causally correct parallel simulation. Recovery is accomplished by undoing the effects of all events that have been processed prematurely by the process receiving the straggler. The net effect of the recovery procedure is that the simulation process *rolls back* in simulated time. The premature execution of an event results in two things that have to be rolled back: (i) the state of the simulation process and (ii) the event messages sent to other processes. A direct consequence of the rollback mechanism is that more cancellation messages may be sent to other processes recursively, and allows all effects of erroneous computation to be eventually annihilated.

3. The Ising spin model

The Ising spin model is a classical model of a system of interacting variables in statistical physics, originally introduced as a model for ferromagnetism.

The model is defined on a grid of interacting "spins" that can either have a state "up" or a state "down". Spins on adjacent, nearest-neighbor sites, interact with each other in a pair-wise manner with a strength J.

At random times, a spin is granted a chance to change the state, a so-called spin flip. The attempted spin flip, or trial, is accepted proportionally to the energy difference ΔE between the new configuration and the old configuration. Using a constant temperature (T) condition the so-called *Metropolis algorithm* can be used to simulate the dynamics. In the Metropolis algorithm, a spin flip is accepted according to the Boltzmann probability distribution $\exp(-\Delta E/kT)$. If the temperature is below a well-defined temperature, known as the critical or *Curie* temperature T_c, a spontaneous non-zero magnetization occurs.

In this paper we use a discrete event simulation for the Metropolis algorithm to simulate the Ising spin model at different temperatures T. The parallel discrete event simulation in APSIS exploits the implicit spatial locality present in the Ising spin model, without altering the update history [8,9].

4. Self-organized criticality

4.1. Criticality versus self-organized criticality

Most of the time, equilibrium systems with short-range interactions, exhibit exponentially decaying correlations. Infinite correlations, i.e. scale invariance, can be achieved by fine-tuning some parameters (e.g., temperature) to a critical value [5]. An example of such a system is the previously mentioned Ising spin model of magnetization. The sudden alignment of the spins as $T \rightarrow T_c^+$ (T_c^+ denotes from high T) results from correlations that are propagated extensively throughout the entire system: everything depends on everything else.

Besides systems exhibiting critical behavior, a large class of non-equilibrium locally interacting, nonlinear systems spontaneously develop scale invariance. Such composite systems with many interacting degrees of freedom may evolve to a critical state in which minor events may trigger a chain reaction that can affect an arbitrary large number of constituents of the system. This state is called *self-organized criticality* (SOC) [6]. The probability of spontaneously generated

structures or events, further called avalanches, of many different sizes s show a power-law distribution $P(s) \sim s^{-\tau}$, where τ is a critical exponent and most other observables of the system have no intrinsic time or length scale. The absence of intrinsic length scale is attributed to SOC, where avalanches of all sizes contribute to keep the system perpetually in a critical state. This critical state is robust with respect to any small change in the rules of the system. The size of an avalanche can be defined in different ways. It can be measured by the number of relaxation steps needed for the chain reaction to stop or the total number of sites involved in the avalanche.

4.2. Self-organized criticality in Time Warp dynamics

An important characteristic of systems exhibiting SOC behavior is a separation of time scales. It is required that such systems are slowly driven, that is *perturbations* occur on a much larger time scale than the *diffusion* or *relaxation* dynamics. The critical state in SOC systems is furthermore characterized by a stationary state where the driving forces balance the cascades. The dynamically stationary state is obtained at the "critical point" where these two effects exactly balance.

A speculative analogy could be made with Time Warp: adding events (i.e. perturbations) causes on the one hand the event rate to grow, on the other hand rollbacks to occur (i.e. relaxation). Our experiments show that in Time Warp, the event rate eventually reaches a kind of stationary state with superimpozed rollback cascade effects. In Fig. 1, the event rate shows a so-called punctuated equilibrium, which is the hallmark of SOC. We define two time scales in Time Warp: *simulation time* and *protocol time*. The simulation time in Time Warp is updated by the rate at which the system is driven. This driving rate is determined by the dynamics of the simulation. The protocol time, needed to process a rollback, is determined by machine specific parameters. A difference with conventional SOC systems and Time Warp is that there is no explicit separation of time scales, asynchronous updates and rollbacks may intervene.

