

Self-Organized Criticality in a Slope Sandpile: Universality, Geometry, and the Scaling of a Filamentary Avalanche Front

Nathan Langley

University of North Carolina at Greensboro

PHY 351 — Independent Summer Research Advisor: Prof. Ian Beatty

June 2026

Abstract

I present a computational study of the continuous-height slope sandpile from Charbonneau’s *Natural Complexity* (Chapter 5), carried from a validated implementation through to a complete scaling theory of its avalanches. After reproducing the chapter’s signatures in one dimension (exact bulk conservation, a stationary angle of repose, and N -independent power-law avalanche statistics), I clarify one point the textbook conflates: the *mean* slope sits about 16% below the critical slope and stays there for any grain size, while only the *peak* slope approaches the threshold as forcing becomes fine. I then build the 2-D generalization as the identical bond rule in both lattice directions, which lets a 1-D-versus-2-D exponent difference be read as a clean effect of dimension alone. Two universality results follow: the slope model’s energy exponent changes with dimension ($\tau_E = 1.03$ in 1-D, 0.87 in 2-D under the same local rule), and the gradient rule is in a different universality class from the canonical Bak-Tang-Wiesenfeld (BTW) sandpile (size exponent 0.89 versus 1.14). Bulk conservation is shown to be necessary for true scale invariance: any dissipation imposes a characteristic avalanche size and the cutoff stops scaling with system size, in both 1-D and 2-D. A list-based, compiled engine (validated bit-for-bit against the reference implementation) then unlocks the larger lattices the central question needs.

The core of the report is a scaling theory of the 2-D avalanche, built from the moment spectrum and conditional exponents — the field’s standard discriminators between simple finite-size scaling and multifractality. Validated first on BTW, the analysis shows the 2-D slope avalanche is a **ballistic, filamentary, multiply-swept front**: its footprint is a constant-width thin filament of mass-radius dimension $D \sim 1$ (thinner than the exactly-solvable directed sandpile at $3/2$, far from compact BTW at 2), it propagates one cell per time step, and it sweeps each footprint bond about as many times as the footprint is long. A tunable stochastic split proves the filament is *caused* by the deterministic rule (its dimension climbs to 1.87 as stochasticity turns on), and the same geometry measured in the model’s native 1-D shows the 2-D filament is literally a 1-D-like object embedded in the plane. The single-scale theory built from this front predicts the entire *spatial* exponent set exactly in the thermodynamic limit, with one quantified exception: duration is a loose proxy for spatial extent, and that looseness does not heal but grows to a true residual of order 0.2. Finally, the avalanche-area distribution is shown to be multifractal asymptotically, not as a finite-size artifact — which reconciles with the single-scale result once one distinguishes the conditional *mean* (single-scale, the typical avalanche) from the distribution’s higher *moments* (multifractal, the largest avalanches). The slope sandpile is single-scale in its means and multifractal in its tails, simultaneously and asymptotically, which places it more precisely against BTW, Manna, and the directed sandpile than any single exponent does.

1 Introduction and Model

Self-organized criticality (SOC) is the observation that some slowly driven dissipative systems settle, with no tuning of a control parameter, into a state with scale-free fluctuations — avalanches whose sizes follow a power law cut off only by the system size. The sandpile is its archetype. Charbonneau’s Chapter 5 develops a particular, continuous-height version, and the work here takes that model as a laboratory for two questions the chapter raises but does not fully resolve: how robust are the critical *exponents* (as opposed to the SOC phenomenon itself), and what does a slope avalanche actually look like?

