

Markus Aschwanden (Ed.)

Self-Organized
Criticality
Systems

Open Academic Press

Markus J. Aschwanden (Editor)

Self-Organized Criticality Systems

Dr. Markus Josef Aschwanden (Editor)
Lockheed Martin, Advanced Technology Center
Solar & Astrophysics Lab, Dept., Org. ADBS, Bldg.252
3251 Hanover Street, Palo Alto, CA 94304, USA
Telephone: 650-424-4001, Fax: 650-424-3994
E-mail: aschwanden@lmsal.com

Published by Open Academic Press Berlin Warsaw 2013
Open Academic Press, Großbeerenstraße 2-10, 12107 Berlin, Germany
Open Academic Press/Villa Europa, ul. Sienna 64, 00-807 Warsaw, Poland
Copyright ©Open Academic Press

All chapters in this book are Open Access distributed under the Creative Commons Attribution 3.0 license, which allows users to download, copy and build upon published articles even for commercial purposes, as long as the author and publisher (Open Academic Press) are properly credited, which ensures maximum dissemination and a wider impact of our publications. After this work has been published by Open Academic Press, authors have the right to republish it, in whole or part, in any publication of which they are the author, and to make other personal use of the work. Any republication, referencing or personal use of the work must explicitly identify the original source.

Contents

1	Introduction	1
	Norma B. Crosby	
1.1	A New Theory Emerges	1
1.2	Frequency distribution: a powerful tool.....	4
1.3	Self-organized criticality and powerlaw behavior	6
1.3.1	Does powerlaw behavior automatically imply SOC?	8
1.3.2	SOC and SOC-like models	9
1.4	Where is SOC observed ?	11
1.4.1	Phenomena on Earth showing SOC behavior	12
1.4.2	Phenomena in space showing SOC behavior	13
1.5	Searching for a common signature: What does it all mean? .	14
	References	19
2	Theoretical Models of SOC Systems	23
	Markus J. Aschwanden	
2.1	Cellular Automaton Models (CA-SOC)	24
2.1.1	Statistical Aspects	25
2.1.2	Physical Aspects	28
2.2	Analytical SOC Models	31
2.2.1	Exponential-Growth SOC Model (EG-SOC)	31
2.2.2	The Fractal-Diffusive SOC Model (FD-SOC)	36
2.2.3	Astrophysical Scaling Laws	45
2.2.4	Earthquake Scaling Laws	47
2.3	Alternative Models Related to SOC	48
2.3.1	Self-Organization Without Criticality (SO)	48
2.3.2	Forced Self-Organized Criticality (FSOC)	49
2.3.3	Brownian Motion and Classical Diffusion.....	50
2.3.4	Hyper-Diffusion and Lévy Flight	52
2.3.5	Nonextensive Tsallis Entropy	54
2.3.6	Turbulence	55
2.3.7	Percolation.....	57

2.3.8	Phase Transitions	58
2.3.9	Network Systems	60
2.3.10	Chaotic Systems	61
2.3.11	Synopsis	64
2.4	References	65
3	SOC and Fractal Geometry	73
	R. T. James McAteer	
3.1	From chaos to order (and back)	74
3.1.1	The coffee table fractal	75
3.1.2	The n-body problem	76
3.1.3	The butterfly effect	77
3.1.4	The critical points	78
3.2	Fractal Properties	79
3.2.1	Dimensionality	80
3.2.2	Self-similarity and Scale-Invariance	81
3.2.3	Generating Fractals	81
3.3	The many flavors of fractal dimension	84
3.3.1	Similarity Dimension of a set of systems	84
3.3.2	Box Counting Dimension	85
3.3.3	The Hölder Exponent	86
3.3.4	The Hurst Exponent	86
3.3.5	Hausdorff Dimension	87
3.4	Multifractals	90
3.4.1	From monofractals to multifractals	90
3.4.2	Generalized Dimensions	91
3.4.3	Connecting forms of multifractality	92
3.4.4	The Devils staircase	94
3.4.5	The Wavelet Transform Modulus Maxima	95
3.5	Future directions	98
	References	99
4	Percolation Models of Self-Organized Critical Phenomena 103	
	Alexander V. Milovanov	
4.1	The Percolation Problem	104
4.1.1	Site and Bond Percolation	105
4.1.2	Percolation Critical Exponents β , ν , and μ	106
4.1.3	Random Walks on Percolating Clusters	107
4.1.4	The Spectral Fractal Dimension	108
4.1.5	The Alexander-Orbach Conjecture	108
4.1.6	Percolation Problem on the Riemann Sphere	109
4.1.7	Summary	111
4.2	The SOC Hypothesis	111
4.2.1	SOC vs. Percolation	112
4.2.2	The Guiding Mechanisms	113
4.3	Going With the Random Walks: DPRW Model	114

4.3.1	Description of the Model	115
4.3.2	Random-Walk Hopping Process	116
4.3.3	Dynamical Geometry of Threshold Percolation	117
4.4	Linear-Response Theory	117
4.4.1	Dynamics and Orderings	117
4.4.2	Frequency-Dependent Conductivity and Diffusion Coefficients	118
4.4.3	Power-Law Power Spectral Density	120
4.4.4	Stretched-Exponential Relaxation and the Distribution of Relaxation Times	121
4.4.5	Consistency Check	124
4.4.6	Fractional Relaxation and Diffusion Equations	125
4.4.7	Derivation of the Fractional Diffusion Equation	126
4.4.8	Dispersion-Relation Exponent	127
4.4.9	The Hurst Exponent	128
4.4.10	Activation-Cluster Size Distribution and the τ -Exponent	129
4.4.11	Occurrence Frequency Energy Distribution and the β -Exponent	129
4.4.12	Values of the Critical Exponents	130
4.5	The Random Walk's Guide to SOC	133
4.5.1	General	133
4.5.2	The Role of Random Walks	134
4.5.3	Universality Class	134
4.6	Self-Organized Turbulence: The "Sakura" Model	135
4.7	Beyond Linear Theories: DANSE Formalism	139
4.7.1	The Roadmap	139
4.7.2	DANSE Equation	141
4.7.3	Coupled Nonlinear Oscillators	142
4.7.4	Chaotic vs. Pseudochaotic Dynamics	143
4.7.5	Nearest-Neighbor Rule	145
4.7.6	Pseudochaotic Dynamics on a Cayley Tree	145
4.7.7	Making Delocalization Transition Self-Organized	148
4.7.8	Asymptotic Spreading of the Hole Wave Function	148
4.7.9	Summary	150
4.8	The Two Faces of Nonlinearity: Instability of SOC	150
4.8.1	Instability Cycle	151
4.8.2	"Fishbone"-Like Instability	154
4.8.3	The Threshold Character of Fishbone Excitation	155
4.8.4	Fractional Nonlinear Schrödinger Equation	156
4.8.5	Mixed SOC-Coherent Behavior	161
4.9	Phase Transitions in SOC Systems	162
4.9.1	Subordination to SOC	163
4.9.2	Generalized Free Energy Expansion	164
4.9.3	Fractional Ginzburg-Landau Equation	164

4.9.4	The q -Exponent	165
4.10	Overall Summary and Final Remarks	166
4.10.1	Finance	169
4.10.2	Climate Dynamics	170
4.11	The Frontier	172
	References	174
5	Criticality and Self-Organization in Branching Processes: Application to Natural Hazards	183
	Álvaro Corral and Francesc Font-Clos	
5.1	The Statistics of Natural Hazards	184
5.1.1	The Gutenberg-Richter Law	184
5.1.2	A First Model for Earthquake Occurrence	187
5.2	Branching Processes	188
5.2.1	Definition of the Galton-Watson Process	189
5.2.2	Generating Functions	189
5.2.3	Distribution of Number of Elements per Generation	192
5.2.4	Expected Number of Elements per Generation	192
5.2.5	The Probability of Extinction	193
5.2.6	The Probability of Extinction for the Binomial Distribution	195
5.2.7	No Stability of the Population	196
5.2.8	Non-Equilibrium Phase Transition	197
5.2.9	Distribution of the Total Size of the Population: Binomial Distribution and Rooted Trees	199
5.2.10	Generating Function of the Total Size of the Population	204
5.2.11	Self-Organized Branching Process	207
5.2.12	Self-Organized Criticality and Sandpile Models	209
5.3	Conclusions	213
	Appendix	214
	References	226
6	Power Laws of Recurrence Networks	229
	Yong Zou, Jobst Heitzig, Jürgen Kurths	
6.1	Introduction	229
6.2	Power-law scaling and singularities of the invariant density .	231
6.2.1	One-dimensional maps: Analytical theory	233
6.2.2	Example: Generalized logistic map	235
6.3	Power-laws and fixed points in 2D flows	237
6.4	Power-law scaling versus fractal dimension	238
6.5	Technical aspects	241
6.5.1	Estimation of scaling exponents	242
6.5.2	Selection of dynamical variable	243
6.6	Conclusions	245
	References	246

7	SOC computer simulations	251
	Gunnar Pruessner	
7.1	Introduction	251
	7.1.1 Observables	255
	7.1.2 Models	256
7.2	Scaling and numerics	272
	7.2.1 Simple scaling	273
	7.2.2 Moment analysis	278
	7.2.3 Statistical errors from chunks	282
7.3	Algorithms and data organisation	284
	7.3.1 Stacks	285
	7.3.2 Sites and Neighbours	290
	7.3.3 Floating Point Precision	294
	7.3.4 Random Number Generators	295
	7.3.5 Output	297
7.4	Summary and conclusion	300
	Appendix: Implementation details for binning	301
	References	303
8	SOC Laboratory Experiments	311
	Gunnar Pruessner	
8.1	Introduction	311
	8.1.1 Identifying SOC in experimental data	312
	8.1.2 Tools and features	314
8.2	Granular Media	317
8.3	Systems with internal disorder	320
8.4	Mechanical instabilities: Fracture and rapture	323
8.5	Biological systems	325
	References	327
9	Self-Organizing Complex Earthquakes: Scaling in Data, Models, and Forecasting	333
	Michael K. Sachs, John B. Rundle, James R. Holliday, Joseph Gran, Mark Yoder, Donald L. Turcotte and William Graves	
9.1	Introduction	334
9.2	Earthquakes	335
9.3	Characteristic Earthquakes	337
9.4	Models of Earthquakes	342
9.5	Forecasting	343
9.6	Results	352
9.7	Summary	354
	References	354

10	Wildfires and the Forest-Fire Model	357
	Stefan Hergarten	
10.1	The Forest-Fire Model	358
10.2	Numerical and Theoretical Results	359
10.3	The Relationship to Real Wildfires	364
10.4	Extensions of the Forest-Fire Model	368
	References	375
11	SOC in Landslides	379
	Stefan Hergarten	
11.1	Landslide Statistics	382
11.2	Mechanical Models	386
11.3	Geomorphic Models	390
	References	399
12	SOC and Solar Flares	403
	Paul Charbonneau	
12.1	Introduction: solar magnetic activity and flares	403
12.2	Parker's coronal heating hypothesis	406
12.3	SOC Models of solar flares	410
12.3.1	The Lu & Hamilton model	410
12.3.2	Sample results	413
12.4	Physical interpretation	416
12.4.1	The lattice and nodal variable	417
12.4.2	The stability criterion	417
12.4.3	Computing the released energy	418
12.4.4	Nodal redistribution as nonlinear diffusion	419
12.4.5	Reverse engineering of discrete redistribution rules .	421
12.5	Beyond the sandpile	423
12.5.1	Numerical simulations	424
12.5.2	SOC in reduced MHD	426
12.5.3	Fieldline-based models	427
12.5.4	Loop-based models	429
12.6	Outlook	432
12.7	References	433
13	SOC Systems in Astrophysics	439
	Markus J. Aschwanden	
13.1	Theory	440
13.1.1	The Scale-Free Probability Conjecture	441
13.1.2	The Fractal-Diffusive Spatio-Temporal Relationship	443
13.1.3	Size Distributions of Astrophysical Observables	445
13.1.4	Scaling Laws for Thermal Emission of Astrophysical Plasmas	448
13.1.5	Scaling Laws for Astrophysical Acceleration Mechanisms	450

13.2	Observations	451
13.2.1	Lunar Craters	452
13.2.2	Asteroid Belt	454
13.2.3	Saturn Ring	455
13.2.4	Magnetospheric Substorms and Auroras	457
13.2.5	Solar Flares	458
13.2.6	Stellar Flares	465
13.2.7	Pulsars	468
13.2.8	Soft Gamma-Ray Repeaters	470
13.2.9	Black Hole Objects	470
13.2.10	Blazars	473
13.2.11	Cosmic Rays	473
13.3	Conclusions	475
13.4	References	478

Chapter 7

SOC computer simulations

Gunnar Pruessner

Abstract The following chapter provides an overview of the techniques used to understand Self-Organised Criticality (SOC) by performing computer simulations. Those are of particular significance in SOC, given its very paradigm, the BTW (Bak-Tang-Wiesenfeld) sandpile, was introduced on the basis of a process that is conveniently implemented as a computer program. The chapter is divided into three sections: In the first section a number of key concepts are introduced, followed by four brief presentations of SOC models which are most commonly investigated or which have played an important part in the development of the field as a whole. The second section is concerned with the basics of scaling with particular emphasis of its rôle in numerical models of SOC, introducing a number of basic tools for data analysis such as binning, moment analysis and error estimation. The third section is devoted to numerical methods and algorithms as applied to SOC models, addressing typical computational questions with the particular application of SOC in mind. The present chapter is rather technical, but hands-on at the same time, providing practical advice and even code snippets (in C) wherever possible.

7.1 Introduction

The concept of Self-Organised Criticality (SOC)¹ was introduced by Bak *et al.* (1987) on the basis of a computer model, the famous BTW Sandpile. The notion of “computer model” and “simulation” used here is subtle and can be misleading. Often the models are not meant to mimic a particular (natural) phenomenon, but are intended to capture merely what is considered

Imperial College London, Department of Mathematics, e-mail: g.pruessner@imperial.ac.uk

¹ A more extensive review on the present subject area can be found in (Pruessner 2012b).

to be the *essential* interaction observed in a natural phenomenon. Per Bak in particular, had the tendency to name models according to their appearance rather than their purpose and so the “Sandpile Model” may not have been envisaged to display the dynamics of a sandpile. The situation is clearer in the case of the “Forest Fire Model” (Bak *et al.* 1990), which was developed as a model of turbulence much more than as a model of fires in woods.

In particular in the early days of SOC modelling, the models were sometimes referred to as “cellular automata” Olami *et al.* (1992), Lebowitz *et al.* (1990), which caused some consternation (*e.g.* Grassberger 1994), as cellular automata normally have discrete states and evolve in discrete time steps according to deterministic rules in discrete space (*i.e.* a lattice). The term “coupled map lattice” (Kaneko 1989) can be more appropriate for some models, such as the Olami-Feder-Christensen Model discussed below (discrete space, continuous state and possibly continuous time).

The terminology of “numerical modelling” has always been somewhat confusing. Many of the models considered in SOC do not model a natural phenomenon and so their numerical implementation is not a “numerical simulation” in the sense that they mimic the behaviour of something else. There are notable exceptions, however, such as the Forest Fire Model (Bak *et al.* 1990) mentioned above and the Oslo ricepile model (Christensen *et al.* 1996). SOC models generally are not “models of SOC”, rather they are algorithmic prescriptions or “recipes” for a (stochastic) process that is believed to exhibit some of the features normally observed in other SOC models. In that sense, the terminology of terms like “SOC models” and “simulation” or even “simulating an SOC model” is misleading — most of these models are not simplified versions or idealisations of some physical process or anything else that is readily identified as “SOC”, but recipes to produce some of the behaviour expected in an SOC system.

To this day, a large fraction of the SOC community dedicate their research to computer models. Initially, the motivation (*e.g.* Zhang 1989, Manna 1991) was to find models displaying the same universal behaviour as the BTW (Bak-Tang-Wiesenfeld) Sandpile. This was followed by an era of proliferation, when many new models, belonging to new universality classes were developed. More recently, in a more reductionistic spirit, new models are mostly developed to isolate the rôle of particular features and to extract and identify their effect (*e.g.* Tadić and Dhar 1997). A lot of numerical research into SOC nowadays happens “en passant”, as SOC is identified in a model for a phenomenon that originally was not considered to be related to SOC (*e.g.* Burridge and Knopoff 1967).

Virtually all SOC (computer) models consist of degrees of freedom interacting with (nearest) neighbours located on a lattice. The degrees of freedom may be parameterised by continuous or discrete variables, in the following denoted $z_{\mathbf{n}}$, where \mathbf{n} is a position vector on the lattice. A **slow, external driving mechanism** (in short, **external drive**) slowly loads the system, *i.e.* the local variables are slowly increased, also referred to as “charging a

site”. That might happen uniformly (sometimes called **global drive**) or at individual lattice sites (sometimes called **point drive**). The driving might happen at randomly chosen points or by random increments, both of which is in the literature referred to as **random driving**. The dynamics of an SOC model is **non-linear**, *i.e.* there is no linear equation of motion that would describe their dynamics.² The response of the system is triggered by a local degree of freedom overcoming a **threshold**, beyond which **relaxation** and thus interaction with other degrees of freedom and the outside world takes place. A site where that happens is said to **topple** and to be **active**. The interaction might lead to one of the neighbours exceeding its threshold in turn, triggering another relaxation event. The totality of the relaxations constitutes an **avalanche**. When the avalanche has finished, *i.e.* there are no active sites left, the system is in a state of **quiescence**. In SOC models, driving takes place only in the quiescent state (separation of time scales, below). If the external drive acts at times when an avalanche is running, it might lead to a continuously running avalanche (e.g. Corral and Paczuski 1999).

In many models the degree of freedom at every site measures a resource that is **conserved** under the dynamics. To balance the external drive, in most models **dissipation** has to take place in some form: **Bulk dissipation** takes place when the resource can get lost in the local interaction. **Boundary dissipation** refers to the situation when the resource is lost only in case a boundary site relaxes. The necessary flux of the resource towards the boundaries has been suggested as some of the key mechanisms in SOC (Paczuski and Bassler 2000b). In some models, such as the Bak-Sneppen Model (Bak and Sneppen 1993) or the Forest-Fire-Models (Henley 1989, Bak *et al.* 1990, Drossel and Schwabl 1992a), no (limited) resource can be identified and therefore the notion of dissipation and conservation is not meaningful.

The question whether conservation is a necessary ingredient of SOC has driven the evolution of SOC models in particular during the 1990s. In fact, early theoretical results by Hwa and Kardar (1989a) suggested that bulk dissipation would spoil the SOC state. Models like the OFC Model (Olami *et al.* 1992, also Bak and Sneppen 1993, Drossel and Schwabl 1992a) questioned that finding. Different theoretical views have emerged over time: Lauritsen *et al.*'s (1996) self-organised branching process (Zapperi *et al.* 1995) contains dissipation as a *relevant* parameter which has a limiting effect on the scaling behaviour. Juanico *et al.* (2007) restored the SOC state of the self-organised branching process by implementing a mechanism that compensates for the non-conservation by a “matching condition” not dissimilar from the mechanism used in the mean-field theory by Pruessner and Jensen (2002b). That, in turn, was labelled by Bonachela and Muñoz (2009) as a form of tuning.

² It is very instructive to ask why a non-linearity is such a crucial ingredient. Firstly, if all interactions were linear, one would expect the resulting behaviour to correspond to that of a solvable, “trivial” system. Secondly, linearity suggests additivity of external drive and response, so responses would be expected to be proportional to the drive, a rather boring behaviour, not expected to result in scale invariance.

More recent field-theoretic work (Pruessner 2012b) points at conservation as a symmetry responsible for the cancellation of mass-generating diagrams, an effect that may equally be achieved by other symmetries.

The external drive, the ensuing sequence of avalanches and the evolution of the model from one quiescent state to the next happen on the **macroscopic time scale**, where time typically passes by one unit per avalanche. As the system size is increased, avalanches are expected to take more and more relaxations to complete. Their duration is measured on the **microscopic time scale**. In the thermodynamic limit, *i.e.* at infinite system size, the infinite duration of an avalanche of the microscopic time scale and the finite driving rate on the macroscopic time scale amount to a complete **separation of time scales**. In general, the separation of time scales is achieved in finite systems provided that no driving takes place when any site is active, because the times of quiescence, measured on the microscopic time scale, can be thought of as arbitrarily long. As a result, the avalanching in these systems becomes **intermittent**.