In the next section, the influence of the spatio-temporal correlation on the average rollback length and the rollback length distribution is studied. The average rollback length and rollback length distribution
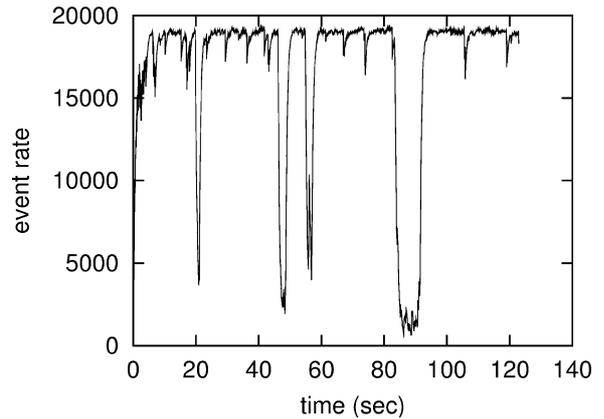

Fig. 1. Event rate of one of the logical processes in a PDES Ising simulation, with number of processors $P = 6$ and Ising temperature $T = 1.0$.

is studied at different temperatures in order to determine the influence of the Ising spin phase transition on the Time Warp protocol. It is expected that around the Ising spin phase transition, the long range spin correlations increase the average rollback sizes. We are interested in rollback length distributions in order to do a first order check of SOC in Time Warp dynamics.

5. Experimental results

The experiments with the parallel simulation of the Ising dynamics use a two-dimensional decomposition of the $L \times L$ square lattice. The decomposition is made up to assure optimal load balance. For all parameter instances of the simulation experiments we measure the average rollback length and rollback length distributions. For the first series of experiments in this section we have fixed the lattice size to $L = 220$ and the number of processors to $P = 12$.

5.1. Influence of the temperature in the correlated system

In Fig. 2 the average rollback lengths are shown for the temperature range 0.1–2.7. Ideally, rollback avalanches are measured instantaneously over the system, that is over all processors used by the parallel simulation. However, the instantaneous measurement requires freezing the forward simulation, and allowing only rollbacks to occur, i.e. the relaxation in Time
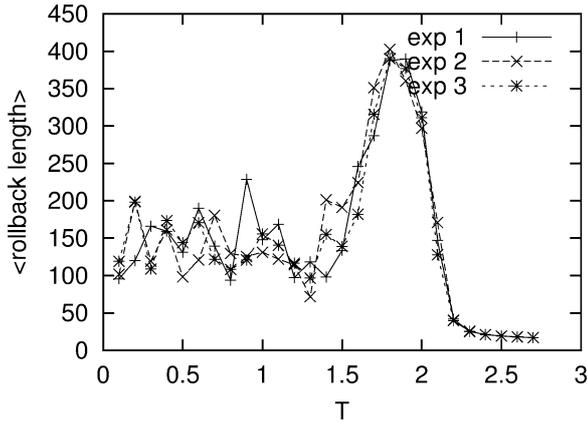
Fig. 2. Average rollback length for different temperatures. For each temperature, the results of 3 experiments are shown. Using the parameters $L = 220$, $P = 12$.

Fig. 3. Rollback distribution for temperatures in the range 0.1–1.5, for $L = 200$ and $P = 12$. The fitted exponent has value -1.3 (± 0.03).

Fig. 4. Rollback distribution for temperatures in the range 1.5–2.7, for $L = 200$, $P = 12$.

Warp. This would fundamentally alter the dynamics of the Time Warp protocol, and is therefore unacceptable intrusive. Hence, each processor records the local rollbacks for analysis. The rollback lengths are averaged over time for all processors. The results of three different runs are depicted in the figure. Close to the Ising phase transition ($T_c \approx 2.27$ for infinite lattices), the expected peak in the average rollback length can be observed. From this figure, three different regimes can be identified: the *physical sub-critical phase* ($< T_c$), the *physical critical phase* ($\approx T_c$) and the *physical super-critical phase* ($> T_c$).

The different phases influence the rollback length distributions. In the physical sub-critical temperature regime 0.1–1.4, power-law scaling is found (see Fig. 3), i.e. the Time Warp dynamics appear to be in a *computational critical* regime (i.e. the Time Warp dynamics reaches a self-organized critical state). As the temperature approaches the physical super-critical regime a transition to exponential scaling can be observed (see Fig. 4). Close to the critical temperature, length distributions with "fat tails" (power-law distributions with exponential cutoff) develop due to the emergence of long-range spin correlations.