Separation of time scales is widely regarded as *the* crucial ingredient of SOC, maybe because it is conceived (and criticised as such) as a substitute of the tuning found in traditional critical phenomena (also Jensen 1998). In numerical models, it normally enters in a rather innocent way — the system is not driven while an avalanche is running. This, however, requires some global supervision, a “babysitter” (Dickman *et al.* 2000) or a “farmer” (Bröker and Grassberger 1999). In some models the separation of time scales can be implemented explicitly (Bak and Sneppen 1993) in the relaxational rule. What makes the separation of time scales very different from other forms of tuning is that it *eliminates* a dimensionful, finite scale, such as the frequency with which an avalanche is triggered.³ In traditional critical phenomena, scaling comes about due to the *presence* of a dimensionful, finite energy scale⁴, where entropic contributions to the free energy compete with those from the internal energy promoting order. In most SOC models, it is pretty obvious that scaling would break down if time scales were not explicitly separated — avalanches start merging and eventually intermittency is no longer observed (Corral and Paczuski 1999).

SOC models are normally studied at **stationarity**, when all correlations originating from the initial state (often the empty lattice) are negligible. Reaching this point is a process normally referred to as **equilibration**. The equilibration time is normally measured as the number of charges by the external drive required to reach stationarity. For some models, exact upper bounds for the equilibration time are known (Dhar *et al.* 1995, Corral 2004a, Dhar 2004, *e.g.*). In deterministic models, a clear distinction exists between **transient** and **recurrent states**, where the former can appear at most once,

³ In the field theory of SOC, the cancellation of diagrams occurs precisely when stationarity is imposed for the density of particles resting (and their correlations) in the limit $\omega \rightarrow 0$, *i.e.* in the long time limit.

⁴ For example $k_B T_c$ in the Ising Model (Stanley 1971).

and the latter with a finite frequency provided the number of states overall is finite. In fact, this frequency is the same for all recurrent states, depending on the driving, which can be at one site only or randomly and independently throughout. A detailed proof of such properties can be cumbersome (Dhar 1999a,b).

The statistics of the avalanches, their size as well as their extent in space and in time, is collected and analysed. SOC is usually said to be found in these models when the statistics displays a **scaling symmetry**, governed by only one upper cutoff which diverges with the system size. In principle, a Gaussian possesses this scaling symmetry,⁵ but not a single important SOC model has a Gaussian event size distribution. On the contrary, the avalanche statistics of all models discussed below deviates dramatically from a Gaussian, thus suggesting that avalanches are not the result of essentially independent patches of avalanching sites creating a bigger overall avalanche. Rather, sites are **strongly interacting**, thereby creating the overall event. The purpose of numerical simulations is to characterise and quantify this interaction and its effect, as well as extracting **universal quantities**, which can be compared with those found in other systems.

7.1.1 Observables

As for the methods of analysis, they have matured considerably over the past decades. The initial hunt for $1/f$ noise in temporal signals has given way to the study of event size distributions. As a matter of numerical convenience, these distributions are often characterised using moments, some of which are known exactly. Since the beginning of computational physics, moments and cumulants have been the commonly used method of choice to characterise critical phenomena (Binder and Heermann 1997). It is probably owed to the time of the late 1980's that memory-intensive observables such as entire distributions became computationally affordable and subsequently the centre of attention in SOC.

To this day, the analysis of moments in SOC is still often regarded as an unfortunate necessity to characterise distributions, which are difficult to describe quantitatively. Apart from the historic explanation alluded to above, there is another, physical reason for that, the **avalanche size exponent** τ . In traditional critical phenomena, the corresponding exponent of the order parameter distribution is fixed at unity in the presence of the Rushbrooke and the Josephson scaling law (Christensen *et al.* 2008). The deviation of τ from unity, which implies that the expected event size does not scale like

⁵ The basic example $\mathcal{P}(s) = s^{-1}\mathcal{G}(s/s_c)$ with $\mathcal{G}(x) = 2x \exp(-x^2)/\sqrt{\pi}$ is normalised and has avalanche size exponent $\tau = 1$, as defined in Eq. (7.3). *Without* the pre-factor x in $\mathcal{G}(x)$ the graph looks surprisingly similar to a PDF as typically found in SOC models.

the characteristic event size, is another distinctive feature of SOC. To some extent, the exponent τ can be extracted from the avalanche size distribution (almost) by inspection. In a moment analysis, on the other hand, it is somewhat “hidden” in the details.

The most important observables usually extracted from an SOC model are thus the scaling exponents, such as τ , D (**avalanche dimension**), α (**avalanche duration exponent**) and z (**dynamical exponent**) discussed below. Here, the two exponents D and z are generally regarded as more universal than τ and α , as the former is often “enslaved” by an exact scaling law related to the average avalanche size, and the latter by a similar scaling law based on the “narrow joint distribution assumption”, discussed in Sec. 7.2. Generally, all observables that are universal or suspected to be are of interest. This includes the scaling function (Sec. 7.2) which is most easily characterised by moment ratios, corresponding to universal amplitude ratios, traditionally studied in equilibrium critical phenomena (Privman *et al.* 1991, Salas and Sokal 2000).

7.1.2 Models

There is wide consensus on a number of general features of SOC models which seem to play a rôle in determining the universality class each belongs to. The very first SOC model, the BTW model, was essentially **deterministic**, *i.e.* there was no randomness in the bulk relaxation. A given configuration plus the site being charged next determines the resulting configuration uniquely. Even in these models, however, there can be a degree of **stochasticity**, namely when the site to be charged by the external drive is chosen at random. Finally, even when this is not the case, *i.e.* external drive and internal relaxation are deterministic, initial conditions are often chosen at random and averaged over.

Deterministic SOC models have the great appeal that they are “autonomous” (in a non-technical sense) or “self-sufficient” in that they do not require an additional source of (uncorrelated) noise. It is difficult to justify the existence of an external source which produces white, Gaussian noise, as that noise correlator, $\langle \eta(t)\eta(t') \rangle = 2\Gamma^2\delta(t-t')$, itself displays a form of scaling $\langle \eta(\alpha t)\eta(\alpha t') \rangle = \alpha^{-1}\langle \eta(t)\eta(t') \rangle$. The presence of an external (scaling) noise source seems to demote an SOC model to a conversion mechanism of scale invariance, which becomes most apparent when the respective model is cast in the language of stochastic equations of motion, *i.e.* **Langevin equations**.

Famous examples of deterministic SOC models, which do not require an external noise source for the relaxation process, are the BTW model with deterministic drive (Bak *et al.* 1987, but Creutz 2004), the OFC model (Olami *et al.* 1992) and, closely related, the train model (de Sousa Vieira 1992).

Of these only the latter has been studied extensively in the absence of all stochasticity.

Most SOC models, however, have a strong stochastic component, *i.e.* there is some randomness in the relaxation mechanism that gives rise to avalanches. In fact, models with some form of built-in randomness seem to give cleaner scaling behaviour, suggesting that deterministic models get “stuck” on some trajectory on phase space, where some conservation law prevents them from exploring the rest of phase space (Bagnoli *et al.* 2003, Casartelli *et al.* 2006). Notably, randomising the BTW model seems to push it into the Manna universality class (Karmakar *et al.* 2005). The latter model is probably the simplest SOC model displaying the most robust and universal scaling behaviour (Huynh *et al.* 2011). Due to the noise, trajectories of particles deposited by the external drive are those of random walkers.

The second dividing line distinguishes **Abelian** and **non-Abelian** models. The term was coined by Dhar (1990) introducing, strictly speaking, the Abelian Sandpile Model, by re-expressing the original BTW Model (Bak *et al.* 1987) in terms of units of slope rather than local particle numbers. This convenient choice of driving and boundary conditions renders the model unphysical as entire rows of particles are added and removed at once. At the same time, however, the model’s final state after two consecutive charges at two different sites becomes independent from the order in which the charges and the subsequent relaxations are carried out. Practically all analytical insight into the BTW model is based on Dhar’s (1990) Abelian version. Because it is easier to implement, it has also favoured in numerical simulations.

The term “Abelian” seems to suggest the existence of a (commutative) group, *i.e.* a set of operators closed under consecutive application, associative and containing inverse and an identity. For most SOC models referred to as Abelian, no such group is known, for example because operators do not exist explicitly, or the associative property makes little sense, similarly for the identity. Crucially, inverse operators rarely exist. To label a model Abelian therefore normally means that the final state does not depend on the order in which external charges are applied, *i.e.* the model updating operators (whether or not they exist), which drive it at various locations, commute. Because the final state is unique only in the case of deterministic models, stochastic models are Abelian provided that the statistics of the final state does not depend on the order in which external charges are applied (Dhar 1999b). The operators, which generally depend on the site the driving is applied to, of deterministic models apply to a model’s state and take it from one quiescent state to the next. The operators in a stochastic model act on the distribution of states, *i.e.* they are the Markov operators. A deterministic model can be cast in the same language, however, the Markov operators then correspond to simple permutation matrices.

While Abelianness originally refers to the evolution of a model on the macroscopic time scale, it is generally used to characterise its behaviour on the microscopic timescale, *i.e.* the step-by-step, toppling-to-toppling update.

It is therefore usually concluded that the properties of avalanches and their statistics is independent from the order of toppling of multiple active sites.

Strictly, however, the Abelian symmetry does not apply to the microscopic time scale, at least for two reasons. Firstly, the Abelian operators apply, a priori, only to the avalanche-to-avalanche evolution, *i.e.* the macroscopic time scale. What is more, they apply to the final state and its statistics, but not necessarily to the observables. Applying charges at two different sites of an Abelian SOC model, starting from the same configuration, results in the same final state (or its statistics) regardless of the order in which the charges were applied, but not necessarily in the same pair of avalanche sizes produced. On the basis of the proof of Abelianness, at least in deterministic models, this limitation is alleviated by the insight that the sum of the avalanche sizes is invariant under a change of the order in which the model is charged.

As for the second reason, many models come with a detailed prescription of the microscopic updating procedure and therefore the microscopic time scale. Strictly, the invariance under a change of order of updates on the microscopic time scale thus applies to different models. The situation corresponds to equating different dynamics in the Ising model: For some observables, Glauber dynamics is different from Heat Bath dynamics, yet both certainly produce the same critical behaviour. In fact, choosing different dynamics (and thereby possibly introducing new conserved symmetries) can lead to different dynamical critical behaviour.

Revisiting the proof of Abelianness, however, generally reveals that the caveats above are overcautious. The very proof of Abelianness on the macroscopic time scale uses and develops a notion of Abelianness on the microscopic time scale. This connection can be made more formally, once it has been established that any configuration, quiescent or not, can be expressed by applying a suitable number of external charges on each site of an empty lattice.

Abelianness generally plays a major rôle in the analytical treatment of SOC models, because it allows significant algebraic simplifications, not least when the dynamics of a model is written in terms of Markov matrices. It applies, generally, equally to recurrent and transient states, where no inverse exists. It remains highly desirable to demonstrate Abelianness on the basis of the algebra, once that is established as a suitable representation of a model's dynamics.

In the following section a few paradigmatic models of SOC are introduced: The BTW Model, the Manna Model, the OFC Model and the Forest Fire Model.

7.1.2.1 The BTW Model

The BTW Model was introduced together with the very concept of SOC (Bak *et al.* 1987), initially to explain the “ubiquity” of $1/f$ noise. Of course, since

then, SOC has been studied very much in its own right. Like virtually all SOC models, the BTW Model consists of a set of rules that prescribe how a local degree of freedom z_i on a d -dimensional lattice with sites i is to be updated. There are two different stages, namely the relaxation and the driving, the latter considered to be slow compared to the relaxation, *i.e.* the relaxation generally is instantaneous and never occurs simultaneously with the driving (separation of time scales). In the Abelian version of the BTW Model (Dhar 1990), the driving consists of adding a single slope unit (Kadanoff *et al.* 1989) to a site, that is normally picked uniformly and at random. The lattice is often initialised with $z_i = 0$ for all i .

If the driving leads to any of the z_i exceeding the the critical slope z^c (also referred to as the critical height or threshold, depending on the view) at a site i a **toppling** occurs whereby z_i is reduced by the coordination number q of the site and z_j of every nearest neighbour j increases by one (sometimes referred to as **charging**). In principle both q and z^c can vary from site to site and such generalisations are trivial to implement. It is common to choose $z^c = q - 1$.

The rules of the BTW Model can be summarised as follows:

Initialisation: All sites i are empty, $z_i = 0$.

Driving: One unit is added at a randomly chosen (or sometimes fixed) site i , *i.e.* $z_i \rightarrow z_i + 1$.

Toppling: A site with $z_i > z^c = q - 1$ (called *active*) distributes one unit to the q nearest neighbouring sites j , so that $z_i \rightarrow z_i - q$ and $z_j \rightarrow z_j + 1$.

Dissipation: Units are lost at boundaries, where toppling site i loses q units, $z_i \rightarrow z_i - q$, yet less than q nearest neighbours exist, which receive a unit.

Time progression: Time progresses by one unit per parallel update, when all active sites are updated at once.

A toppling can trigger an avalanche, as charged neighbours might exceed the threshold in turn, possibly by more than one unit. Strictly, the BTW Model is updated in parallel, all sites topple at once whose local degree of freedom exceeds the threshold at the beginning of a time step. Microscopic time then advances by one unit. This way, z_i might increase far beyond z^c before toppling itself. As long as $z_i > z^c$ for any site i , the sites in the model carry on toppling. The totality of the toppling events is an avalanche. In the Abelian BTW model as refined by Dhar (1990), the final state of the model does not depend on the order in which external charges are applied. In the process of the proof of this property, it turns out that the order of processing any charges during the course of an avalanche neither affects the final state nor the size of the avalanche triggered. Using a parallel updating scheme or not therefore does not change the avalanche sizes recorded. As the order of updates defines the microscopic time scale, a change in the updating procedure, however, affects all observables dependent on that time, such as avalanche duration or correlations on the fast time scale.

To keep the prescription above consistent with the notion of boundary sites, where toppling particles are to be lost to the outside, boundary sites have to be thought of as having the same number of nearest neighbours as any other, equivalent site in the bulk, except that some of their neighbours are not capable of toppling themselves. For numerical purposes it is often advisable to embed a lattice in some “padding” (a neighbourhood’s “halo”, see Sec. 7.3.2.2, p. 292), *i.e.* sites that cannot topple but are otherwise identical to all other sites.

The sum of the slope units residing on a given site i and those residing on its nearest neighbours remains unchanged by the toppling of site i , *i.e.* the bulk dynamics in the BTW are conservative. Dissipation occurs exclusively at the boundary and every slope unit added to the system in the bulk must be transported to the boundary in order to leave the system.

The original version of the BTW model is defined in terms of local heights, so that the height differences give rise to the slope z_i , which has to reach q in order to trigger an a toppling. While this is a perfectly isomorphic view of the BTW, *driving* it in terms of height units has a number of unwanted implications. In particular, it loses its Abelianness. For that reason, the original version of the BTW is rarely studied numerically nowadays.

The BTW Model is **deterministic** apart from the driving, which can be made deterministic as well, simply by fixing the site that receives the external charge that triggers the next avalanche. Even when slope units do not move independently at toppling, a randomly chosen slope unit being transported through a BTW system describes the trajectory of a random walker trajectories (Dhar 1990), essentially because every possible path is being realised (just not independently, but all with the correct weight). As a result, the average avalanche size $\langle s \rangle$ can be calculated exactly; The number of moves a slope unit makes on average from the time of being added by the external drive to the time it leaves the system through an open boundary is equal to the expected number of charges it causes. The expected number of charges (caused by the movement of all slope units taking part in an avalanche) per slope unit added is thus exactly equal to the expected number of moves a slope unit makes until it leaves the system, *i.e.* its escape time. If the avalanche size is measured by the number of topplings, which is more common, the expected number of moves has to be divided by the number of moves per toppling, q in the present case. Higher moments of the avalanche size, or, say, the avalanche size conditional to non-zero size (*i.e.* at least one site toppling in every avalanche), cannot be determined using the random walker approach, as they are crucially dependent on the interaction of toppling sites.

Due to the random walker property of the slope units added, the scaling of the average avalanche size thus merely depends on the particularities of the driving. If the driving is random and uniform, then $\langle s \rangle \propto L^2$ for any d -dimensional hypercubic lattice and (Ruelle and Sen 1992)

$$\langle s \rangle = \frac{1}{12}(L+1)(L+2) \quad (7.1)$$

in one dimension with two open boundaries, where the avalanche size is the number of topplings per particle added. However, the dynamics of the BTW Model in one dimension is trivial, so that the model is usually studied only in $d = 2$ and beyond.

Because (or despite of) its deterministic nature, a large number of analytical results are known, in one dimension (Ruelle and Sen 1992) but more importantly in two dimensions (Majumdar and Dhar 1992), not least on the basis of (logarithmic) conformal field theory (e.g. Majumdar and Dhar 1992, Ivashkevich 1994, Mahieu and Ruelle 2001, Ruelle 2002, Jeng 2005). Unfortunately, to this day, the scaling of the avalanche size distribution in dimensions $d \geq 2$ remains somewhat unclear. Numerically, results are inconclusive, as different authors quote widely varying results for $d = 2$ (Vespignani and Zapperi 1995, Chessa *et al.* 1999a, Lin and Hu 2002, Bonachela 2008, e.g.), possibly due to logarithmic corrections (Manna 1990, Lübeck and Usadel 1997, Lübeck 2000)

A major insight into the *collective* dynamics of toppling sites was the decomposition of avalanches into **waves** (Ivashkevich *et al.* 1994), which was later used by Priezzhev *et al.* (1996) to conjecture $\tau = 6/5$ for the avalanche size exponent in two dimensions. No site in an avalanche can topple more often than the site at which the avalanche was triggered. Not allowing that first site to topple therefore stops the avalanche from progressing any further and each toppling of the first site thus defines a wave of toppling.

While the BTW Model has been crucial for the formation of the field of SOC as a whole, its poor convergence beyond one dimension has made it fall in popularity. One may argue that the determinism of the dynamics is to blame, as found in other models (Middleton and Tang 1995). Indeed, adding some stochasticity makes the BTW Model display the universal behaviour of the Manna Model discussed in the next section (Černák 2002, Černák 2006).

The exponents reported for the BTW Model vary greatly. In two dimensions, the value of τ found in various studies ranges from 1 (Bak *et al.* 1987) to 1.367 (Lin and Hu 2002) and that for D from 2.50(5) (De Menech *et al.* 1998) to 2.73(2) (Chessa *et al.* 1999a). Similarly α is reported from 1.16(3) (Bonachela 2008) to 1.480(11) (Lübeck and Usadel 1997) and z from 1.02(5) (De Menech and Stella 2000) to 1.52(2) (Chessa *et al.* 1999a). Using comparatively large system sizes, Dorn *et al.* (2001) found exponents that seem to vary systematically with the system size with little or no chance to identify an asymptotic value.

The first exactly solved SOC model was the Dhar-Ramaswamy Model (Dhar and Ramaswamy 1989) which is the **directed** variant of the BTW Model. The directedness means that during an individual avalanche, sites are never re-visited, which effectively suppresses spatial correlations. Random drive of the model results in a product state, where sites taking part in an avalanche form a “compact” patch (*i.e.* they have no holes), which is delimited by boundaries describing a random walk. The exponents in $d = d_{\perp} + 1$ dimensions are given analytically by $D = 1 + d_{\perp}/2$, $D(2 - \tau) = 1$, $z = 1$

and $D(\tau - 1) = z(\alpha - 1)$, which implies $\alpha = D$ and $\tau = 2 - 1/D$ (Dhar and Ramaswamy 1989, Christensen 1992, Christensen and Olami 1993, Tadić and Dhar 1997, Kloster *et al.* 2001). For example, in $d = 1+1$ dimensions (directed square lattice), exponents are $D = 3/2$, $\tau = 4/3$, $z = 1$ and $\alpha = 3/2$. Mean-field exponents apply at $d = 2 + 1$ and above.

7.1.2.2 The Manna Model

The Manna (1991) Model was originally intended as a simplified version of the BTW Model but has since then acquired the status of the paradigmatic representative of the largest (and maybe the only) universality class in SOC, generally referred to as the Manna, Oslo (Christensen *et al.* 1996) or C-DP (conserved directed percolation, Rossi *et al.* 2000) universality class.

The Manna Model displays robust, clean critical behaviour in any dimension $d \geq 1$, characterised by non-trivial exponents below $d = 4$ (Lübeck and Heger 2003b). Originally, it is defined as follows: The external drive adds particles at random chosen sites i , *i.e.* the local degree of freedom increases by one, $z_i \rightarrow z_i + 1$. If a site exceeds the threshold of $z^c = 1$ it topples, so that *all* its particles are redistributed to the nearest neighbours, which are chosen independently at random. After the toppling of site i , the local degree of freedom is therefore set to $z_i = 0$, while the total increase of the z_j at the nearest neighbours j of i maintains conservation. Again, as in the BTW model, non-conservation at boundary sites can be thought of as been implemented by sites that never topple themselves.

Charging neighbours might push their local degree of freedom beyond the threshold and they might therefore topple in turn. When a site topples, all particles present there at the time of toppling are transferred to its neighbour (maybe to a single one) and it is therefore crucial to maintain the order of (parallel) updates. The model is thus non-Abelian. In fact, the notion of Abelianness was initially restricted to deterministic models (Milshtein *et al.* 1998). However, Dhar (1999a) introduced a version of the Manna Model which is Abelian in the sense that the statistics of the final state remains unchanged if two consecutive external charges (by the driving) are carried out in reverse order. In that version of the Manna Model, a toppling site redistributes only 2 of its particles, *i.e.* the number of particles redistributed at a toppling does not depend on z_i itself. The difference between the BTW Model and the Manna Model lies thus merely in the fact that only two particles are re-distributed when a site topples in the Manna Model (irrespective of the coordination number of the site) and that the receiving sites are picked at random.

In summary, the rules of the Abelian Manna Model are:

Initialisation: All sites i are empty, $z_i = 0$.

Driving: One unit is added at a randomly chosen (or sometimes fixed) site i , *i.e.* $z_i \rightarrow z_i + 1$.

Toppling: A site with $z_i > z^c = 1$ (called *active*) distributes one unit to 2 randomly and independently chosen nearest neighbouring sites j , so that $z_i \rightarrow z_i - 2$ and $z_j \rightarrow z_j + 1$.

Dissipation: Units are lost at boundaries, where the randomly chosen nearest neighbour might be outside the system.

Time progression: Originally, time progresses by one unit per parallel update, when all active sites are updated at once.

That the scaling in one dimension is not as clean as in higher dimension may be caused by logarithmic corrections (Dickman and Campelo 2003). Nevertheless, it has been possible to extract consistent estimates for exponents in dimensions $d = 1$ to $d = 5$ (Lübeck and Heger 2003b, Huynh *et al.* 2011, Huynh and Pruessner 2012). Because some of its exponents are so similar to that of the directed percolation universality class (Janssen 1981, Grassberger 1982, Hinrichsen 2000) there remains some doubt whether the Manna Model really represents a universality class in its own right (Muñoz *et al.* 1999, Dickman *et al.* 2002). The problem is more pressing in the **fixed energy** version (Dickman *et al.* 1998, Vespignani *et al.* 1998) of the Manna Model (Basu *et al.* 2012), where dissipation at boundaries is switched off by closing them periodically, thereby studying the model at a fixed amount of particles. The term “fixed energy sandpile” was coined to stress the conserved nature of the relevant degree of freedom (which may be called “energy”) and to suggest a similar distinction as in the change of ensemble from canonical to micro-canonical. Bonachela and Muñoz (2007) suggested to study the model with different boundary conditions which have an impact on the Manna Model that is distinctly different from that on models in the directed percolation universality class.

Because of its fixed energy version, the Manna Model is frequently studied for its links to absorbing state (AS) phase transitions (Dickman *et al.* 1998, Vespignani *et al.* 1998, Hinrichsen 2000, Henkel *et al.* 2008). In fact, it has been suggested that SOC is due to the self-organisation to the critical point of such an AS phase transition (Tang and Bak 1988, Dickman *et al.* 1998, Vespignani *et al.* 1998), whereby strong activity leads to a reduction of particles by dissipation, making the system in-active, while quiescence leads to activity due to the external drive. One may argue that such a linear mechanism cannot produce the desired universal critical behaviour without finely tuning the relevant parameters (Pruessner and Peters 2006, 2008, Alava *et al.* 2008).

A number of theoretical results are available for the Manna Model (Vespignani *et al.* 1998, 2000, Rossi *et al.* 2000, van Wijland 2002, Ramasco *et al.* 2004), yet an ϵ -expansion (Le Doussal *et al.* 2002) for the Manna universality class is available only via the mapping (Paczuski and Boettcher 1996, Pruessner 2003) of the Oslo Model (Christensen *et al.* 1996), which is the same universality class (Nakanishi and Sneppen 1997) as the Manna Model, to the quenched Edwards-Wilkinson equation (Bruinsma and Aeppli 1984, Koplik and Levine 1985, Nattermann *et al.* 1992, Leschhorn *et al.* 1997). Quenched

noise and disorder are, however, notoriously difficult to handle analytically. It is thus highly desirable to develop a better theoretical understanding of the Manna Model in its own right, including its mechanism of self-organisation, and to derive an ϵ -expansion for its exponents.

Although the Manna Model is more frequently studied in one dimension, for comparison with the BTW Model above, the exponents listed in the following were determined numerically in two dimensions for the Abelian and the non-Abelian (original) variant of the Manna Model. For τ they range from 1.25(2) (Biham *et al.* 2001) to 1.28(2) (Manna 1991, Lübeck and Heger 2003a), for D from 2.54 (Ben-Hur and Biham 1996) to 2.764(10) (Lübeck 2000), for α from 1.47(10) (Manna 1991) to 1.50(3) (Chessa *et al.* 1999b, Lübeck and Heger 2003a) and for z from 1.49 (Ben-Hur and Biham 1996) to 1.57(4) (Alava and Muñoz 2002, Dickman *et al.* 2002), generally much more consistent than in the BTW Model.

As in the BTW Model, various directed variants of the Manna Model which are exactly solvable for similar reasons as in the deterministic case have been extensively studied (Pastor-Satorras and Vespignani 2000b,a, Hughes and Paczuski 2002, Pan *et al.* 2005, Jo and Ha 2008). They have been characterised in detail by Paczuski and Bassler (2000b) and related to the deterministic directed models by Bunzarova (2010). Exponents generally follow $D = 3/2 + d_{\perp}/4$, which can be interpreted as the diffusive exploration of a random environment. Again, correlations are suppressed as sites are never re-visited in the same avalanche. As in the deterministic case, $z = 1$ and $D(2 - \tau) = 1$ and $D(\tau - 1) = z(\alpha - 1)$ result in $D = \alpha$. In $d = 1 + 1$ exponents are $\tau = 10/7$, $D = 7/4$, $\alpha = 7/4$ and $z = 1$.

7.1.2.3 The Forest Fire Model

The Forest Fire Model has an interesting, slightly convoluted history. Two distinct versions exist, which share the crucial feature that the bulk dynamics is not conservative. In the original version introduced by Bak *et al.* (1990) sites i , most frequently organised in a (two-dimensional) square lattice with periodic boundary conditions, can be in one of three states $\sigma_i \in \{T, F, A\}$, corresponding to occupation by a **T**ree, by **F**ire or by **A**sh. As time t advances in discrete steps, the state changes cyclically under certain conditions: A **T**ree turns into **F**ire at time $t + 1$ if a nearest neighbouring site was on **F**ire at time t . In turn, a **F**ire at time t becomes **A**sh in time $t + 1$, and a site covered in **A**sh at time t might become occupied by a **T**ree at time $t + 1$ due to a repeated Bernoulli trial with (small) probability p . Starting from a lattice covered in trees, a single site is set on fire and the system evolves under the rules described. The key observable is the number of sites on fire as a function of time.

Initialisation: All (many) sites i contain a tree (otherwise ash), $\sigma_i = T$, and (at least) one site is on fire, $\sigma_i = F$.




Fig. 7.1: Realisation of the original Forest Fire Model by Bak *et al.* (1990). Ash is marked by a white site, Trees are black and Fires grey.

Driving: With (small) probability p , a site i containing ash at the beginning of time step t contains a tree, $\sigma_i = A \rightarrow T$ at time $t + 1$.

Toppling: A site i that contains a tree at beginning of time step t and has at least one nearest neighbour on fire, turns into fire as well, $\sigma_i = T \rightarrow F$. Simultaneously, a site on fire at t turns into ash, $\sigma_i = F \rightarrow A$.

Dissipation: trees grow slowly in Bernoulli trials and are removed in the “toppling”. Their number is not conserved under any of the updating.

Time progression: Time progresses by one unit per parallel update.

The original Forest Fire Model (FFM) just described possesses an **absorbing state** from which it cannot recover within the rules given. If the fire stops spreading because the last site on fire is surrounded by ash, the only transition that can and will take place eventually occupies every site by a tree. Bak *et al.* (1990) originally suggested that occasional re-lightning might be necessary — in fact, if p is large enough, on sufficiently large lattices, there will always be tree to burn available. This, however, points to a fundamental shortcoming, as quantified by Grassberger and Kantz (1991), namely that the lengthscale of the relevant features of the FFM are determined by p . Typically, at small p , some large spiral(s) of fire keeps sweeping across the lattice. If p is chosen too small, the spatial extent of the spiral becomes too large compared to the size of the lattice and the fire eventually goes out. However, if a control parameter determines the characteristic length scale of the phenomenon, it cannot be *bona fide* SOC (e.g. Bonachela and Muñoz 2009). Figure 7.1 shows an example of the structures, most noticeable the fire fronts, developing.

The name “Forest Fire Model” should be taken as a witty aide-memoire. Bak *et al.* (1990) designed the model to understand scale free dissipation with uniform driving as observed in turbulent flow. The model should therefore be considered much more as a model of turbulence that happened to look like fires spreading in a forest. In the present model, perpetual fires spread across trees as they re-grow, which is a rather unrealistic picture; most fires in real forests are shaped by fire brigades, geographical and geological features and other environmental characteristics, as well as policies. Nevertheless, the original FFM as well as the version by Drossel and Schwabl (1992a), attracted significant attention as an actual model of forest fires, as well as other natural and sociological phenomena (Turcotte 1999).

There are two distinguishing features that set the FFM apart from many other SOC models. Firstly, the separation of time scales is incomplete, because driving the system by supplying new trees is a process running in parallel to the burning as fire spreads. Although the time scale of tree growth, parameterised by p , can in principle be made arbitrarily slow, the fire has to be constantly fed by new trees and cannot be allowed to go out, because there is no explicit re-lighting. In other words, the tree growth rates that still sustain fire are bounded from below. As a result, there are no distinct avalanches, as found in the BTW and the Manna Models.

More importantly, however, the FFM is different from other models because it is non-conservative at a fundamental level. No quantity is being transported to the boundaries and the local degree of freedom changes without any conservation.⁶ At the time of the introduction of the FFM, it challenged Hwa and Kardar’s (1989a) suggested mechanism of SOC that relied on a conservation law to explain the absence of a field-theoretic mass in the propagator.

Other dissipative models, like the SOC version of the “Game of Life” (Bak *et al.* 1989a), the OFC model discussed in the next section (Olami *et al.* 1992) and the Bak-Sneppen Model (Bak and Sneppen 1993) chipped away from the conservation argument put forward by Hwa and Kardar (1989a, 1992), Grinstein *et al.* (1990) and Socolar *et al.* (1993). The latter seem to have been caught by surprise by the advent of a variant of the FFM by Drossel and Schwabl (1992a) discussed in the following.

The Drossel-Schwabl Forest Fire Model (DS-FFM), as it is now normally referred to, was originally introduced by Henley (1989). It changes the original Forest Fire Model in two very important points: Firstly, the separation of time scales between burning and growing is completed, so that patches of (nearest neighbouring) trees are burned down instantly compared to all other

⁶ It is difficult to make the statement about non-conservation strict. After all, the state of each site is meant to change and allowing for that, it is always possible to trace the appearance and the disappearance of something back to some influxes and outfluxes. Here is an attempt in the present case: While the increase in the number of trees can be thought of as being due to a corresponding influx, they can disappear with an enormous rate by spreading fire without explicit outflux *on that timescale*.




Fig. 7.2: Realisation of the Drossel-Schwabl Forest Fire Model (Drossel and Schwabl 1992a). Ash is marked by a white site, Trees are black.

processes. Because fires therefore burn down completely before anything else can happen, fires are set, secondly, explicitly by random, independent uniform lightning. The key-observables of the DS-FFM are the geometrical features of the clusters burned down, such as the number of occupied sites (the mass) and the radius of gyration.

While trees grow with rate p on every empty site (*i.e.* one containing ash), lightning strikes with much lower rate f on every site. If it contains a tree, the fire eradicates the entire cluster of trees connected to it by nearest neighbour interactions. In summary:

Initialisation: All sites i contain ash, $\sigma_i = A$.

Driving: With (small) probability p , a site i containing ash at the beginning of time step t contains a tree, $\sigma_i = A \rightarrow T$ at time $t + 1$.

Toppling: With probability $f \ll p$, a site containing a tree at the beginning of time step t and the entire cluster of trees connected to it by nearest neighbour interactions is changed to ash, $\sigma_i = T \rightarrow A$.

Dissipation: trees grow slowly in Bernoulli trials and are removed in the “toppling”. They are not conserved in any of the updates.

Time progression: Time progresses by one unit per parallel update, toppling is instantaneous relative to growing trees.

As a result entire patches of forest disappear at a time, which are reforested with the same Poissonian density p . This process results in a patchy structure with individual islands having roughly homogeneous tree-density, Figure 7.2.

In a change of perspective, the processes parameterised by p and f are tree growth attempts and lightning attempts which fail if the site is already occupied by a tree or does not contain one, respectively. The original definition by Drossel and Schwabl (1992a) still used discrete time, so that both p and f were probabilities, rather than Poissonian rates, which can be recovered

by rescaling p and f simultaneously. However, it is common (e.g. Clar *et al.* 1996) to rescale time so that $p = 1$ (enforced growth on randomly picked empty sites) and to attempt p/f times to grow a tree before attempting to set one alight. In order to see scale-free cluster size distributions, a **second separation of timescales** is needed, whereby the ratio p/f diverges.

Many of the properties of the DS-FFM are percolation-like. If it were not for the correlations in the tree-density, which develop because of “synchronous, patchy re-forestation”, i.e. if the tree-density was homogeneous, then the DS-FFM would be a form of percolation. In particular, the cluster size distribution (of the patches removed and the totality of all patches present) was given by that of (well-known) static percolation.

The DS-FFM does not suffer from the same short-coming as the original FFM of having a well-understood typical (spiral) structure, whose size is determined by the single control parameter p , yet it still has one control parameter which needs to be finely tuned in accordance with the system size. This parameter is p/f — if it is too large, then the lattice will be densely filled with trees before lightning strikes and removes almost all of them, leaving behind essentially a clean sheet with a few remaining (small) islands of trees. If p/f is too small, then no dense forest ever comes into existence and the cluster size distribution has a cutoff not determined by the system size, but by that parameter.

In extensive numerical studies (Grassberger 2002, Pruessner and Jensen 2002a, 2004), the system sizes were chosen big enough for each p/f that finite size effects were not visible, *i.e.* for each p/f convergence of the cluster size distribution $\mathcal{P}(s; L)$ in the system size L was achieved. However, these studies revealed that the DS-FFM does not display simple scaling in $s_c = s_c(p/f)$, Eq. (7.3) (Sec. 7.2.1). While $\mathcal{P}(s)/s^{-\tau}$ converges in the thermodynamic limit (as it should, trivially) for any τ , there is no choice of τ so that the remaining functional profile depends only on the ratio $s/s_c(p/f)$. Instead, $\mathcal{P}(s)/s^{-\tau}$ depends explicitly on both s and $s_c(p/f)$, or, for that matter, p/f . The only feature that may display some convergence (Pruessner and Jensen 2002a) is the bump in the probability density function (PDF) towards large s . For some choice of τ , there is a small region, say $[s_c(p/f)/2, s_c(p/f)]$, where $\mathcal{P}(s)/s^{-\tau}$ traces out a very similar graph, as if the lower cutoff s_0 itself was a divergent multiple of the upper cutoff.⁷

One may hope that finite size scaling can be recovered, taking the limit of large p/f and considering $\mathcal{P}(s)/s^{-\tau}$ as a function of L . However, it is clear that the PDF trivialises in this limit,

$$\lim_{p/f \rightarrow \infty} \mathcal{P}(s; p/f, L) = s^{-1} \delta\left(\frac{s}{L^d}\right) \quad (7.2)$$

⁷ If $s_c(p/f)$ marks roughly the maximum of the bump, the PDF drops off beyond it so quickly, that next to nothing is known of $\mathcal{P}(s)$ beyond s_c . In principle, however, if there is approximate coincidence on $[s_c(p/f)/2, s_c(p/f)]$, there should also be approximate coincidence on $[s_c(p/f)/2, \infty)$.

as the lattice is completely covered in trees before they all get completely removed in a singly lightning.

Interestingly, the lack of scaling in finite $s_c(p/f)$ is not visible in the scaling of the moments $\langle s^n \rangle$ because they are sensitive to large event sizes (at any fixed $n > \tau - 1$), rather than the smaller ones around the lower cutoff, whose divergence violates simple scaling.

As in the BTW Model, exponents reported for the DS-FFM (if they are reported at all) display a fairly wide spread. In two dimensions, they are τ from 1 (Drossel and Schwabl 1992a) to 1.48 (Patzlaff and Trimper 1994) and D from 1 (Drossel and Schwabl 1992a) to 1.17(2) (Henley 1993, Honecker and Peschel 1997).

7.1.2.4 The OFC Model

To this day, the Olami-Feder-Christensen Model (OFC Model Olami *et al.* 1992) is one of the most popular and spectacular models of SOC. It is a simplified version of the Burridge-Knopoff Model (Burridge and Knopoff 1967) of earthquakes, it has a *tunable* degree of non-conservation (including a conservative limit) with a clear physical meaning, it has been extensively analysed, both in time and space, for the effect of different boundary conditions (Middleton and Tang 1995), and its one-dimensional variant (de Sousa Vieira 1992) has been linked to the Manna universality class (Paczuski and Boettcher 1996, Chianca *et al.* 2009). After the definition of the model, the discussion below focuses on the model's rôle in earthquake modelling and the attention it received for the spatio-temporal patterns it develops.

The OFC Model is at home on a two-dimensional square lattice. As in the models above, each site i has a local degree of freedom $z_i \in \mathbb{R}$ (called the local “pulling force”), which is, in contrast to the models above, however, real-valued. As in the BTW Model, there are two clearly distinct stages of external driving and internal relaxation. During the driving *all* sites in the system receive the same amount of force (sometimes referred to as “continuous” or better “uniform” drive) until one site exceeds the threshold $z^c = 1$, which triggers a relaxation during which no further external driving is applied. In a relaxation or toppling, a site re-distributes a fraction of *all* pulling force evenly among its nearest neighbours which may in turn exceed the threshold. The force z_i at the toppling site i is set to 0 and the amount arriving at each neighbour is αz_i , where α is the **level of conservation**. At coordination number q , a level conservation less than $1/q$ means that the bulk dynamics is dissipative. Boundary sites lose force αz_i (at corners multiples thereof) to the outside. Because the force re-distributed depends on the amount of pulling force present at the site at the time of the re-distribution, the order of updates matters greatly, *i.e.* the OFC Model is not Abelian. If $\alpha < 1/q$ periodic boundary conditions can be applied without losing the possibility of a stationary state, yet normally the boundaries are open. The OFC Model is

normally initialised by assigning random, independent forces from a uniform distribution.

Sites to topple are identified at the beginning of a timestep and only those have been relaxed by the end of it (parallel updates). Unless more than one site exceeds the threshold (degenerate maximum) at the beginning of an avalanche, toppling sites therefore reside on either of the two next nearest neighbour sublattices of a square lattice.

Again, a separation of time scales is applied, where the relaxation becomes infinitely fast compared to an infinitesimally slow drive. In an actual implementation, however, the driving is applied instantaneously and the relaxation takes up most (computational time): The driving can be completed in a single step by keeping track of the site, say i^* with the largest pulling force acting on it. The amount of force added throughout the system is thus simply $z^c - z_{i^*}$, triggering the next avalanche.

Because sweeping the lattice in search of the maximum is computationally very costly,⁸ the main computational task in the OFC Model is to keep track of the site exposed to the maximum pulling force. This is a classic computational problem (Cormen *et al.* 1996), which also occurs in other models, such as the Bak-Sneppen Model (Bak and Sneppen 1993). The traditional solution is to organise data in a tree-like structure and devise methods that allow fast updating and determination of the maximum. However, in the OFC Model updating a site's force is much more frequent than determination of the maximum and thus a fast algorithm focuses on the optimisation of the former at the expense of the latter, *i.e.* a slightly slower procedure to determine the maximum.

Grassberger (1994) pointed out a number of improvements to a naïve, direct implementation of the OFC Model. Firstly, instead of driving the system uniformly, thereby having to sweep the lattice to increase the force on every site by $z^c - z_{i^*}$, the threshold z^c is to be lowered; the amount of force re-distributed at toppling is obviously to be adjusted according to the new offset. The second major improvement Grassberger (1994) suggested was the organisation of forces in “boxes” (sometimes referred as **Grassberger's box-technique**), which splits the range of forces present in the system in small enough intervals that the search for the maximum force succeeds very quickly, yet keeps the computational effort to a minimum when re-assigning a box after an update. Other improvements suggested was maintaining a stack (Sec. 7.3.1) of active sites, and the use of a scheme to determine neighbouring sites suitable to the programming language at hand.