The scaling exponent α in the physical sub-critical regime seems to be universal for all temperatures in this regime. From the experimental data a power law with exponent $\alpha = -1.3$ is fitted. Because the rollback length distribution obeys power-law scaling,
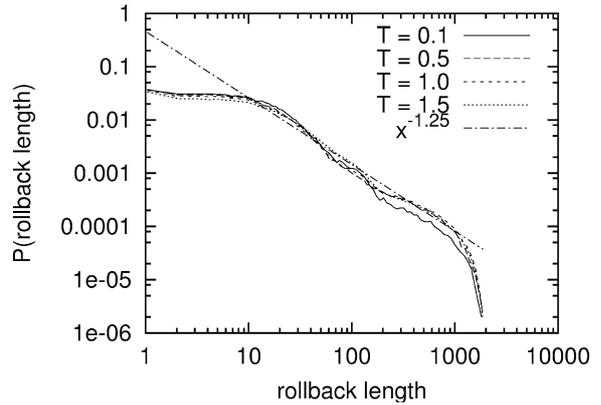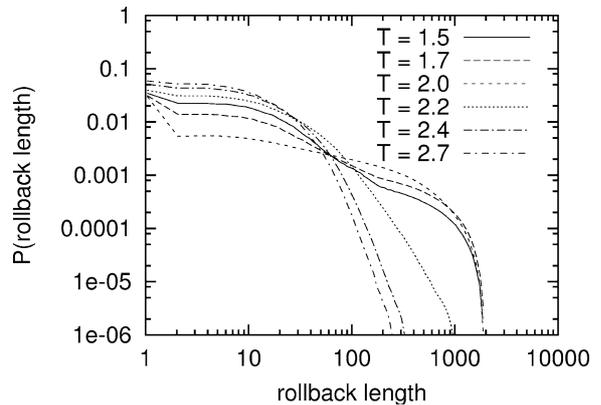
we conjecture that the rollback dynamics are in a SOC regime.

The slope of the power function flattens out in the range of its lower magnitude rollbacks with length 1–10. This flattening, rather than remaining double log-linear, occurs because discreteness effects of the rollback lengths come into play. As the rollback cascades approach the size of a single event, it becomes impossible for the fractal pattern to repeat at this scale [6].

5.2. Varying the number of processors

To study the influence of the number of processors on the rollback length distribution in the SOC regime,

a series of experiments with $P = \{4, 8, 12, 24\}$ using a 2D decomposition has been performed. The lattice size has been fixed to $L = 220$. The rollback distributions of Ising spin simulations at $T = 1.0$ are shown in Fig. 5. Similar results are seen for other temperatures T in the SOC regime. The results indicate that for increasing P the rollback length distributions converge. Again, the scaling exponent α is not influenced by increasing P.

The average rollback lengths for different processors in the range $T = [0.1 - 2.7]$ are shown in Fig. 6. For 8, 12 and 24 processors, again, a peak around T_c can be distinguished. For $P = 4$, this peak is not present. Apparently, the critical Ising spin dynamics does not reduce the performance of the Time Warp protocol if only 4 processors are used.
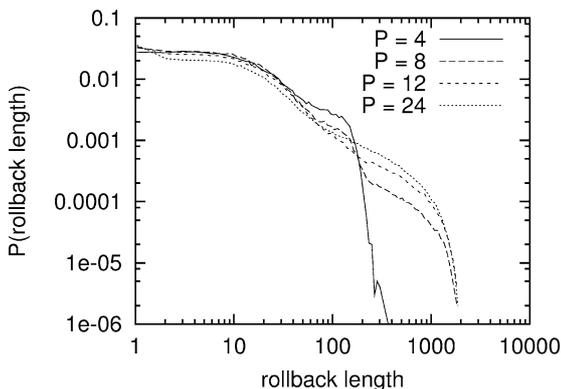

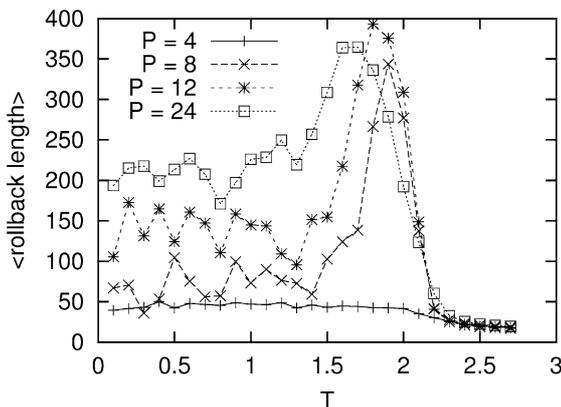Fig. 5. Rollback distributions for $P = \{4, 8, 12, 24\}$ at $T = 1.0$ for $L = 220$.