The adjustment of z^c outlined above has some rather unexpected effects depending on the numerical precision (Sec. 7.3.3) used in the simulation (Pruessner 2012b). As pointed out by Drossel (2002), the OFC Model is ex-

⁸ Not only is the very searching *across all sites* costly, most of the memory occupied by the lattice will not reside in a cache line (as for example most “local” data) and thus has to be fetched through a comparatively slow bus.

tremely sensitive to a change of precision; a lower precision seems to enhance or favour phase-locking, discussed in the following.

Most of the studies of SOC models focuses on large-scale statistical features, large both in time and space. The analysis of the OFC Model by Socolar *et al.* (1993) Middleton and Tang (1995) and Grassberger (1995) therefore ventured into uncharted territory as they studied the evolution towards stationarity in the OFC Model on a microscopic scale, analysing the patchy structure of the forces on the lattice.

Firstly, periodic boundary conditions inevitably lead to periodic behaviour in time. Below $\alpha \approx 0.18$ in a two-dimensional square lattice, (almost) every avalanche has size unity. In that extreme case, the period is strictly $1 - q\alpha$, because discounting the external drive, this is the amount of force lost from every site after every site has toppled exactly once, as the system goes through one full period.

The periodicity is broken once open boundaries are introduced. Sites at the edge of the lattice have fewer neighbours that charge them, so if every site in the system topples precisely once, the force acting on a boundary site is expected to be lower. While open boundaries indeed break temporal periodicity, they form, at the same time, seeds for (partially) synchronised patches, which seem to grow from the outside towards the inside, increasing in size. Middleton and Tang (1995) introduced the term **marginal (phase) locking** to describe this phenomenon.

The temporal periodicity might similarly be broken by introducing inhomogeneities or disorder, effective even at very low levels (Grassberger 1994, Ceva 1995, 1998, Torvund and Frøyland 1995, Middleton and Tang 1995, Mousseau 1996). That a spatial inhomogeneity helps forming synchronised patches in space can also be attributed to marginal phase locking.

Because the OFC Model is so sensitive to even the smallest amount of disorder and inhomogeneity, its statistics is often taken from very big samples with extremely long transients. Many authors also average over the initial state. Drossel (2002) suggested that despite these precautions, some of the statistical behaviour allegedly displayed by the OFC Model might rather be caused by numerical “noise”, also a form of inhomogeneity or disorder entering into a simulation. In practise, it is difficult to discriminate genuine OFC behaviour from numerical shortcomings and one may wonder whether some of these shortcomings are not also present in the natural phenomenon the OFC Model is based on.

That SOC may be applicable in seismology had been suggested by Bak *et al.* (1989b, also Bak and Tang 1989, Sornette and Sornette 1989, Ito and Matsuzaki 1990) at a very early stage. The breakthrough came with the OFC Model, which is based on the Burridge-Knopoff Model of earthquakes (or rather fracturing rocks). The latter is more difficult to handle numerically, with a “proper” equation of motion taking care of the loading due to spring-like interaction much more carefully. The OFC Model, on the other hand,

is much easier to update, almost like a cellular automaton.⁹ The context of SOC provided an explanatory framework of the scale-free occurrence of earthquakes as described by the **Gutenberg-Richter** law (Gutenberg and Richter 1954, Olami *et al.* 1992). Even though exponents both in the real-world as well as in the OFC Model seem to lack universality, certain scaling concepts, motivated by studies in SOC, have been applied successfully to earthquake catalogues (Bak *et al.* 2002).

It is fair to say that the OFC Model, to this day, is widely disputed as a *bona fide* model of earthquakes. Its introduction has divided the seismology community, possibly because of the apparent disregard of their achievements by the proponents of SOC (Bak and Tang 1989). One of the central claims made initially is that earthquakes are unpredictable if they are “caused” by SOC, which questions the very merit of seismology. However, given that SOC is a framework for the understanding of natural phenomena on a long time and length scale, providing a mechanism for the existence of long temporal correlations, SOC indicates precisely the opposite of unpredictability. This point is discussed controversially in the literature to this day (Corral 2003, 2004c,b, Davidsen and Paczuski 2005, Lindman *et al.* 2005, Corral and Christensen 2006, Lindman *et al.* 2006, Werner and Sornette 2007, Davidsen and Paczuski 2007, Sornette and Werner 2009). Older reviews (Turcotte 1993, Carlson *et al.* 1994) help to understand the historical development of the dispute. Hergarten (2002) and more recently Sornette and Werner (2009) have put some of the issues in perspective.

There is not a single set of exponents for the OFC Model, as they are generally expected to vary with the level of conservation (Christensen and Olami 1992). Because authors generally do not agree on the precise value of α to focus on, results are not easily comparable across studies. Even in the conservative limit, $\alpha = 1/q$, little data is available, suggesting $\tau = 1.22(5) - 1.253$ and $D = 3.3(1) - 3.01$ (Christensen and Olami 1992, Christensen and Moloney 2005).

7.2 Scaling and numerics

As a rule of thumb, SOC models are **SDIDT** systems Jensen (1998): **S**lowly **D**riven **I**nteraction **D**ominated **T**hreshold systems. The driving implements a separation of time scales and thresholds lead to highly non-linear interaction, which results in avalanche-like dynamics, the statistics of which displays scaling, a continuous symmetry. Ideally, the scaling behaviour of an SOC model can be related to some underlying continuous phase transition, which is triggered by the system self-organising to the critical point.

⁹ Strictly, the OFC Model generally is not a cellular automaton, because the local states z_i are continuous.

The critical behaviour can be characterised by (supposedly) universal critical exponents, the determination of which is the central theme of the present section. At the time of the conception of SOC, critical exponents were extracted directly from probability density function, (PDFs), often by fitting the data to a straight line in double-logarithmic plot. Frequently, such scaling is referred to as “power law behaviour”. Very much to the detriment of the entire field, some authors restrict their research to the question whether an observable displays the desired behaviour, without attempting to determine its origin and without considering the consequences of such behaviour. Power law behaviour therefore has become, in some areas, a mere curiosity.

7.2.1 Simple scaling

While studying power laws in PDFs can be instructive, there are far superior methods to quantify scaling behaviour. In recent years, most authors have focused on an analysis of the moments of the PDFs, as traditionally done in the study of equilibrium statistical mechanics. Not only is this approach more efficient, it also is more accurate and mathematically better controlled. Moreover, it is concerned directly with an observable (or rather, arithmetic means of its powers), rather than its accumulated histogram.

Nevertheless the starting point of a scaling analysis in SOC, is the **simple (finite size) scaling assumption**,

$$\mathcal{P}(s) = as^{-\tau}\mathcal{G}(s/s_c) \text{ for } s \gg s_0, \quad (7.3)$$

where $\mathcal{P}(s)$ is the (normalised) probability density function of an observable, s in this case, a is a (non-universal) **metric factor** present to restore dimensional consistency and accounting for the (microscopic) details of the model, τ is a **universal scaling (or critical) exponent**, \mathcal{G} is a **universal scaling function**, s_c is the **upper cutoff** and s_0 the **lower cutoff**. If s is the avalanche size, then τ is known as the **avalanche size exponent**, when s is the duration, then τ is traditionally replaced by α and called the **avalanche duration exponent**.

The two cutoffs signal the onset of new physics: Below s_0 some microscopic physics prevails, often a lattice spacing or some other minimal length below which discretisation effects take over. Above s_c some large finite length scale becomes visible, which in SOC is normally controlled by the size of the lattice, so that Eq. (7.3) is referred to as *finite size* scaling. In traditional critical phenomena, s_c is controlled by the correlation length, beyond which distant parts of the system can be thought of as being independent, suggesting the validity of the central limit theorem.

Strictly, SOC models should always tune themselves to a critical point, so that the algebraic, critical behaviour is cut off only by the system size.

All scaling in SOC therefore is finite size scaling. There are a handful of established SOC models, which violate that strict rule, however, such as the Drossel-Schwabl Forest Fire Model Drossel and Schwabl (1992a), where an additional parameter has to be tuned simultaneously with the system size.

The physical origin of the scales contained in the metric factor a and the lower cutoff s_0 often is the same, yet even with these length scales present, $\mathcal{P}(s)$ has an arbitrarily wide region where it displays a power-law dependence on s and whose width is controlled by s_c ; if $s_0 \ll s \ll s_c$, then $\mathcal{P}(s) = a s^{-\tau+\alpha} s_c^{-\alpha} \tilde{\mathcal{G}}_0$, provided

$$\lim_{x \rightarrow 0} x^{-\alpha} \mathcal{G}(x) = \tilde{\mathcal{G}}_0 . \quad (7.4)$$

Typically, however, $\alpha = 0$ so that the intermediate region of $\mathcal{P}(s)$ displays a power law dependence with exponent τ , which can in principle be extracted as the negative slope of $\mathcal{P}(s)$ in a double logarithmic plot. However, because it is *a priori* unclear whether the scaling function $\mathcal{G}(s/s_c)$ can be approximated sufficiently well by a constant $\tilde{\mathcal{G}}_0$, “measuring” the exponent τ by fitting the intermediate region of a double logarithmic plot to a straight line (sometimes referred to as the **apparent exponent**) is very unreliable. If the scaling function displays a power law dependence on the argument, $\alpha \neq 0$, the effective exponent in the intermediate region is $\tau - \alpha$. One can show that α is non-negative, $\alpha \geq 0$, and $\tau = 1$ if $\alpha > 0$ (Christensen *et al.* 2008).




Fig. 7.3: Example of a double logarithmic plot of the PDF of the avalanche size in an SOC model (Data from Pruessner 2012b).

Figure 7.3 shows a typical double-logarithmic plot of the PDF in an SOC model. The power law region is marked by two dashed lines. The lower cutoff is at around $s_0 = 50$ and the features below that value are expected to be essentially reproduced by that model irrespective of its upper cutoff. The

spiky structure visible in that region is not noise and may, to some extent, be accessible analytically, similar to the lattice animals known in percolation (Stauffer and Aharony 1994). The power law region between the two dashed lines can be widened arbitrarily far by increasing the upper cutoff s_c . Running the same model with increasing s_c will reproduce this almost straight region beyond which the bump in the data indicates the onset of the upper cutoff.

The upper cutoff in SOC models supposedly depends only on the system size and does so in a power-law fashion itself,

$$s_c(L) = bL^D \quad (7.5)$$

where b is another metric factor and D is the **avalanche dimension**. The exponent describing the same behaviour for the upper cutoff of the avalanche duration is the **dynamical exponent** z . The four exponents τ , D , α and z are those most frequently quoted as the result of a numerical study of an SOC model.

The simple scaling ansatz Eq. (7.3) as well the scaling of the upper cutoff, Eq. (7.5), both describe *asymptotic* behaviour in large s_c and L respectively. When determining exponents in computer simulations of SOC models, corrections have to be taken into account in a systematic manner. While sub-leading terms are difficult to add to the simple scaling ansatz Eq. (7.3), this is routinely done in the case of the upper cutoff,

$$s_c(L) = bL^D(1 + c_1L^{-\omega_1} + c_2L^{-\omega_2} \dots) \quad (7.6)$$

Corrections of this form are referred to as **corrections to scaling** (Wegner 1972) or confluent singularities. These corrections are discussed further in the context of moment analysis, Sec. 7.2.2.

Although some very successful methods of analysis exist (Clauset *et al.* 2009), Eq. (7.3) does not lend itself naturally to a systematic quantitative analysis for fixed s_c . Often, a **data collapse** is performed in order to demonstrate the consistency of the data with simple scaling. According to Eq. (7.3) the PDF $\mathcal{P}(s)$ for different cutoffs s_c produces the same graph by suitable rescaling, in particular by plotting $\mathcal{P}(s)s^\tau$ against $x = s/s_c$, which gives $\mathcal{G}(x)$. Deviations are expected for small values of s/s_c , namely for s around s_0 , where Eq. (7.3) does not apply. Figure 7.4 shows an example of such a collapse using the same model as in Figure 7.3.

Provided $\lim_{x \rightarrow 0} \mathcal{G}(x) = \mathcal{G}_0 \neq 0$, the region where $\mathcal{P}(s)$ displays (almost) a power law translates into a horizontal, (nearly) constant section in the rescaled plot. The graph terminates in a characteristic **bump**, where the probability density of some larger event sizes exceeds that of some large, but smaller ones. This counter-intuitive feature is normally interpreted as being caused by system spanning events which were terminated prematurely by the boundaries of the system. Had the system been larger, the events would have developed further. In the PDF of a larger system thus make up the regular, straight power law region, where the smaller system's PDF displays a bump.




Fig. 7.4: Data collapse of three different data sets similar to the data shown in Figure 7.3. The upper cutoff s_c is solely controlled by the system size L (Data from Pruessner 2012b).

Even when the total probability contained in the bump is finite but very small, it is enough to account for all events contained beyond it in the power law region of an infinite system.

A data collapse is not unique, as plotting $\mathcal{P}(s) s^\tau f(s/s_c)$ produces $\mathcal{G}(x)f(x)$ for any function $f(x)$. In the literature, $f(x)$ is often chosen as $f(x) = x^{-\tau}$ so that $\mathcal{P}(s) s^\tau f(s/s_c) = \mathcal{P}(s) s_c^\tau$. Plotting that data has the fundamental disadvantage that $\mathcal{P}(s) s_c^\tau$ usually spans many orders of magnitude more across the ordinate compared to $\mathcal{P}(s) s^\tau$, so that details in the terminal bump are less well resolved.

7.2.1.1 Binning

A clean, clear dataset like the one shown in Figure 7.3 is the result of **binning**. For numerical studies of SOC models this is a necessary procedure in order to smoothen otherwise rather rugged histograms. The reason for that ruggedness is the strong dependence of the probability density on the event size, with very few large events occurring. Because of the power law relationship between event size and frequency, their total numbers decrease even on a logarithmic scale. As a result, statistical noise visibly takes over, often clearly before the onset of the terminal bump. Statistics for large event size is sparse and often little more than a muddle of seemingly unrelated data points is visible in the raw data for large events.

The noise can be reduced by averaging the data for increasingly large event sizes over increasingly large “bins”, hence the name binning. This is

normally done in **post-processing** of the raw data produced in a numerical simulation, by summing over all events within a bin and dividing by its size. In principle, the bin sizes could be chosen to fit the data; if the bin ranges are $[b_i, b_{i+1})$, then a pure power law $\mathcal{P}(s) = as^{-\tau}$ would deposit

$$\int_{b_i}^{b_{i+1}} ds as^{-\tau} = \frac{a}{\tau - 1} (b_i^{1-\tau} - b_{i+1}^{1-\tau}) \quad (7.7)$$

events in each bin i . This number can be made constant by choosing $b_i = (B_0 - B_1 i)^{1/(1-\tau)}$. Similarly, one might chose the bin boundaries b_i “on the fly”, *i.e.* successively increase the bin size until roughly a given number of entries have been collected. While those two choices lead to uniformly low statistical errors (assuming constant correlations), they both suffer from significant shortcomings. Firstly, the exponent τ to be estimated from the data should not enter into the very preparation of the data that is meant to produce the estimate. This problem is mitigated by the fact that τ may be determined through a separate, independent procedure. Secondly and more importantly, both procedures will lead to an increasingly wide spacing of data points, which becomes unacceptable towards large event sizes, because the abscissa will no longer be defined well enough — if b_{i+1} and b_i are orders of magnitude apart, which s does Eq. (7.7) estimate. Last but not least, to make PDFs of different system sizes comparable, the same b_i should be used for all datasets.

The widely accepted method of choice is **exponential binning** (sometimes also referred to as **logarithmic binning**), where $b_i = B_0 \exp(\beta i)$. Such bins are equally spaced on the abscissa of a double logarithmic plot. Because the width of exponential bins is proportional¹⁰ to their limits, Eq. (7.7), sparse data can cause a surprising artefact, whereby single events spuriously produce a probability density which decays inversely with the event size, $\mathcal{P}(s) \propto s^{-1}$, suggesting an exponent of $\tau = 1$. A typical problem with exponential bins occurs at the small end of the range when used for integer valued event sizes, because in that case the $b_{i+1} - b_i$ should not be less than 1. It is then difficult to control the number of bins and thus the resolution effectively, because decreasing β increases the number of minimally sized bins and has highly non-linear knock-on effects on all bin boundaries. The problem is obviously much less relevant for non-integer event sizes, such as the avalanche duration. However, it is rather confusing to use non-integer bin boundaries for integer valued event sizes, because bins may remain empty and the effective bin size cannot be derived from $b_{i+1} - b_i$. For example a bin spanning $b_{i+1} - b_i = 0.9$ may not contain a single integer, whereas $b_{i+1} - b_i = 1.1$ may contain two.

It is obviously advantageous to perform as much as possible of the data manipulation as post-processing of raw simulation data. Efficiency and memory limitations, however, normally require a certain level of binning at the

¹⁰ For integer valued bin boundaries, strictly, this holds only approximately.

simulation stage. When event sizes and frequencies spread over 10 orders of magnitude a simple line of code¹¹

```
    histogram[size]++; /* one count for size in the histogram */
```

would require `histogram` to have a precision of more than 32 bits. Normally such counters are implemented as integers, which would need to be a `long long int` in the present case. The memory required for 10^{10} of these 64 bit numbers (about 75 GB) exceeds by far the memory typically available in computers in common use at the time of writing this text (2012). Writing every event size in a list, eventually to be stored in a file, is rarely an alternative, again because of the enormous memory requirements and because of the significant amount of computational time post-processing would take.

Consequently, some form of binning must take place at the time of the simulation. In principle, any sophisticated binning method as used during post-processing can be deployed within the simulation, yet the risk of coding errors ruining the final result and the computational effort renders this approach unfeasible. The established view that complicated floating point operations such as `log` or `pow` are too expensive to be used regularly in the course of a numerical simulation has experienced some revision over the last decade or so, as techniques like hyperthreading and out-of-order execution are commonly used even in the FPU. Nevertheless, integer manipulation, often doable within a single CPU cycle, remains computationally superior compared to floating point manipulation. Even some of the rather archaic rules remain valid, such as multiplications being computationally more efficient than divisions, as they can be performed within a short, fixed number of cycles. Further details can be found in the appendix at the end of the chapter.

7.2.2 Moment analysis

By far the most powerful technique to extract universal features of an SOC model is a moment analysis (De Menech *et al.* 1998). Traditionally, the numerical investigation of critical phenomena has focused on moments much more than on the underlying PDF, even when the former are often seen as the “result” of the latter. Mathematically, no such primacy exists and one can be derived from the other under rather general conditions (Feller 1966, Carleman’s theorem in). In general one expects that a finite system produces only finite event sizes, *i.e.* that finite systems have a sharp cutoff of the “largest possible event size”. While very physical, this rule finds its exception in residence times, when particles get “buried” in a “pile” over long periods. In the

¹¹ All explicit examples in this chapter are written in C, which is the most widely used programming language for numerical simulations, as long as they are not based on historic Fortran code.

Oslo Model, some of these waiting time distributions seem to be moderated by scaling functions that are themselves power laws and may possess upper cutoffs exponential in the system size (Dhar and Pradhan 2004, Pradhan and Dhar 2006, 2007, Dhar 2006).

Assuming, however, that all moments

$$\langle s^n \rangle = \int_0^\infty ds s^n \mathcal{P}(s) \quad (7.8)$$

exist, *i.e.* are finite, then for $n + 1 > \tau$

$$\langle s^n \rangle \simeq a s_c^{1+n-\tau} g_n \quad (7.9)$$

where \simeq is used to indicate equivalence *to leading order in large s_c* . Moments with $n < \tau - 1$ are not dominated by the scaling in s_c , *i.e.* they are convergent in large s_c . The (asymptotic) amplitudes g_n are defined as

$$g_n = \int_0^\infty x^{n-\tau} \mathcal{G}(x) \quad (7.10)$$

expected to be finite for all $n \geq 0$. There is an unfortunate confusion in the literature about the (spurious) consequences of $\langle s^0 \rangle = 1$ scaling like $s_c^{1-\tau} g_0$. If $\tau > 1$, then the leading order of $\langle s^0 \rangle$ is not given by Eq. (7.9).

The only scaling in SOC is finite size scaling, *i.e.* the upper cutoff is expected to diverge with the system size, Eq. (7.5), so that moments scale like

$$\langle s^n \rangle \simeq a b^{1+n-\tau} L^{D(1+n-\tau)} g_n . \quad (7.11)$$

Neither a nor b are universal and neither are the g_n unless one fixes some features of $\mathcal{G}(x)$ such as its normalisation and its maximum. To extract universal characteristics of $\mathcal{G}(x)$, moment ratios can be taken for example

$$\frac{\langle s^{n-1} \rangle \langle s^{n+1} \rangle}{\langle s^n \rangle^2} = \frac{g_{n-1} g_{n+1}}{g_n^2} + \text{corrections} \quad (7.12)$$

or

$$\frac{\langle s^n \rangle \langle s \rangle^{n-2}}{\langle s^2 \rangle^{n-1}} = \frac{g_1^{n-2}}{g_2^{n-1}} g_n + \text{corrections} , \quad (7.13)$$

which is particularly convenient because of its very simple form when fixing $g_1 = g_2 = 1$ by choosing the metric factors a and b appropriately.

The most important result of a moment analysis, however, are the universal exponents D and τ and corresponding pairs for avalanche duration (z and α respectively), as well as the area (normally D_a and τ_a) *etc.*. This is done in a three step process. Firstly, the SOC model is run with different systems sizes, typically spaced by a factor 2, or 2, 5, 10. It can pay to use slightly “incommensurate” system sizes to identify systematic effects, for example

due to boundary effects being particularly pronounced in system sizes that are powers of 2. A typical simulation campaign would encompass 10 to 15 system sizes, of which maybe only 6 to 10, stretching over two to four orders of magnitude¹² will be used to produce estimates of exponents. The result of the simulation are estimates for the moments of the relevant observables together with their error (see below).

Secondly, the moments of the event sizes distribution, $\langle s^n \rangle$, are fitted against a power law in L (which is the parameter controlling s_c) with corrections,

$$\langle s^n \rangle = A_0 L^{\mu_n} + A_1 L^{\mu_n - \omega_1} + \dots \quad (7.14)$$

with positive exponents ω_i , known as confluent singularities; in particular $\mu_n - \omega_1$ is sometimes referred to as a sub-dominant exponent. The introduction of such **corrections to scaling** goes back to Wegner (1972), who applied it in equilibrium critical phenomena. The Levenberg-Marquardt algorithm (Press *et al.* 2007) is probably the fitting routine most frequently employed for matching the estimates (with their error bars) from the simulation to the fitting function Eq. (7.14). There are a number of problems that can occur:

- Unless the result is purely qualitative, a good quality fit cannot be achieved without good quality numerical data, that includes a solid estimate of the numerical error, *i.e.* the estimated standard deviation of the estimate.
- The very setup of fitting function Eq. (7.14) (sometimes referred to as “the model”) can introduce a systematic error; after all it is only a hypothesis.
- If $n > \tau - 1$ is very small, corrections due to the presence of the lower cutoff (s_0 , Eq. (7.3)) can be very pronounced.
- The error stated for the fitted exponents alone can be misleading. If Eq. (7.14) is very constraining, the error will be low, but so will the goodness-of-fit.
- Too many fitting parameters allow for a very good goodness of fit, but also produce very large estimated statistical errors for the exponents.
- Fitting against a function with many parameters often is highly dependent on the initial guess. In order to achieve good convergence and systematic, controlled results, it may pay off to fit the data against Eq. (7.14) step-by-step, using the estimates obtained in a fit with fewer corrections as initial guesses for a fit with more corrections.
- In most cases, there is little point in having as many parameters as there are data points, as it often produces a seemingly perfect fit (goodness-of-fit of unity), independent of the input data.

¹² One might challenge the rule of thumb of the linear system size L having to span at least three orders of magnitude — in higher dimensions, say $d = 5$, spanning three orders of magnitude in linear extent leads to 15 orders of magnitude in volume, which might be the more suitable parameter to fit against.

- Extremely accurate data, *i.e.* estimates for the moments with very small error bars, may require a large number of correction terms.
- It can be difficult to force the corrections ω_i to be positive. It is not uncommon to fix them at certain reasonable values such as $\omega_i = i$ or $\omega_i = i/2$. Alternatively, they can be introduced differently, writing them, for example, in the form $\omega_i = i + |\tilde{\omega}_i|$.
- If finite size scaling applies, the relative statistical error for any moment scales like $\langle s^{2n} \rangle / \langle s^n \rangle^2 \propto L^{D(\tau-1)}$, assuming that $\sigma^2(s^n)$ scales like $\langle s^{2n} \rangle$, which it certainly does for $\tau > 1$. At $\tau = 1$ the scaling of $\sigma^2(s^n)$ may be slower than that of $\langle s^n \rangle^2$. While $L^{D(\tau-1)}$ does not depend on n , the amplitude of the moments does, leading normally to an increase of the relative error with n .

In some models the first moment of the avalanche size displays anti-correlations and thus faster numerical convergence as found in a mutually independent sample (Welinder *et al.* 2007). In many models, the average avalanche size $\langle s \rangle$ is known exactly, in one dimension often including the confluent singularities (Pruessner 2012a). These exact results can provide a test for convergence in numerics and also provide a **scaling relation**

$$D(2 - \tau) = \mu_1 \quad (7.15)$$

If μ_1 is known exactly ($\mu_1 = 2$ for bulk driving Manna, Oslo and Abelian BTW Models, $\mu_1 = 1$ for boundary drive), then Eq. (7.15) gives rise to a **scaling relation**. Normally, there are no further, strict scaling relations. However, the assumption of narrow joint distributions suggests $D(\tau - 1) = z(\alpha - 1)$ *etc.* (Christensen *et al.* 1991, Chessa *et al.* 1999a, Pruessner and Jensen 2004). If the exponent μ_1 is given by a mathematical identity and $\langle s \rangle$ serves as an analytically known reference in the numerical simulation, then μ_1 should not feature in the numerical analysis to extract the scaling exponents D and τ . Rather, when fitting μ_n versus $D(1 + n - \tau)$, this should be replaced by $D(n - 1) + \mu_1$.

Fitting μ_n in a linear fit (without corrections) against $D(1 + n - \tau)$ (or against $D(n - 1) + \mu_1$ if μ_1 is known exactly) is, in fact, the third step in the procedure described in this section. In principle, the $n > \tau - 1$ do not need to be integer valued. They have to be large enough to avoid a significant corrections due to the lower cutoff, and small enough to keep the relative statistical error small. Non-integer n can be computationally expensive, as they normally require at least one library call of `pow`.

While each estimate μ_n for every n should be based on the entire ensemble, considering them together in the same fit to extract the exponents D and τ introduces correlations, which are very often unaccounted for. As a result both goodness-of-fit as well as the statistical error for the exponents extracted are (unrealistically) small.

There are a number of strategies to address this problem. The simplest is to up-scale the error of the μ_n as if every estimate was based on a separate,

independent set of raw data. Considering M moments simultaneously, their error therefore has to be scaled up by a factor \sqrt{M} (Huynh *et al.* 2011). In a more sophisticated approach, one may extract estimates from a series of sub-samples (Efron 1982, Berg 1992, 2004).

It often pays to go through the process of extracting the exponents D and τ at an early stage of a simulation campaign, to identify potential problems in the data. Typical problems to watch out for include

- Corrections are too strong because system sizes are too small.
- Results are too noisy because sample sizes are too small, often because the system sizes are too big.
- Results have so little noise that fitting functions need to contain too many free parameters.
- Too few data points (*i.e.* too few different system sizes L or different moments n).
- Large event sizes suffer from integer overflow, resulting in seemingly negative or very small event sizes.
- Data identical in supposedly different runs, because of using the same seed for the random number generator.
- Transients chosen too short.

7.2.3 Statistical errors from chunks

One of the key-ingredients in the procedures described above is a reliable estimate for the statistical error of the estimates of the individual moments. The traditional approach is to estimate the variance, $\sigma^2(s^n) = \langle s^{2n} \rangle - \langle s^n \rangle^2$ of each moment, so that the statistical error of the estimate of $\langle s^n \rangle$ is estimated by $\sigma^2(s^n)/\sqrt{N/(2\tau+1)}$, where $N/(2\tau+1)$ is the number of effectively independent elements in the sample with correlation time τ .

This approach has a number of significant drawbacks. Firstly, each moment $\langle s^n \rangle$ requires a second moment, $\langle s^{2n} \rangle$, to be estimated as well. Considering a range of moments, this might (almost) double the computational effort. Rather dissatisfyingly, the highest moment estimated itself cannot be used to extract its finite size scaling exponent μ_n , because its variance is not estimated. Furthermore, because of their very high powers, the moments entering the estimates of the variances and thus the variances themselves have large statistical errors and are prone to integer overflow.

Estimating the effective number of independent elements in the sample is a hurdle that can be very difficult to overcome. Usually, it is based on an estimate of the correlation time τ . If $\langle s_i s_j \rangle - \langle s \rangle^2 = \sigma^2(s) \exp(-|i-j|/\tau)$, then the variance of the estimator

$$\bar{s} = \frac{1}{N} \sum_i^N s_i \quad (7.16)$$

of $\langle s \rangle$ for $N \gg \tau$ is

$$\sigma^2(\bar{s}) = \frac{1}{N^2} \sum_{ij}^N \left(\langle s_i s_j \rangle - \langle s \rangle^2 \right) \approx \frac{1 + \exp(-1/\tau)}{N(1 - \exp(-1/\tau))} \sigma^2(s) \approx \frac{2\tau + 1}{N} \sigma^2(s) \quad (7.17)$$

as if the sample contained only $N/(2\tau + 1)$ independent elements.

The main difficulty of this strategy is a reliable estimate of τ which often cannot be easily extracted from $\langle s_i s_j \rangle - \langle s \rangle^2$ because of noise and the presence of other exponential contributions, of which $\exp(-|i - j|/\tau)$ is the slowest decaying one. Moreover, in principle τ has to be measured for each observable separately (even when it makes physically most sense to assume that the system is characterised by a single correlation time).

To avoid these difficulties, one may resort to a simple sub-sampling plan. As discussed below (also Sec. 7.3.5), it is a matter of mere convenience and efficiency to repeatedly write estimates of moments based on a comparatively small sample into the output stream of a simulation and reset the cumulating variables. In the following these raw estimates based on a small sample are referred to as **chunks**. If their sample size is significantly larger than the correlation time, then each of these estimates can be considered as independent and the overall estimates based on it has its statistical error estimated accordingly. For example, if m_i with $i = 1, 2, \dots, M$ are estimates of $\langle s^n \rangle$ all based on samples of the same size N , say $m_i = \sum_j^N s_{ij}^n$ with s_{ij} the j th element of the i sample, then the overall unbiased and consistent estimator (Brandt 1998) of $\langle s \rangle$ is

$$\bar{m} = \frac{1}{M} \sum_i^M m_i \quad (7.18)$$

which has an estimated standard deviation of $(\overline{m^2} - \bar{m}^2)/(M - 1)$ where

$$\overline{m^2} = \frac{1}{M} \sum_i^M m_i^2. \quad (7.19)$$

One crucial assumption above is that the m_i are independent, which can always be achieved by merging samples. As long as M remains sufficiently large, one may be generous with the (effective) size of the individual samples (Flyvbjerg and Petersen 1989).

Chunks also allow a more flexible approach to determining and discarding transient behaviour from the sample supposedly taken in the stationary state. The transient can be determined as a (generous) multiple of the time after which (ideally all or several) observables no longer deviate more from the asymptotic or long time average than their characteristic variance. Where observables are known exactly (e.g. the average avalanche size Pruessner 2012a), they can be used as a suitable reference. Figure 7.5 shows the transient behaviour of the average avalanche size in a realisation of the Manna Model. A




Fig. 7.5: Example of the transient behaviour of an observable (here the average avalanche size in the one-dimensional Manna Model with $L = 65536$) as a function of the chunk index in a log-lin plot (data from Huynh *et al.* 2011). The straight dashed line shows the exact expected average $\langle s \rangle$, Eq. (7.1). The arrow indicates the chunk from where on stationarity is roughly reached. A generous multiple of that time should be taken as the number of chunks to discard in order to ensure that correlations (and thus dependence on the initial setup) are essentially overcome.

more cautious strategy is to consider a series of different transients and study the change in the final estimates (with their estimated error) as a function of the transient discarded.

7.3 Algorithms and data organisation

In the following, a range of numerical and computational procedures are discussed that are commonly used in the numerical implementation of SOC models (for a more extensive review see Pruessner 2012b). Some of them are a matter of common sense and should be part of the coding repertoire of every computational physicist. However, it is not always entirely obvious how these “standard tricks” are used for SOC models.

In the following, the focus is on computational performance, which often comes with the price of lower maintainability of the code. The amount of real time spent on writing code and gained by making it efficient, should account for the time spent on debugging and maintaining it.

Most of the discussion below is limited to algorithmic improvements. The aim is produce code that communicates only minimally with the “outside

world”, because in general, interaction with the operating system, as required for writing to a file, is computationally expensive and extremely slow. The UN*X operating system family (including, say, Linux and Mac OS X) distinguishes two different “modes” by which an executable keeps the CPU busy: By spending time on the (operating) system and by spending it in “user mode”. Roughly speaking, the former accounts for any interaction between processes, with external controls or peripherals, including writing files. The latter accounts for the computation that takes place solely on the CPU (ALU, FPU, GPU, *etc.*) and the attached RAM. Tools like `time` and library functions like `getrusage` provide an interface to assess the amount of various resources used, while being themselves or resulting in systems calls.

Of course, the literature of computational physics in general is vast. Reviews and texts that are of particular use in the present context include Kernighan and Ritchie (1988), Cormen *et al.* (1996), Knuth (1997), Newman and Barkema (1999), Berg (2004), Landau and Binder (2005), Press *et al.* (2007).

7.3.1 Stacks

The definition of most SOC models makes no reference to the method to identify active sites, *i.e.* sites that are due to be updated. In principle, an implementation of an SOC model could therefore repeatedly scan the entire lattice to find the relevant sites. This is, however, very inefficient and therefore should be avoided. Instead, the index of active sites (or their coordinates) should be organised in a list. Every site in that list is subsequently updated. Moreover, it is often very important to know whether a site is maintained in the list or not. Sometimes this can be determined implicitly (for example, when a site is guaranteed to reside on the list from the moment its height exceeds the threshold), sometimes this is done explicitly by means of a flag associated with the site. The following contains a more detailed discussion of the various techniques available.

The most commonly used form of a list is a **stack**, called so, because this is how it appears to be organised. It consists of a vector, say `int stack[STACK_SIZE]`, of pre-defined size `STACK_SIZE`. It must be large enough to accommodate the maximum number of simultaneously active sites. Simulating large lattices, a balance has to be struck between what is theoretically possible and what is happening in practise.

The type of the stack, vector of `int` in the example above, is determined by the type it is meant to hold. If it holds the index of active sites, it is likely to be `int`, but it may also hold more complex objects, say, coordinates of active particles (but see below). The number of objects currently held by the stack is stored in `int stack_height`.

If `STACK_SIZE` is smaller than the theoretical maximum of active sites, `int stack_height` has to be monitored as to prevent it from exceeding `STACK_SIZE`. The outcome of the simulation is undefined if that happens, because the exact position in memory of `stack[STACK_SIZE]` is *a priori* unknown. If therefore `stack_height` exceeds `STACK_SIZE`, memory has to be extended one way or another. For example, one may use `realloc()`, which assumes, however, that enough memory is actually available. Modern operating systems all provide virtual memory which is transparently supplemented by a swap file residing on the (comparatively slow) hard drive. This is to be avoided because of the computational costs associated. It may thus pay off for the process itself to make use, temporarily, of a file to store active sites. The alternative to abandon the particular realisation of the simulation introduces a bias away from rare events which is likely to have significant effect on observables. The same applies obviously if activity is suppressed if it reaches the maximum level.

There are two fundamental operations defined on a stack,

```
#define PUSH(a) stack[stack_height++]=a)
#define POP(a) (a)=stack[--stack_height]
```

where `PUSH(a)` places `(a)` on the stack and `POP` takes an element off. The underlying idea is literally that of a stack: When a site becomes active, its index goes on a pile (`PUSH`) so that each index number on that pile represents a site waiting to be updated. When that happens, it is removed from the pile (`POP`).

It simplifies the code greatly if all objects on the stack are, in a sense, equivalent. For example, all sites on a stack are active. Guaranteeing this is not necessarily trivial, because the manipulation of one item on the stack may affect the state (and thus the eligibility) of another item on the stack. It is therefore advisable to ensure that all elements on the stack are distinct. In SOC models that means that active sites enter the stack exactly once, namely when the *turn* active. If an active site is charged again by a toppling neighbour, a new copy of its index is *not* placed on the stack. In the Manna Model, for instance, the single line of code to place objects on the stack could be

```
if (z[i]++==1) {PUSH(i);}
```

so that the index `i` of a site enters when it is charged while its height `z` is *at* the critical value z^c . The line should *not* read `if (z[i]++>=1)PUSH(i);`.

Unfortunately, the very data structure of a stack, which in the present context may better be called a **LIFO** (last in, first out), suggests a particular procedure to explore active sites, namely a depth first search (DFS); Whenever a toppling site activates its neighbours, one of them will be taken off first by the next call of `POP`, toppling in turn. Activity thus spreads very far very quickly, then returning, then spreading far again, rather than “burning locally”. In fact, in DS-FFM a DFS is probably the simplest way of exploring a cluster of trees.

The alternative, a breadth first search (BFS) requires slightly greater computational effort because it normally makes use of a **FIFO** (first in, first out). The last object to arrive on a FIFO is the last one to be taken off, exactly the opposite order compared to a stack. Naively, this may be implemented by removing items from the front, `stack[0]`, and using `memmove()`¹³ to feed it from the end, lowering `stack_height`. This approach, however, is computationally comparatively costly. A faster approach is to organise the stack in a queue, organised in a ring (circular buffer) to keep it finite, where a string of valid data grows at the end while retreating from the front.

In Abelian models, where the statistics of static features of avalanches, such as size and area, do not depend on the details of the microscopic dynamics¹⁴, working through the stack using `POP` may be acceptable. Where temporal features are of interest too, the microscopic dynamics must implement a suitable microscopic time scale. Often the microscopic timescale is given by Poissonian updates, for example by active sites toppling with a Poissonian unit rate.

In principle that means that waiting times between events (sites toppling) are themselves random variables. If a faithful representation of the microscopic time is desired, then the random waiting times can be generated by taking the negative logarithm of a random number drawn from a uniform distribution on $(0, 1]$. If an approximate representation of the Poisson processes is acceptable (which, in fact converges to the exact behaviour in the limit of large numbers of active sites, see Liggett 2005), then elements are taken off the stack at random and time is made to progress in steps of `1./stack_height`. If `stack_height` remains roughly constant, then on average `stack_height` events occur per unit time as expected in a Poisson process. A simple implementation reads

```
int rs_pos;
#define RANDOM_POP(a) rs_pos=rand() % stack_height; (a)=stack[rs_pos
↳   ]; POP(stack[rs_pos])
```

where the last operation, `POP(stack[rs_pos])` overwrites the content of `stack[rs_pos]` by `stack[stack_height-1]` decrementing `stack_height` at the same time. When selecting the random position on the stack via `rs_pos=rand()%stack_height` a random number generator has to be used (Sec. 7.3.4), which only for illustrative purposes is called `rand()` here.

One consequence of the constraint of distinct objects on the stack is that a site may need to topple several times before being allowed to leave the stack. In Abelian models some authors circumvent that by placing a copy of the site index on the stack every time a pair of particles has to be toppled from it, which can be implemented easily by removing an appropriate number of

¹³ Dedicated library functions like `memmove` and `memcpy` are generally much faster than naive procedures based on loops, although the latter can be subject to significant optimisation by the compiler.

¹⁴ But note the strict definition of Abelianness discussed on p. 257.

particles from the site each time it enters the stack. As a result, however, stacks may become much larger, *i.e.* a greater amount of memory has to be allocated to accommodate them.

Depending on the details of the microscopic dynamics, an possible alternative is to relax a site completely after it has been taken off the stack, for example in the Manna Model:

```

while (stack_height) {
    RANDOM_POP(i);
    do {
        topple(i); /* Site i topples, removing two particles from i. */
        avalanche_size++; /* avalanche_size counts the number of topplings. */
    } while (z[i]>1);
}

```

where `topple(i)` reduces `z[i]` by 2 each time. If the avalanche size counts the number of topplings performed, `avalanche_size` has to be incremented within the loop. Counting only *complete* relaxations would spoil the correspondence with exact results.