Fig. 6. Average rollback lengths for $P = \{4, 8, 12, 24\}$ for varying T and $L = 220$.

In the low temperature regime the average rollback length increases with the number of processors. This is not valid anymore in the high T regime, where the rollback length distributions approximately collapse to the same exponentially decreasing distribution. As a direct consequence, the average rollback lengths will collapse (see Fig. 6). In the high T regime the rollback lengths are not influenced by P as in the low T regime. The processors are synchronized frequently in this regime, due to a high acceptance ratio of flipped spins. Therefore, there is hardly any real time to build large simulation time differences between the processors, resulting in only small rollback lengths.

6. Conclusions

In this paper we have studied the influence of spatio-temporal correlations on the dynamical behavior of the Time Warp protocol for parallel discrete event simulations. As a simulation case we considered the Ising spin model, which is basically a cellular automata model.

For the rollback dynamics three different phases can be distinguished: *physical sub-critical*, *physical critical*, and *physical super-critical* rollback length scaling behavior. In the sub-critical regime the scaling behavior appears to behave like a power-law, with exponents independent of the temperature. In this regime we conjecture that *computational critical* (SOC) behavior appears. Around the critical phase large rollback lengths become more abundant due the long-range spin correlations. Here the computational complexity and the physical complexity are entangled and contribute both to the runtime and rollback behavior in a nonlinear way. In the physical super-critical phase a negative exponential distribution of the rollback lengths is observed.

Obviously a lot of work remains to be done in the study of physical- and computational critical behavior in Time Warp. Other series of experiments with different parameters for the Ising lattice size and bounded time window, support our findings. Complementary studies are necessary, such as the simulation of the dynamic behavior of Time Warp and the Ising spin simulation. The simulation of the model of Time Warp and Ising spin simulation allows for a full statistical analysis of the runtime behavior.

The results presented in this paper are, to our knowledge, the first series of experiments that have ever been conducted to study the influence and the appearance of critical behavior in Time Warp. The entanglement of the computational and physical complexity, and their non-trivial contribution to the runtime behavior might have consequences for other optimistic simulations.

References

[1] D.R. Jefferson, Virtual time, ACM Trans. Programming Languages Systems 7 (3) (1985) 404–425.

[2] B.J. Overeinder, P.M.A. Sloot, Application of Time Warp to parallel simulations with asynchronous cellular automata, in: Proceedings of the 1993 European Simulation Symposium, Delft, The Netherlands, 1993, pp. 397–402.

[3] P.M.A. Sloot, B.J. Overeinder, Time-Warped automata: Parallel discrete event simulation of asynchronous CA's, in: Proceedings of the Third International Conference on Parallel Processing and Applied Mathematics, Kazimierz Dolny, Poland, 1999, pp. 43–62.

[4] B.A. Huberman, N.S. Glance, Evolutionary games and computer simulations, Proc. Nat. Acad. Sci. USA 90 (16) (1993) 7716–7718.

[5] H.E. Stanley, Scaling, universality, and renormalization: Three pillars of modern critical phenomena, Rev. Modern Phys. 71 (2) (1999) S358–S366.

[6] P. Bak, C. Tang, K. Wiesenfeld, Self-organized criticality, Phys. Rev. A 38 (1) (1988) 364–374.

[7] B.J. Overeinder, Distributed event-driven simulation: Scheduling strategies and resource management, Ph.D. thesis, Department of Computer Science, University of Amsterdam, Amsterdam, The Netherlands, November 2000.

[8] B.D. Lubachevsky, Efficient parallel simulation of dynamic Ising spin systems, J. Comput. Phys. 75 (1) (1988) 103–122.

[9] G. Korniss, Z. Toroczkai, M.A. Novotny, P.A. Rikvold, From massively parallel algorithms and fluctuating time horizons to non-equilibrium surface growth, Phys. Rev. Lett. 84 (6) (2000) 1351–1354.