An alternative approach with different microscopic time scale is to topple a site on the stack only once, and take it off only once it is fully relaxed. This approach requires some “tempering” with the stack:

```

while (stack_height) {
    i=rand() % stack_height;
    topple(stack[i]);
    if (z[i]<=1) POP(stack[i]);
}

```

In systems with parallel update, where all sites at the beginning of a time step have to be updated concurrently before updating the generation of sites that have been newly activated, a red-black approach (Dowd and Severance 1998) can be adopted. This requires the use of two stacks, which have to be swapped after completing one:

```

int *stack, stack_height=0;
int rb_stack[2][STACK_SIZE], next_stack_height;
int current_stack, next_stack;

#define NEXT_PUSH(a) rb_stack[next_stack][next_stack_height++]=a
#define NEXT_POP(a) (a)=rb_stack[next_stack][--next_stack_height]

...
current_stack=0;
next_stack=1;
stack=rb_stack[current_stack];
...
PUSH(i);
...
for (;;) {
    while (stack_height) {
        ...
        POP(i);
    }
}

```

```

...
    NEXT_PUSH(j);
...
}
if (next_stack_height==0) break;
/* Swap stacks. */
stack_height=next_stack_height;
next_stack_height=0;
current_stack=next_stack;
stack=rb_stack[current_stack];
next_stack=1-next_stack;
}
/* Both stacks are empty. */

```

The use of the pointer `stack` is solely for being able to use the macros `PUSH` and `POP` defined earlier. Otherwise, it might be more suitable to define macros `CURRENT_PUSH` and `CURRENT_POP` corresponding to `NEXT_PUSH` and `NEXT_POP`.

A stack should also be used when determining the area of an avalanche, *i.e.* the number of distinct sites toppled (or visited, *i.e.* charged). To mark each site that has toppled during an avalanche and to avoid double counting, a flag has to be set, say `visited[i]=1` or `site[i].visited=1` (see Sec. 7.3.2). Counting how often the flag has been newly visited then gives the avalanche area. However, in preparation for the next avalanche, the flags have to be reset. This is when a stack comes handy, say

```

int area_stack[SYSTEM_SIZE];
int area_stack_height=0;
#define AREA_PUSH(a) area_stack[area_stack_height++]=(a)
#define AREA_POP(a) (a)=area_stack[--area_stack_height]
...
/* For each toppling site. */
if (visited[i]==0) {
    visited[i]=1;
    AREA_PUSH(i);
}
...
/* After the avalanche has terminated.
 * area_stack_height is the avalanche area. */
...
/* Re-initialise */
while (area_stack_height) {
    AREA_POP(i);
    visited[i]=0;
}
...

```

In the example above, the area is tracked implicitly in `area_stack_height`. The re-initialisation can be further improved using `while (area_stack_height) visited[area_stack[--area_stack_height]]=0`.

7.3.2 Sites and Neighbours

In SOC models, every site has a number of properties, most importantly the local degree of freedom, but also (statistical) observables which are being measured and updated as the simulation progresses. Other information associated with each site are flags (such as the one mentioned above to indicate whether a site had been visited) and even the neighbourhood (discussed below). In fact, the site itself may be seen as the **key** associated with all that information. That key might represent information in its own right, say, the coordinate, it might be an index of a vector, or a pointer.

7.3.2.1 Pointers and structures

A word of caution is in order with regard to pointers. The programming language C lends itself naturally to the use of pointers. However, code on the basis of pointers is difficult to optimise automatically at compile time. Depending on the quality of the compiler and the coding an index based implementation (which is also more portable) may thus result in faster code than the seemingly more sophisticated implementation based on pointers.

That said, in theory placing pointers on the stack, which gives immediately access to a relevant object should be faster than using indices, which are effectively an offset relative to a base: `b=z[stack[i]]` might result in machine code of the form `b=(z+(stack+i))` which contains one more addition than `b=*stack[i]` resulting in `b=*(stack+i)` if `stack` is a vector of pointers.

Similar considerations enter when using structures, which provide very convenient and efficient ways of organising and encapsulating data associated with each site. For example

```

struct site_struct {
    int height;
    char visited;
};

```

defines a structure with two members, `height` and `visited`. Declaring a variable `struct site_struct site[10]` allows the individual elements to be accessed in a structured way, say `site[i].height++`, `site[i].visited=1`. There are a number of computational drawbacks, which are, however, normally outweighed by the better maintainability of the code.

- Depending on the platform and the compiler, padding might become necessary, *i.e.* some empty space is added to the structure (Sec. 7.3.2.2, p. 292). The memory requirements of the structure is thus greater than the memory requirements for each variable when defined individually.
- Again depending on the platform as well as the compiler, without padding some operations on some types may require more CPU cycles (in particular when floating point types are used).

- Members within the structure are accessed similar to elements in a vector, namely by adding an offset. Access to the first member (where no offset is needed, `site[i].height` in the example above) can thus be faster than access to the other members (`site[i].visited` above). Because of that additional addition, the approach is often slower than using separate vectors for each member of the structure.

7.3.2.2 Neighbourhood information

It can be convenient, in particular for complicated topologies or when the neighbourhood information is computed or supplied externally, to store information about the local neighbourhood in a site structure, for example:

```

struct site_struct {
    ...
    int neighbour[MAX_NEIGHBOURS];
    int num_neighbours;
};

```

Because of the significant memory requirements, this is often not viable for large lattices. Again, instead of addressing neighbours by their index, pointers can be used, which often produces very efficient and elegant code.

The neighbours of each site thus are calculated and stored at the site only once. The strategy of pre-calculated neighbourhoods goes back to the very beginning of computational physics, when access to memory was much faster than doing such calculations on-the-fly.¹⁵ This, however, has changed. It can be *much* faster to determine a neighbourhood on-the-fly than looking it up, unless, of course, the topology is so complicated that it becomes computationally too costly. Unfortunately, it is often difficult to try out different implementations (lookup tables and calculation on the fly), as the setup of a neighbourhood is at the heart of a lattice simulation.

As for calculating neighbourhoods, in one dimension the index of a site, which is strictly only a key to access the information, is often associated with its position on a one-dimensional lattice. Actual computation takes place only at boundaries. If the right neighbour of site `i` in the bulk is `i+1`, it may not exist on the right boundary or be 0 if periodic boundary conditions (PBC) apply in an implementation in C where the index of a vector of size `LENGTH` can take values from 0 to `LENGTH-1`. Similarly, the left neighbour is `i-1` in the bulk and `LENGTH-1` at `i=0` in case of periodic boundaries. Those are most easily implemented in the form `left=(i+LENGTH-1)%LENGTH` and `right=(i+1)%LENGTH` respectively using a modulo operation. The shift by `LENGTH` in the former avoids problems with negative indices at `i=0`.

A less elegant but often faster implementation is to determine whether a site is at the boundary before assigning the value for the neighbour, such as

¹⁵ Back in the days when lookup tables for modulo operations were in fashion.

```

if (i==0) left=LENGTH-1;
else left=i-1;

```

or just `left=(i==0)?LENGTH-1:i-1`, which is more readable. This method is also more flexible with respect to the boundary condition implemented. Reflecting boundary conditions, for example are implemented by `left=(i==0)? 1 : i-1`. Open boundary conditions, on the other hand, might require special attention. If possible, they are best implemented using **padding**, *i.e.* by pretending that a neighbouring site exists, which, however, cannot interact with the rest of the lattice, for example, by making sure that it never fulfils the criterion to enter the stack. Such a site may need to be “flushed” occasionally to prevent it, for example, from fulfilling the criterion due to integer overflow. One might either assign one special site, say the variable `dump` in `left=(i==0)? dump : i-1` or allocate memory for `LENGTH+2` sites with an index from 0 to `LENGTH+1`, with valid sites ranging from 1 to `LENGTH` with sites 0 and `LENGTH+1` receiving charges without toppling in turn. This procedure also allows a very efficient way to determine the number of particles leaving the system, the **drop number** (Kadanoff *et al.* 1989).

Usually only in higher dimensions, one distinguishes reflecting boundary conditions, where the particle offloaded is moved to another site (normally the mirror image of the “missing” site), and “closed” boundary conditions, where the number of nearest neighbours is reduced and shed particles are evenly re-distributed among them.

Most of the above techniques remain valid in higher dimension, where the data can be organised in either a one-dimensional vector or a multidimensional vector. The former strategy makes use of macros of the form

```

#define COORDINATE2INDEX(x,y,z) ((x)+(LENGTH_X*((y)+
↳   LENGTH_Y*(z)))
#define INDEX2COORDINATE(i,x,y,z) z=(i)/(LENGTH_X*LENGTH_Y),
↳   y=((i)/LENGTH_X)%LENGTH_Y,x=(i)%LENGTH_X

```

The use of the coma operator in the second macro helps to avoid errors when omitting curly brackets in expressions like `if (1)INDEX2COORDINATE(i,x,y,z);`. Where stacks are used to hold coordinates, the multiple assignments needed to store and fetch all of them may computationally outweigh the benefit of not having to calculate coordinates based on a single index.

The two biggest problem with the use of multi-dimensional vectors is their ambiguity when used with fewer indices and the consistency when passing them to functions. Both subtleties arise because of the logical difference between a vector of pointers to a type and the interpretation of a lower-dimensional variant of a multi-dimensional vector. While C makes that distinction, there is no syntactical difference between the two. For example

```

int a[2][10];

a[0][5]=7;

```

is a multi-dimensional vector using up $2 \cdot 10 \cdot \text{sizeof}(\text{int})$ sequential bytes of memory. Each $\text{a}[i]$ is the starting address of each row $i = 0, 1$. On the other hand

```

int *a[2];
int row1[10], row2[10];
a[0]=row1; a[1]=row2;

a[0][5]=7;

```

makes a a vector of pointers, using up $2 \cdot \text{sizeof}(* \text{int})$ bytes of memory, while each row uses $10 \cdot \text{sizeof}(\text{int})$ bytes. Both snippets of code declare a to be completely different objects, yet, for all intents and purposes in both cases a will behave like a two-dimensional array. That is, until it is to be passed as an argument to another function. In the first case, that function can be declared by `function(int array[2][10])`, informing it about the dimensions of the array, and subsequently called using `function(a)`. The two-dimensional vector a will behave as in the calling function. In fact, the function will even accept any other vector, lower dimensional or not, passed on to it as an argument (even when the compiler may complain).

In the second case, a is a vector of pointers to `int`, and so a function taking it as an argument must be declared in the form `function(int **a)`, using additional arguments or global constants (or variables) to inform it about the size of the vector. The two versions of the functions are incompatible, because a two-dimensional vector is really a one-dimensional vector with a particularly convenient way of addressing its components. In particular, the two-dimensional vector cannot be passed to the function designed for the second case using, say, `function(&a)` or `function((int **)a)`.

While these issues normally are resolved at the time of coding they can cause considerable problems when the memory allocation mechanism for the vector is changed. This happens, in particular, when lattice sizes are increased during the course of a simulation campaign. Initially, one might be tempted to define a lattice globally (stored in BSS or data segment) or as automatic variables taken from the stack, choosing a multi-dimensional array for convenience. Later on, they make be taken from the (usually much bigger) heap using `malloc()`, at which point the way they are accessed may have to be changed. The latter approach is the most flexible but possibly not the most convenient way of allocating memory for large items.

Finally, it is advisable to scan sites (when sweeping the lattice is unavoidable or scanning through a local neighbourhood) in a way that is local in memory and thus cache. The first option, declaring a two-dimensional vector in a single step, makes that more feasible than the second option, where different rows might end up at very different regions of memory. Not using higher dimensional vectors at all, however, is probably the best performing option.

7.3.3 Floating Point Precision

Very little and at times too little attention is being paid to the effect of limited floating point precision. Most SOC models can be implemented fully in integers even when their degrees of freedom are meant to be real valued, such as the Zhang Model (Zhang 1989), the Bak-Sneppen Model (Bak and Sneppen 1993) or the Olami-Feder-Christensen Model (Olami *et al.* 1992). In case of the latter, floating point precision has been found to significantly affect the results (Drossel 2002).

Where **random floating point** numbers are drawn, they might in fact contain much fewer random bits than suggested by the size of their mantissa. In that case, an implementation in integers is often not only faster but also “more honest”. Where rescaling of variables cannot be avoided and occurs frequently, multiplying by a constant inverse often produces faster code than division.

Over the last decade or so, the floating point capabilities of most common CPUs have improved so much, however, that the difference in computational costs between integers and floating point arithmetics is either negligible or not clear-cut. The most significant disadvantage of the latter is the limited control of precision that is available on many platforms.

The levels of precision as defined in the IEEE standard 754 that are very widely used are single, double and extended. They refer to the number of bits in the mantissa determined when floating point operations are executed, *i.e.* they are the precision of the floating point unit (FPU). The precision the FPU is running at depends on platform, environment, compiler, compiler switches and the program itself. Some operating systems offer an IEEE interface, such as `fpsetprec()` on FreeBSD, and `feenv` on Linux.

Results of floating point arithmetics are stored in variables that may not offer the same level of precision the FPU is running at and in fact it is possible that none of the data types available matches a particular level of precision set on the FPU. Crucially, the precision setting of the FPU normally affects *all* floating point operations on *all* floating point variables, regardless of type, *e.g.* information is lost when results are calculated with extended precision and stored in variables offering only single precision. A notorious error observed on systems which default to extended precision, in particular Linux on x86, occurs when comparisons between variables produce different outcomes depending on the position in the code — at one point the result calculated may still reside on the FPU and thus offer extended precision, whereas at a later point the result is truncated after being written to memory. This can lead to serious inconsistencies when data is held in an ordered tree. Compiler switches like `-ffloat-store` for `gcc` help in these cases.

The commonly used `gcc` compiler offers three basic floating point types, `float`, `double` and `long double`, matching the three levels of precision mentioned above. The very nature of SOC means that observables span very many orders of magnitude. If variables that accumulate results, such as moments,

are too small (*i.e.* have a mantissa that is too small), smaller events may not accumulate at all any more once the variable has reached a sufficiently large value. This can skew estimates considerably where very large events occur very rarely. The macros `FLT_EPSILON`, `DBL_EPSILON` and `LDBL_EPSILON` in `float.h` give a suggestion of the relative scale of the problem. It can be mitigated by frequently “flushing” accumulating variables (see Sec. 7.3.5).

7.3.4 *Random Number Generators*

Random Number Generators (RNGs) are a key ingredient in many areas of computational physics, in particular in Monte-Carlo and Molecular Dynamics simulations. The vast majority of them, strictly, are not random, but follow instead a deterministic but convoluted computational path. RNGs are constantly being improved and evaluated, not least because of their use in cryptography. An introduction into the features of a good RNG can be found in the well-known Numerical Recipes (Press *et al.* 2007), with further details to be found in the review by Gentle (1998).

A “good” random number generator is one that offers a reasonable compromise between two opposing demands, namely that of speed and that of quality. In most stochastic SOC models, the RNG is used *very* often and thus typically consumes about half of the overall CPU time. Improving the RNG is thus a particularly simple way of improving the performance of an implementation. Because the variance (square of the standard deviation) of an estimate vanishes inversely proportional with the sample size it is based on, the performance of an implementation is best measured as the product of variance and CPU time spent “for it”. However, one is ill-advised to cut corners by using a very fast RNG which has statistical flaws. The resulting problem may be very subtle and might not show until after a very detailed analysis.

One of the problems is the period of an RNG. Because RNGs generally have a finite state, they are bound to repeat a sequence of random numbers after a sufficient number of calls, at which point the simulation using the random numbers produces only copies of previous results. With improving hardware the RNG must therefore be re-assessed. A “good RNG” is a function of time, and very much a function of perception, as a mediocre RNG might appear to be a fantastic improvement over a poor RNG. It is good practise to use more than one random number generator to derive the same estimates and compare the results.

The C library’s implementation of `rand()` is legendary for being unreliable and can be very poor. At the very least, it is essentially uncontrolled, although, of course, standards exist, which are, however, not always adhered to. It is fair to say that pure linear congruential RNGs are somewhat (out-)dated and indeed rarely used. They are, however, sometimes combined or

enhanced with more sophisticated techniques. In recent years, the Mersenne Twister (Matsumoto and Nishimura 1998, Matsumoto 2008) has become very widely used, yet, criticised by Marsaglia (2005) who proposed in turn KISS (Marsaglia 1999, but see Rose 2011), which is a remarkably simple RNG. The GNU Scientific Library (Galassi *et al.* 2009) contains an excellent collection of random number generators.

Somewhat more specific to the use of RNGs in SOC models is the frequent demand for random bits, for example in order to decide about the direction a particle is taking. Because every acceptable RNG is made up of equally random bits, each and everyone of them should be used for random booleans. These bits can be extracted one-by-one, by bit-shifting the random integer or by shifting a mask across, as in

```
#define RNG_MT_BITS (32)
#define RNG_TYPE unsigned long
RNG_TYPE mt_bool_rand=0UL;
RNG_TYPE mt_bool_mask=1UL<<(RNG_MT_BITS-1);
#define RNG_MT_BOOLEAN ( ( mt_bool_mask==(1UL<<(
↳   RNG_MT_BITS-1)) ) ? ((mt_bool_mask=1UL, mt_bool_rand=
↳   genrand_int32()) & mt_bool_mask) : (mt_bool_rand & (mt_bool_mask
↳   +=mt_bool_mask)) )
```

based on the Mersenne Twister. In general, bit shifts to the left using `a+=a` instead of `a<<=1` are faster, because the latter requires one more CPU cycle to write the constant 1 into the CPU's register.

More generally, integer random numbers have to be chosen uniformly from the range $\{0, 1, \dots, n-1\}$ suggesting the use of the modulo operation, `r=rand()%n`. However, if `rand()` produces random integers uniformly from 0 up to and including `RAND_MAX`, then the modulo operation skews the frequencies with which random number occurs towards smaller values if `RAND_MAX+1` is not an integer multiple of n . The effect is of order $n/(RAND_MAX+1)$ and thus is negligible if n is significantly smaller than `RAND_MAX`. However, picking a site at random on a very large lattice or an element from a very large stack, this effects becomes a realistic concern. In that case, the modulo operation can be used on a random number drawn uniformly among integers from 0 up to and including $R-1$, where R is a multiple of n and ideally the largest multiple of n less or equal to `RAND_MAX+1`:

```
const long long int n=...;
/* The constant multiple_minus_1 is made to have type as the return
 * value of rand(). */
const int multiple_minus_1=(n*(((long long int)RAND_MAX) + 1LL)/n)
↳   )-1LL;
int r;
#define RANDOM(a) while ((r=rand())>multiple_minus_1); (a)=r%n
```

where `multiple_minus_1` plays the rôle of $R-1$. When determining the maximum multiple, it is crucial that the operation `RAND_MAX+1` is performed using a type where the addition does not lead to rounding or integer overflow. The latter is also the reason why one is subtracted in the expression for

`multiple_minus_1`, which otherwise might not be representable in the same type as the return value of `rand()`, which is necessary to avoid any unwanted type casting at run time.¹⁶

The initial seed of the RNG needs to be part of the output of the programme it is used in, so that the precise sequence of events can be reproduced in case an error occurs. Some authors suggest that the initial seed itself should be random, based, for example, on `/dev/random`, or the library functions `time()` or `clock()`,¹⁷ and that the RNG carries out a “warm-up-cycle” of a few million calls (Jones 2012). After that, it is sometimes argued, chances are that one sequence of (pseudo) random numbers is independent from another sequence of random numbers generated by the same RNG based on a different seed. Fortunately, some RNGs, in particular those designed for use on parallel machines, offer a facility to generate sequences that are guaranteed to be independent. Where poor-man’s parallel computing (many instances of the same simulation running with different seeds) takes place, independent sequences are of much greater concern than in situations where different parameter settings are used in different instances. In the former case the data of all instances will be processed as a whole, probably under the assumption that it is actually independent. In the latter case, the results will enter differently and using even an identical sequence of random numbers will probably not have a noticeable effect. All these caveats are put in perspective by the fact that most SOC models fed by a slightly differing sequences of pseudo random numbers take “very different turns in phase space” and thus will display very little correlations.

7.3.5 Output

As mentioned above, it is generally advisable to output and flush data frequently in **chunks**, resetting accumulating variables afterwards. Even when output occurs every second, the overhead in terms of the CPU and real time spent by the system is likely to be negligibly small.

Where data is written to a file in large quantities or frequently, **buffered I/O** as provided by `stdio` through the `printf`-family of library calls is usually much faster than writing immediately to the file using `unistd`’s `write`. There are two caveats to this approach: Firstly, depending on the size of the buffer and thus the frequency of writing, a significant amount of CPU time may be

¹⁶ This is one of the many good reasons to use constants rather than macros (van der Linden 1994, Kernighan and Pike 2002).

¹⁷ Both functions are bad choices on clusters where several instances of the same programme are intended to run in parallel. The function `time()` changes too slowly (returning the UN*X epoch time in seconds) and the function `clock()` wraps after about 36 minutes, so that neither function guarantees unique seeds. In general, seeding is best done explicitly.

lost if the program terminates unexpectedly. To avoid corrupt data, `fflush()` should be used rather than allowing the buffer to empty whenever it reaches its high-water mark. Secondly, if buffering I/O has a significant impact on the computational performance, the data may better be processed on-the-fly rather than storing it in a file.

In the following, `stdio` is used for its convenient formatting capabilities, provided by the plethora of flags in the formatting string of a `printf` call. To avoid the problems mentioned above, buffers are either flushed after each chunk by means of `fflush`, or buffering is switched to buffering line by line, using `setlinebuf`.

To avoid unexpected interference of the operating system with the simulation, operations should be avoided that can potentially fail because the environment changes. This applies, in particular, to read and write access to files. In any case, such operations need to be encapsulated in an `if` condition that catches failing system calls and triggers a suitable remedy.

Output of chunks should therefore happen through the `stdout` stream which is by default open at the time of the program start. As the output is usually used in post-processing it needs to be retained, which can be achieved by re-directing `stdout` into a file. In the typical shell syntax this can be done in the command line by, say, `./simulation > output.txt`. To allow easy post-processing, every line should contain all relevant simulation parameters, such as the system size, the number of the chunk (a counter), the number of events per chunk, the initial seed of the random number generator (RNG), in fact, everything that is needed to reproduce that line from scratch or to plot the relevant (derived) data. Typical examples are moments to be plotted against the system size and moment ratios, involving different moments of the same observable. Using post-processing tools to wade through vast amounts of data to find the missing piece of information to amend a line of data can require significant effort and is highly error-prone.

Repeating the same output (system size, RNG seed etc) over and over seemingly goes against the ethos of avoiding redundant information, which should be applied when setting up a computer simulation (to avoid clashes), but is wholly misplaced when it comes to data output. In fact, redundancy in output is a means to measure consistency and a matter of practicality as almost all basic post-processing tools are line-oriented.

In some rare cases, an action by the simulation or an event on the system can result in a **signal** being sent to the running instance of the program. In response the program suspends the current operation, executes a signal handler and continues where it left off. In principle, the signal should not lead to inconsistent data or behaviour; in fact, it is probably the most basic but also a very convenient way to communicate with a running program. For example

```
#include <signal.h>
...
void sighup_handler(int signo);
```

```

...
signal(SIGHUP, sighup_handler);
...
void sighup_handler(int signo)
{
    finish_asap=1;
}

```

assigns the signal handler `sighup_handler` to deal with the signal `SIGHUP`, which can be sent to the program using `kill -HUP`.

There is a rare situation when the signal interrupts in a way that it leads to unexpected behaviour, namely when it arrives while a “slow system call” is executed, *i.e.* an operation that is performed by the kernel on behalf of the programme, but which can take a long time to complete, such as `pause`, `sleep`, but also `write` to so-called pipes. Without discussing the technical details of the latter, it can lead to inconsistencies in the output which might not be detected in the post-processing. For example, a chunk may contain truncated lines and thus may lack certain information or data, which the post-processing tools might treat as zeroes. Apart from a graphical inspection of the data, two measures may therefore be advisable: Firstly, output can be encapsulated in calls of `sigprocmask` which allows temporary suspension of the delivery of signals. Secondly, a chunk can be terminated by a single line containing a keyword to indicate the successful completion of the output (*i.e.* without catching an error, in particular not an “interrupted system call”, `EINTR`), such as the tag (see below) `#Completed`. Simply counting the number of occurrences of that tag and comparing to (supposed) the number of valid chunks can pick up inconsistencies. In large scale simulations, where disk space can be a problem leading to truncated files as the system runs out of file space, this is particularly advisable.

After a chunk has been written out, variables collecting data have to be reset. Where PDFs are estimated, sweeping across the entire histogram can become expensive and therefore performing all relevant steps simultaneously is advantageous for the overall performance. Using one of the examples above (Sec. 7.2.1.1):

```

long long total=0;
for (i=0; i<SMALL2MEDIUM_THRESHOLD; i++)
    if (histo_small[i]) {
        printf(...);
        total+=histo_small[i];
        histo_small[i]=0;
    }
...
printf("out_of_range:_%i\n", histo_out_of_range);
total+=histo_out_of_range;
histo_out_of_range=0;
printf("total:_%lli", total);

```

The final line allows the user to compare the number of events collected in the histogram to the number of events expected. It is a computationally cheap additional check for data consistency.

To distinguish different types of output, such as moments of different observables, data should be **tagged** by short keys that are easily filtered out in post processing. For example, if every line containing moments of avalanche sizes is tagged by `#M_SIZE` at the beginning, all relevant lines can be extracted very easily for example using `grep '^#M_SIZE' output.txt`. To strip off the tags, one either appends `|sed 's/#M_SIZE//'` or includes the functionality of `grep` in the `sed` command,

```
sed -n 's/^#M_SIZE//p' output.txt > output.txt.M_SIZE
```

storing all relevant lines in `output.txt.M_SIZE` for further processing by other tools. One very simple, but particularly powerful one is `awk`. For example, the average across the seventh column starting with the 101st chunk (stored in the first column) can be calculated using

```
awk ' { if ($1>100) {m0++; m1+= $7;} } END { printf ("%i %10.20g\n",
↳    m0, m1/m0); } ' output.txt.M_SIZE
```

All of this is very easily automated using powerful scripting languages (in particular shell scripts, `awk`, `sed` and `grep`), and more powerful (interpreted) programming languages, such as `perl` or `python`, which provide easy access to line-oriented data. In recent years, XML has become more popular to store simulation parameters as well as simulation results.

7.4 Summary and conclusion

The early life of SOC was all about computer models that showed the desired features of SOC: Intermittent behaviour (slow drive, fast relaxation) displaying scale invariance as observed in traditional critical phenomena without the need to tune a control parameter to a critical value. After many authors had (mostly with little success) attempted to populate the universality class of the BTW Sandpile, a range of SOC models was proposed firstly as a paradigm of alternative universality classes and later to highlight specific aspects of SOC, such as non-conservation (as for example in the Forest-Fire Model), non-Abelianness (as for example in the Olami-Feder-Christensen Model) and stochasticity (as for example in the Manna Model).

Many of these models have been studied extensively, accumulating hundreds of thousands of hours of CPU time in large-scale Monte Carlo simulations. A finite size scaling analysis of the data generally produces a set of two to eight exponents, which are supposedly universal. It turns out, however, that very few models display clean, robust scaling behaviour in the event size distribution, although it is remarkably broad for many models.

Of the models discussed above, the Manna Model displays the clearest signs of scale invariance. There is wide consensus that it is the same universality class as the Oslo Model (Christensen *et al.* 1996, Nakanishi and Sneppen 1997). In the conservative limit and in the near-conservative regime, the Olami-Feder-Christensen Model also displays convincing moment scaling, but less so for smaller values of the level of conservation. Numerical artefacts may play a significant rôle in its scaling (Drossel 2002).

The Forest Fire Models is widely acknowledged for failing to display finite size scaling in the event size distribution (Grassberger 2002, Pruessner and Jensen 2002a), although its moments still display some scaling (Pruessner and Jensen 2004). The contrast is even sharper in the Bak-Tang-Wiesenfeld Model: Some scaling is known analytically (Majumdar and Dhar 1992, Ivashkevich 1994, Ivashkevich *et al.* 1994, Dhar and Manna 1994), yet the event size distribution seems at best be governed by multiscaling (Tebaldi *et al.* 1999, Drossel 1999, 2000, Dorn *et al.* 2001)

While analytical approaches receive increasing attention, numerical techniques remain indispensable in the development and analysis of models which are tailor-made to display specific features or to mimic experimental systems. Models developed more recently are usually implemented in C, producing numerical data in Monte-Carlo simulations. It is fair to say that the careful data analysis requires as much attention to detail as the implementation of the model in the first place.

While the classic data-collapse and more immediate tests for scaling dominated the early literature of SOC, more recently the finite size scaling of moments (Tebaldi *et al.* 1999) has become the predominant technique for the extraction of scaling exponents. Apart from identifying the mechanism of SOC, the main purpose of the numerical work is to establish universality and universality classes among models, as well as their relation to natural phenomena. One may hope that these efforts will eventually help to uncover the necessary and sufficient conditions for SOC.

Acknowledgements The author gratefully acknowledges the kind support by EPSRC Mathematics Platform grant EP/I019111/1.

Appendix: Implementation details for binning

To implement binning in computer simulations of SOC models it is advisable to perform simple bit manipulations on basic, integer-valued observables. It often suffices to implemented three levels of coarse graining or less, for example

```
#define SMALL2MEDIUM.THRESHOLD (1LL<<15)
long long histo_small[SMALL2MEDIUM.THRESHOLD]={0LL};
#define MEDIUM2LARGE.THRESHOLD (1LL<<30)
```

```

#define MEDIUM_SHIFT (12)
long histo_medium[(MEDIUM2LARGE_THRESHOLD-
↳ SMALL2MEDIUM_THRESHOLD)>>MEDIUM_SHIFT]={0L};
#define LARGE_THRESHOLD (1LL<<45)
#define LARGE_SHIFT (27)
int histo_large[(LARGE_THRESHOLD-MEDIUM2LARGE_THRESHOLD)
↳ >>LARGE_SHIFT]={0};
int histo_out_of_range=0;
long long int s; /* event size */

...

if (s<SMALL2MEDIUM_THRESHOLD) histo_small[s]++;
else if (s<MEDIUM2LARGE_THRESHOLD) histo_medium[(s-
↳ SMALL2MEDIUM_THRESHOLD)>>MEDIUM_SHIFT]++;
else if (s<LARGE_THRESHOLD) histo_large[(s-
↳ MEDIUM2LARGE_THRESHOLD)>>LARGE_SHIFT]++;
else histo_out_of_range++;

```

Here the event size to be tallied is `s`. In the block of `if` statements, it is compared to various thresholds before it is rescaled and counted into a histogram. Because vectors in many programming languages start with index 0, a shift an offset is subtracted as well. It can pay off to re-arrange the `if` statements as to test against the most frequent case as early as possible. One case, in the present example the last one, counts the number of times the counter overflows, here `histo_out_of_range`.

Some subtleties of the above implementation are worth discussing. Firstly, the types used for the histogram typically decrease in size with increasing event size while the size of the type needed to represent the event size at the respective thresholds increases. This is because normally the frequency is an *inverse* power law of the event size. Great care must be taken to avoid unnecessary typecasts and undesired outcomes, as some languages, in particular C, are rather idiosyncratic when it comes to (integer) type-promotion in comparisons, in particular when they involve signs.

In the above examples, automatic vector variables are used and initialised by assigning `{0}`, which is expanded by the compiler to a suitable size by adding zeroes. Initialisation of vectors in C has been further simplified in the C99 standard.

Secondly, it is important to choose the thresholds together with the planned bit-shifts, in order to avoid an *off-by-one* error. The problem is that, say,

`s<MEDIUM2LARGE_THRESHOLD`, does not imply

$$\begin{aligned}
 & (s - \text{SMALL2MEDIUM_THRESHOLD}) / ((1 \ll \text{MEDIUM_SHIFT}) < (\\
 \hookrightarrow & \text{MEDIUM2LARGE_THRESHOLD} - \text{SMALL2MEDIUM_THRESHOLD}) \\
 \hookrightarrow & / (1 \ll \text{MEDIUM_SHIFT})
 \end{aligned}$$

because for some `s<MEDIUM2LARGE_THRESHOLD` their bitshifted value `s>>MEDIUM_SHIFT` in fact equals `MEDIUM2LARGE_THRESHOLD>>MEDIUM_SHIFT`, namely precisely when `MEDIUM2LARGE_THRESHOLD` is not an integer multiple of

1<<MEDIUM_SHIFT. It is therefore a matter of defensive programming to write the thresholds for the macros in this form:

```
#define MEDIUM2LARGE_THRESHOLD ((1LL<<18) * (1LL<<
↳ MEDIUM_SHIFT))
```

As for a rudimentary output routine

```
for (i=0; i<SMALL2MEDIUM_THRESHOLD; i++)
  if (histo_small[i]) printf("%i_%i_%li_%i\n", i, i, histo_small[i], 1);
for (i=0; i<((MEDIUM2LARGE_THRESHOLD-
↳ SMALL2MEDIUM_THRESHOLD)>>MEDIUM_SHIFT); i++)
  if (histo_medium[i]) printf("%i_%i_%li_%i\n", ((long)
↳ SMALL2MEDIUM_THRESHOLD)+(((long)(i))<<MEDIUM_SHIFT),
↳ i, histo_medium[i], 1<<MEDIUM_SHIFT);
for (i=0; i<((LARGE_THRESHOLD-MEDIUM2LARGE_THRESHOLD)
↳ >>LARGE_SHIFT); i++)
  if (histo_large[i]) printf("%li_%i_%i_%i\n", ((long long)
↳ MEDIUM2LARGE_THRESHOLD)+(((long long)(i))<<
↳ LARGE_SHIFT), i, histo_large[i], 1<<LARGE_SHIFT);
printf("out_of_range:_%i\n", histo_out_of_range);
```

care must again be taken that the formatting of the output is in line with the type of the data and does not spoil it. Fortunately, most modern compilers spot clashes between the formatting string used in `printf` and the actual argument. As discussed below, it is generally advisable to have only one output stream, namely `stdout`, and to use `tags` to mark up data for easy fetching by post-processing tools. In the example above, the bins have not been rescaled by their size which instead has been included explicitly in the output. A sample of the PDF can be inspected by plotting the third column divided by the fourth against the first.

References

- Alava, M., and M. A. Muñoz, 2002, Phys. Rev. E **65**(2), 026145 (pages 8).
 Alava, M. J., L. Laurson, A. Vespignani, and S. Zapperi, 2008, Phys. Rev. E **77**(4), 048101 (pages 2), comment on (Pruessner and Peters 2006), reply (Pruessner and Peters 2008).
 Bagnoli, F., F. Cecconi, A. Flammini, and A. Vespignani, 2003, Europhys. Lett. **63**(4), 512.
 Bak, P., K. Chen, and M. Creutz, 1989a, Nature **342**(6251), 780.
 Bak, P., K. Chen, and C. Tang, 1990, Phys. Lett. A **147**(5–6), 297.
 Bak, P., K. Christensen, L. Danon, and T. Scanlon, 2002, Phys. Rev. Lett. **88**(17), 178501 (pages 4).
 Bak, P., and K. Sneppen, 1993, Phys. Rev. Lett. **71**(24), 4083.
 Bak, P., and C. Tang, 1989, J. Geophys. Res. **94**(B11), 15635.
 Bak, P., C. Tang, and K. Wiesenfeld, 1987, Phys. Rev. Lett. **59**(4), 381.
 Bak, P., C. Tang, and K. Wiesenfeld, 1989b, in *Cooperative Dynamics in Complex Physical Systems, Proceedings of the Second Yukawa International Symposium, Kyoto, Japan, August 24–27, 1988*, edited by H. Takayama (Springer-Verlag, Berlin, Germany), volume 43 of *Springer Series in Synergetics*, pp. 274–279.

- Basu, M., U. Basu, S. Bondyopadhyay, P. K. Mohanty, and H. Hinrichsen, 2012, *Phys. Rev. Lett.* **109**, 015702.
- Ben-Hur, A., and O. Biham, 1996, *Phys. Rev. E* **53**(2), R1317.
- Berg, B. A., 1992, *Comp. Phys. Comm.* **69**(1), 7.
- Berg, B. A., 2004, *Markov Chain Monte Carlo Simulations and Their Statistical Analysis* (World Scientific, Singapore).
- Biham, O., E. Milshtein, and O. Malcai, 2001, *Phys. Rev. E* **63**(6), 061309 (pages 8).
- Binder, K., and D. W. Heermann, 1997, *Monte Carlo Simulation in Statistical Physics* (Springer-Verlag, Berlin, Germany), 3rd edition.
- Bonachela, J. A., 2008, *Universality in Self-Organized Criticality*, Ph.D. thesis, Departamento de Electromagnetismo y Física de la Materia & Institute Carlos I for Theoretical and Computational Physics, University of Granada, Granada, Spain, accessed 12 Sep 2009, URL <http://hera.ugr.es/tesisugr/17706312.pdf>.
- Bonachela, J. A., and M. A. Muñoz, 2007, *Physica A* **384**(1), 89, ISSN 0378-4371, proceedings of the International Conference on Statistical Physics, Raichak and Kolkata, India, Jan 5–9, 2007.
- Bonachela, J. A., and M. A. Muñoz, 2009, *J. Stat. Mech.* **2009**(09), P09009 (pages 37).
- Brandt, S., 1998, *Data Analysis* (Springer-Verlag, Berlin, Germany).
- Bröker, H.-M., and P. Grassberger, 1999, *Physica A* **267**(3–4), 453, ISSN 0378-4371.
- Bruinsma, R., and G. Aeppli, 1984, *Phys. Rev. Lett.* **52**(17), 1547.
- Bunzarova, N. Z., 2010, *Phys. Rev. E* **82**(3), 031116 (pages 14).
- Burridge, R., and L. Knopoff, 1967, *Bull. Seismol. Soc. Am.* **57**(3), 341.
- Carlson, J. M., J. S. Langer, and B. E. Shaw, 1994, *Rev. Mod. Phys.* **66**(2), 657.
- Casartelli, M., L. Dall’Asta, A. Vezzani, and P. Vivo, 2006, *Eur. Phys. J. B* **52**(1), 91.
- Černák, J., 2002, *Phys. Rev. E* **65**(4), 046141 (pages 6).
- Ceva, H., 1995, *Phys. Rev. E* **52**(1), 154.
- Ceva, H., 1998, *Phys. Lett. A* **245**(5), 413.
- Chessa, A., H. E. Stanley, A. Vespignani, and S. Zapperi, 1999a, *Phys. Rev. E* **59**(1), R12, numerics may not be independent from (Chessa *et al.* 1999b).
- Chessa, A., A. Vespignani, and S. Zapperi, 1999b, *Comp. Phys. Comm.* **121–122**, 299, proceedings of the Europhysics Conference on Computational Physics CCP 1998, Granada, Spain, Sep 2–5, 1998. Numerics presented may not be independent from that in (Chessa *et al.* 1999a).
- Chianca, C. V., J. S. Sá Martins, and P. M. C. de Oliveira, 2009, *Eur. Phys. J. B* **68**(4), 549.
- Christensen, K., 1992, *Self-Organization in Models of Sandpiles, Earthquakes, and Flashing Fireflies*, Ph.D. thesis, Institute of Physics and Astronomy, University of Aarhus, DK-8000 Aarhus C, Denmark.
- Christensen, K., 1993, *Phys. Rev. Lett.* **71**(8), 1289, reply to comment (Klein and Rundle 1993).
- Christensen, K., Á. Corral, V. Frette, J. Feder, and T. Jøssang, 1996, *Phys. Rev. Lett.* **77**(1), 107.
- Christensen, K., N. Farid, G. Pruessner, and M. Stapleton, 2008, *Eur. Phys. J. B* **62**(3), 331.
- Christensen, K., H. C. Fogedby, and H. J. Jensen, 1991, *J. Stat. Phys.* **63**(3/4), 653.
- Christensen, K., and N. R. Moloney, 2005, *Complexity and Criticality* (Imperial College Press, London, UK).
- Christensen, K., and Z. Olami, 1992, *Phys. Rev. A* **46**(4), 1829.
- Christensen, K., and Z. Olami, 1993, *Phys. Rev. E* **48**(5), 3361.
- Clar, S., B. Drossel, and F. Schwabl, 1996, *J. Phys.: Condens. Matter* **8**(37), 6803.
- Clauset, A., C. R. Shalizi, and M. E. J. Newman, 2009, *SIAM Rev.* **51**(4), 661.
- Cormen, T. H., C. E. Leiserson, and R. L. Rivest, 1996, *Introduction to Algorithms* (The MIT Press, Cambridge, MA, USA).

- Corral, Á., 2003, Phys. Rev. E **68**(3), 035102(R) (pages 4).
- Corral, Á., 2004a, Phys. Rev. E **69**(2), 026107 (pages 12).
- Corral, Á., 2004b, Phys. Rev. Lett. **92**(10), 108501 (pages 4).
- Corral, Á., 2004c, Physica A **340**(4), 590.
- Corral, Á., and K. Christensen, 2006, Phys. Rev. Lett. **96**(10), 109801 (pages 1), comment on (Lindman *et al.* 2005), reply (Lindman *et al.* 2006).
- Corral, Á., and M. Paczuski, 1999, Phys. Rev. Lett. **83**(3), 572.
- Creutz, M., 2004, Physica A **340**(4), 521, ISSN 0378-4371, proceedings of the symposium Complexity and Criticality: in memory of Per Bak (1947–2002), Copenhagen, Denmark, Aug 21–23, 2003.
- Daidsen, J., and M. Paczuski, 2005, Phys. Rev. Lett. **94**(4), 048501 (pages 4), comment (Werner and Sornette 2007).
- Daidsen, J., and M. Paczuski, 2007, Phys. Rev. Lett. **99**(17), 179802, reply to comment (Werner and Sornette 2007).
- De Menech, M., and A. L. Stella, 2000, Phys. Rev. E **62**(4), R4528.
- De Menech, M., A. L. Stella, and C. Tebaldi, 1998, Phys. Rev. E **58**(3), R2677.
- de Sousa Vieira, M., 1992, Phys. Rev. A **46**(10), 6288.
- Dhar, D., 1990, Phys. Rev. Lett. **64**(14), 1613.
- Dhar, D., 1999a, Physica A **263**(1–4), 4, proceedings of the 20th IUPAP International Conference on Statistical Physics, Paris, France, Jul 20–24, 1998, overlaps with (Dhar 1999b).
- Dhar, D., 1999b, Studying self-organized criticality with exactly solved models, arXiv:cond-mat/9909009, arXiv:cond-mat/9909009.
- Dhar, D., 2004, Physica A **340**(4), 535.
- Dhar, D., 2006, Physica A **369**(1), 29, proceedings of the 11th International Summer-school on 'Fundamental Problems in Statistical Physics', Leuven, Belgium, Sep 4 – 17, 2005; updated from (Dhar 1999b).
- Dhar, D., and S. S. Manna, 1994, Phys. Rev. E **49**(4), 2684.
- Dhar, D., and P. Pradhan, 2004, J. Stat. Mech. **2004**(05), P05002 (pages 12), includes erratum.
- Dhar, D., and R. Ramaswamy, 1989, Phys. Rev. Lett. **63**(16), 1659.
- Dhar, D., P. Ruelle, S. Sen, and D.-N. Verma, 1995, J. Phys. A: Math. Gen. **28**(4), 805.
- Dickman, R., and J. M. M. Campelo, 2003, Phys. Rev. E **67**(6), 066111 (pages 5).
- Dickman, R., M. A. Muñoz, A. Vespignani, and S. Zapperi, 2000, Braz. J. Phys. **30**(1), 27.
- Dickman, R., T. Tomé, and M. J. de Oliveira, 2002, Phys. Rev. E **66**(1), 016111 (pages 8).
- Dickman, R., A. Vespignani, and S. Zapperi, 1998, Phys. Rev. E **57**(5), 5095.
- Dorn, P. L., D. S. Hughes, and K. Christensen, 2001, On the avalanche size distribution in the btw model, preprint from http://www.cmth.ph.ic.ac.uk/kim/papers/preprints/preprint_btw.pdf, accessed 19 Oct 2010.
- Dowd, K., and C. Severance, 1998, *High Performance Computing* (O'Reilly, Sebastopol, CA, USA), 2nd edition.
- Drossel, B., 1999, An alternative view of the Abelian sandpile model, arXiv:cond-mat/9904075v1.
- Drossel, B., 2000, Phys. Rev. E **61**(3), R2168.
- Drossel, B., 2002, Phys. Rev. Lett. **89**(23), 238701 (pages 4).
- Drossel, B., and F. Schwabl, 1992a, Phys. Rev. Lett. **69**(11), 1629, largely identical to proceedings article (Drossel and Schwabl 1992b).
- Drossel, B., and F. Schwabl, 1992b, Physica A **191**(1–4), 47, proceedings of the International Conference on Fractals and Disordered Systems, Hamburg, Germany, Jul 29–31, 1992.

- Efron, B., 1982, *The Jackknife, the Bootstrap and Other Resampling Plans* (SIAM, Philadelphia, PA, USA).
- Feller, W., 1966, *An Introduction to Probability Theory and its Applications*, volume II (John Wiley & Sons, New York, NY, USA).
- Flyvbjerg, H., and H. G. Petersen, 1989, *J. Chem. Phys.* **91**(1), 461.
- Galassi, M., J. Davies, J. Theiler, B. Gough, G. Jungman, P. Alken, M. Booth, and F. Rossi, 2009, *GNU Scientific Library Reference Manual* (Network Theory Ltd.), 3rd (v1.12) edition, <http://www.network-theory.co.uk/gsl/manual/>, accessed 18 Aug 2009.
- Gentle, J. E., 1998, *Random Number Generation and Monte Carlo Methods* (Springer-Verlag, Berlin, Germany).
- Grassberger, P., 1982, *Z. Phys. B* **47**, 365.
- Grassberger, P., 1994, *Phys. Rev. E* **49**(3), 2436.
- Grassberger, P., 1995, *Phys. Lett. A* **200**(3-4), 277.
- Grassberger, P., 2002, *New J. Phys.* **4**(1), 17 (pages 15).
- Grassberger, P., and H. Kantz, 1991, *J. Stat. Phys.* **63**(3-4), 685.
- Grinstein, G., D.-H. Lee, and S. Sachdev, 1990, *Phys. Rev. Lett.* **64**(16), 1927.
- Gutenberg, B., and C. F. Richter, 1954, *Seismicity of the earth and associated phenomena* (Princeton University Press, Princeton, NJ, USA), 2nd edition.
- Henkel, M., H. Hinrichsen, and S. Lübeck, 2008, *Non-Equilibrium Phase Transitions* (Springer-Verlag, Berlin, Germany).
- Henley, C. L., 1989, *Bull. Am. Phys. Soc.* **34**(3), 838, abstract of talk M18.2, 23 March 1989, of the 1989 March Meeting of The American Physical Society, St. Louis, MO, USA, Mar 20-24, 1989.
- Henley, C. L., 1993, *Phys. Rev. Lett.* **71**(17), 2741.
- Hergarten, S., 2002, *Self-Organized Criticality in Earth Systems* (Springer-Verlag, Berlin, Germany).
- Hinrichsen, H., 2000, *Adv. Phys.* **49**, 815.
- Honecker, A., and I. Peschel, 1997, *Physica A* **239**(4), 509.
- Hughes, D., and M. Paczuski, 2002, *Phys. Rev. Lett.* **88**(5), 054302 (pages 4).
- Huynh, H. N., and G. Pruessner, 2012, *Phys. Rev. E* **85**, 061133.
- Huynh, H. N., G. Pruessner, and L. Y. Chew, 2011, *J. Stat. Mech.* **2011**(09), P09024.
- Hwa, T., and M. Kardar, 1989a, *Phys. Rev. Lett.* **62**(16), 1813, identical to proceedings article (Hwa and Kardar 1989b).
- Hwa, T., and M. Kardar, 1989b, *Physica D* **38**(1-3), 198, identical to (Hwa and Kardar 1989a); Proceedings of a conference held in honour of Benoit B. Mandelbrot's 65th birthday, Les Mas d'Artigny (Vence), France, Oct 1 - 4, 1989.
- Hwa, T., and M. Kardar, 1992, *Phys. Rev. A* **45**(10), 7002.
- Ito, K., and M. Matsuzaki, 1990, *J. Geophys. Res.* **95**(B5), 6853.
- Ivashkevich, E. V., 1994, *J. Phys. A: Math. Gen.* **27**(11), 3643.
- Ivashkevich, E. V., D. V. Ktitarev, and V. B. Priezzhev, 1994, *Physica A* **209**(3-4), 347.
- Janssen, H. K., 1981, *Z. Phys. B* **42**, 151.
- Jeng, M., 2005, *Phys. Rev. E* **71**(3), 036153 (pages 17).
- Jensen, H. J., 1998, *Self-Organized Criticality* (Cambridge University Press, New York, NY, USA).
- Jo, H.-H., and M. Ha, 2008, *Phys. Rev. Lett.* **101**(21), 218001 (pages 4).
- Jones, D., 2012, Good practice in (pseudo) random number generation for bioinformatics applications, as of 7 May 2010, available from <http://www0.cs.ucl.ac.uk/staff/d.jones/GoodPracticeRNG.pdf>, accessed 21 Aug 2012.
- Juanico, D. E., C. Monterola, and C. Saloma, 2007, *Phys. Rev. E* **75**(4), 045105(R) (pages 4).
- Kadanoff, L. P., S. R. Nagel, L. Wu, and S.-m. Zhou, 1989, *Phys. Rev. A* **39**(12), 6524.

- Kaneko, K., 1989, *Physica D* **37**(1–3), 60.
- Karmakar, R., S. S. Manna, and A. L. Stella, 2005, *Phys. Rev. Lett.* **94**(8), 088002 (pages 4).
- Kernighan, B. W., and R. Pike, 2002, *The Practice of Programming* (Addison-Wesley, Boston, MA, USA).
- Kernighan, B. W., and D. M. Ritchie, 1988, *The C programming language* (Prentice Hall, Englewood Cliffs, NJ), 2nd edition.
- Klein, W., and J. Rundle, 1993, *Phys. Rev. Lett.* **71**(8), 1288, comment on (Olami *et al.* 1992), reply (Christensen 1993).
- Kloster, M., S. Maslov, and C. Tang, 2001, *Phys. Rev. E* **63**(2), 026111 (pages 4).
- Knuth, D. E., 1997, *The Art of Computer Programming Volumes 1–3* (Addison-Wesley, Reading, MA, USA).
- Koplik, J., and H. Levine, 1985, *Phys. Rev. B* **32**(1), 280.
- Landau, D. P., and K. Binder, 2005, *A Guide to Monte Carlo Simulations in Statistical Physics* (Cambridge University Press, Cambridge, UK), 2nd edition.
- Lauritsen, K. B., S. Zapperi, and H. E. Stanley, 1996, *Phys. Rev. E* **54**(3), 2483.
- Le Doussal, P., K. J. Wiese, and P. Chauve, 2002, *Phys. Rev. B* **66**(17), 174201 (pages 34).
- Lebowitz, J. L., C. Maes, and E. R. Speer, 1990, *J. Stat. Phys.* **59**, 117, ISSN 0022-4715.
- Leschhorn, H., T. Nattermann, S. Stepanow, and L.-H. Tang, 1997, *Ann. Physik* **6**, 1.
- Liggett, T. M., 2005, *Stochastic Interacting Systems: Contact, Voter and Exclusion Processes* (Springer-Verlag, Berlin, Germany).
- Lin, C.-Y., and C.-K. Hu, 2002, *Phys. Rev. E* **66**(2), 021307 (pages 12).
- Lindman, M., K. Jónsdóttir, R. Roberts, B. Lund, and R. Bødvarsson, 2005, *Phys. Rev. Lett.* **94**(10), 108501 (pages 4), comment (Corral and Christensen 2006).
- Lindman, M., K. Jónsdóttir, R. Roberts, B. Lund, and R. Bødvarsson, 2006, *Phys. Rev. Lett.* **96**(10), 109802 (pages 1), reply to comment (Corral and Christensen 2006).
- Lübeck, S., 2000, *Phys. Rev. E* **61**(1), 204.
- Lübeck, S., and P. C. Heger, 2003a, *Phys. Rev. E* **68**(5), 056102 (pages 11).
- Lübeck, S., and P. C. Heger, 2003b, *Phys. Rev. Lett.* **90**(23), 230601 (pages 4).
- Lübeck, S., and K. D. Usadel, 1997, *Phys. Rev. E* **55**(4), 4095.
- Mahieu, S., and P. Ruelle, 2001, *Phys. Rev. E* **64**(6), 066130 (pages 19).
- Majumdar, S. N., and D. Dhar, 1992, *Physica A* **185**(1–4), 129.
- Manna, S. S., 1990, *J. Stat. Phys.* **59**(1/2), 509.
- Manna, S. S., 1991, *J. Phys. A: Math. Gen.* **24**(7), L363.
- Marsaglia, G., 1999, Random numbers for c: The end?, newsgroup posting 20 Jan 1999, 9:00am, accessed 10 May 2011, URL <https://groups.google.com/group/sci.math.num-analysis/msg/eb4ddde782b17051>.
- Marsaglia, G., 2005, Mersenne twister, newsgroup posting 14 Jul 2005, 12:04am, available from http://groups.google.com/group/sci.crypt/browse/_thread/thread/305c507efbe85be4, accessed 10 May 2011.
- Matsumoto, M., 2008, Mersenne twister home page, available from <http://www.math.sci.hiroshima-u.ac.jp/~m-mat/MT/mt.html>, accessed 9 Oct 2008.
- Matsumoto, M., and T. Nishimura, 1998, *ACM Trans. Model. Comp. Sim.* **8**(1), 3.
- Middleton, A. A., and C. Tang, 1995, *Phys. Rev. Lett.* **74**(5), 742.
- Milshstein, E., O. Biham, and S. Solomon, 1998, *Phys. Rev. E* **58**(1), 303.
- Mousseau, N., 1996, *Phys. Rev. Lett.* **77**(5), 968.
- Muñoz, M. A., R. Dickman, A. Vespignani, and S. Zapperi, 1999, *Phys. Rev. E* **59**(5), 6175.
- Nakanishi, H., and K. Sneppen, 1997, *Phys. Rev. E* **55**(4), 4012.

- Nattermann, T., S. Stepanow, L.-H. Tang, and H. Leschhorn, 1992, *J. Phys. II (France)* **2**, 1483.
- Newman, M. E. J., and G. T. Barkema, 1999, *Monte Carlo Methods in Statistical Physics* (Oxford University Press, New York, NY, USA).
- Olami, Z., H. J. S. Feder, and K. Christensen, 1992, *Phys. Rev. Lett.* **68**(8), 1244, comment (Klein and Rundle 1993).
- Paczuski, M., and K. E. Bassler, 2000a, *Phys. Rev. E* **62**(4), 5347, according to (Hughes and Paczuski 2002), this(?) version "was published by mistake. The correct version is [...] arXiv:cond-mat/0005340". This probably means arXiv:cond-mat/0005340v2.
- Paczuski, M., and K. E. Bassler, 2000b, Theoretical results for sandpile models of soc with multiple topplings, according to (Hughes and Paczuski 2002), this is probably the correct version of (Paczuski and Bassler 2000a)., [arXiv:cond-mat/0005340v2](https://arxiv.org/abs/cond-mat/0005340v2).
- Paczuski, M., and S. Boettcher, 1996, *Phys. Rev. Lett.* **77**(1), 111, .
- Pan, G.-J., D.-M. Zhang, Z.-H. Li, H.-Z. Sun, and Y.-P. Ying, 2005, *Phys. Lett. A* **338**(3–5), 163, ISSN 0375-9601.
- Pastor-Satorras, R., and A. Vespignani, 2000a, *Phys. Rev. E* **62**(5), 6195, see (Pastor-Satorras and Vespignani 2000b).
- Pastor-Satorras, R., and A. Vespignani, 2000b, *J. Phys. A: Math. Gen.* **33**(3), L33, see (Pastor-Satorras and Vespignani 2000a).
- Patzlaff, H., and S. Trimper, 1994, *Phys. Lett. A* **189**(3), 187.
- Pradhan, P., and D. Dhar, 2006, *Phys. Rev. E* **73**(2), 021303 (pages 12).
- Pradhan, P., and D. Dhar, 2007, *J. Phys. A: Math. Theor.* **40**(11), 2639.
- Press, W. H., S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, 2007, *Numerical Recipes* (Cambridge University Press, Cambridge, UK), 3rd edition.
- Priezzhev, V. B., D. V. Kvitarev, and E. V. Ivashkevich, 1996, *Phys. Rev. Lett.* **76**(12), 2093.
- Privman, V., P. C. Hohenberg, and A. Aharony, 1991, in *Phase Transitions and Critical Phenomena*, edited by C. Domb and J. L. Lebowitz (Academic Press, New York, NY, USA), volume 14, chapter 1, pp. 1–134.
- Pruessner, G., 2003, *Phys. Rev. E* **67**(3), 030301(R) (pages 4).
- Pruessner, G., 2012a, The average avalanche size in the manna model and other models of self-organised criticality, [arXiv:1208.2069](https://arxiv.org/abs/1208.2069), [arXiv:1208.2069](https://arxiv.org/abs/1208.2069).
- Pruessner, G., 2012b, The field theory of self-organised criticality i: General structure and tree level, in preparation.
- Pruessner, G., 2012c, *Self-Organised Criticality* (Cambridge University Press, Cambridge, UK).
- Pruessner, G., and H. J. Jensen, 2002a, *Phys. Rev. E* **65**(5), 056707 (pages 8).
- Pruessner, G., and H. J. Jensen, 2002b, *Europhys. Lett.* **58**(2), 250.
- Pruessner, G., and H. J. Jensen, 2004, *Phys. Rev. E* **70**(6), 066707 (pages 25).
- Pruessner, G., and O. Peters, 2006, *Phys. Rev. E* **73**(2), 025106(R) (pages 4), comment (Alava *et al.* 2008).
- Pruessner, G., and O. Peters, 2008, *Phys. Rev. E* **77**(4), 048102 (pages 2), reply to comment (Alava *et al.* 2008).
- Ramasco, J. J., M. A. Muñoz, and C. A. da Silva Santos, 2004, *Phys. Rev. E* **69**(4), 045105(R) (pages 4).
- Rose, G., 2011, Kiss: A bit too simple, preprint from <http://eprint.iacr.org/2011/007.pdf> accessed 20 Aug 2012.
- Rossi, M., R. Pastor-Satorras, and A. Vespignani, 2000, *Phys. Rev. Lett.* **85**(9), 1803.
- Ruelle, P., 2002, *Phys. Lett. B* **539**(1), 172.
- Ruelle, P., and S. Sen, 1992, *J. Phys. A: Math. Gen.* **25**(22), L1257.
- Salas, J., and A. D. Sokal, 2000, *J. Stat. Phys.* **98**(3–4), 551.
- Socolar, J. E. S., G. Grinstein, and C. Jayaprakash, 1993, *Phys. Rev. E* **47**(4), 2366.
- Sornette, A., and D. Sornette, 1989, *Europhys. Lett.* **9**(3), 197.

- Sornette, D., and M. J. Werner, 2009, in *Encyclopedia of Complexity and Systems Science*, edited by R. A. Meyers (Springer-Verlag, New York, NY, USA), volume 9, pp. 7872–7891, [arXiv:0803.3756v2](#).
- Stanley, H. E., 1971, *Introduction to Phase Transitions and Critical Phenomena* (Oxford University Press, New York, NY, USA).
- Stauffer, D., and A. Aharony, 1994, *Introduction to Percolation Theory* (Taylor & Francis, London, UK).
- Tadić, B., and D. Dhar, 1997, *Phys. Rev. Lett.* **79**(8), 1519.
- Tang, C., and P. Bak, 1988, *Phys. Rev. Lett.* **60**(23), 2347.
- Tebaldi, C., M. De Menech, and A. L. Stella, 1999, *Phys. Rev. Lett.* **83**(19), 3952.
- Torvund, F., and J. Frøyland, 1995, *Phys. Scripta* **52**, 624.
- Turcotte, D. L., 1993, *Fractals and chaos in geology and geophysics* (Cambridge University Press, Cambridge, UK).
- Turcotte, D. L., 1999, *Rep. Prog. Phys.* **62**, 1377.
- van der Linden, P., 1994, *Expert C Programming* (Sunsoft Press, A Prentice Hall Title, Mountain View, CA, USA).
- van Wijland, F., 2002, *Phys. Rev. Lett.* **89**(19), 190602 (pages 4).
- Černák, J., 2006, *Phys. Rev. E* **73**, 066125.
- Vespignani, A., R. Dickman, M. A. Muñoz, and S. Zapperi, 1998, *Phys. Rev. Lett.* **81**(25), 5676.
- Vespignani, A., R. Dickman, M. A. Muñoz, and S. Zapperi, 2000, *Phys. Rev. E* **62**(4), 4564.
- Vespignani, A., and S. Zapperi, 1995, *Phys. Rev. E* **51**(3), 1711, see (Lin and Hu 2002).
- Wegner, F. J., 1972, *Phys. Rev. B* **5**(11), 4529.
- Welinder, P., G. Pruessner, and K. Christensen, 2007, *New J. Phys.* **9**(5), 149 (pages 18).
- Werner, M. J., and D. Sornette, 2007, *Phys. Rev. Lett.* **99**(17), 179801, comment on (Davidsen and Paczuski 2005), reply (Davidsen and Paczuski 2007).
- Zapperi, S., K. B. Lauritsen, and H. E. Stanley, 1995, *Phys. Rev. Lett.* **75**(22), 4071.
- Zhang, Y.-C., 1989, *Phys. Rev. Lett.* **63**(5), 470.