The 1-D model (eqs 5.1-5.10). Nodal heights $S[0..N-1]$ sit on a line. Slow forcing adds a grain of size drawn uniformly from $(0, \text{eps})$ at a random node, but only on a fully stable iteration — the stop-and-go timescale separation that keeps loading and relaxation from overlapping. A nodal pair is unstable when its slope $z = |S[j+1] - S[j]|$ reaches the critical value Z_c ; an unstable pair halves its slope, moving $z/4$ of sand downslope (eq 5.4). Updates are synchronous. The right boundary is open and drains ($S[N-1] = 0$); the left is a wall. Default parameters are $N = 100$, $\text{eps} = 0.1$, $Z_c = 5$.

The 2-D model — an owned design choice. The 1-D model is a *bond* model: every adjacent pair carries a slope, and an unstable bond halves it. The natural 2-D generalization keeps that exact rule but lets bonds run in both lattice directions — each site couples to its right and upper neighbors by an x-bond and a y-bond, every bond obeying the identical 1-D pair rule, with all four edges open. I chose this over the book’s hinted “2x2-block slope” precisely because it reduces to the 1-D model exactly when one dimension is collapsed. The local dynamics are unchanged across dimension, so any 1-D-versus-2-D exponent difference is a clean effect of dimensionality rather than of also changing the rule. A toppling-count “size” series (number of bond topplings) was added so the model can be compared to the BTW size exponent, which counts site topplings.

All code lives in `sandpile/`, self-contained and importing nothing from the rest of the project. Figures are `figures/sandpile_*.png`; run logs are `outputs/sandpile_*.txt`. The detailed per-finding record (the S-series, S1-S18) is in `findings_sandpile.md`; this report is the synthesis.

2 Validation and the Angle of Repose

The implementation is sound (S1). Before exploring anything new, I checked the model against the chapter’s predictable outcomes. Over a fully instrumented run, grains added minus sand drained at the open edge equals the final pile mass to a residual of $3.3\text{e-}11$ — redistribution moves sand without creating or destroying it, and the boundary is the only sink. A single isolated unstable pair has its slope reduced by exactly a factor of two. From an empty $N = 100$ pile the total mass grows through a transient and saturates to a stationary plateau interrupted by avalanche discharges (`sandpile_validate_mass.png`), avalanches fall inside the energy-duration wedge bounded by slopes $+1$ and $+2$ (`sandpile_validate_ET.png`), and the avalanche-energy PDF is a power law whose logarithmic slope is essentially independent of lattice size ($-0.99, -1.00, -1.01$ for $N = 100, 300, 1000$; `sandpile_validate_pdf.png`). That size-independence is the defining SOC signature. Started from an empty pile, a triangle at repose, and a uniformly loaded pile, the lattice converges to the same stationary slope and the same PDF slope (Exercise 3) — the SOC state is an attractor independent of initial conditions.

Mean slope versus peak slope (S2). The chapter states the stationary slope settles a few percent below Z_c (it quotes $\sim 7\%$ for $\text{eps} = 0.1$), approaching Z_c only as $\text{eps} \rightarrow 0$. My pile shows a

stationary *mean* bond slope about 16% below Z_c , essentially flat across two decades of ϵ s, with the deficit if anything widening slightly as grains get finer — the opposite of the stated trend. This is not an implementation error. The book’s own mass figure (5.4A, $M \sim 2.1e4$ at $N = 100$) back-calculates through $\text{mass} = \text{slope} * N(N-1)/2$ to a slope about 15% below Z_c , agreeing with me, not with the “7%” in the text. The resolution (`sandpile_repose_peak.png`) is that mean and peak slope are different quantities with opposite ϵ -behavior: the redistribution rule *halves* an unstable bond, removing far more sand than the small threshold overshoot a single grain produces, so the time-mean slope is dragged about 16% below Z_c and stays there for any grain size; the *peak* (steepest sustained) slope tracks Z_c , its overshoot being $O(\epsilon)$ and vanishing as grains get finer. The book’s “approaches Z_c as $\epsilon \rightarrow 0$ ” is correct only for the peak. Conflating the two is what makes its 7% claim look inconsistent with its own mass figure. A small but genuine clarification of the chapter.

3 Critical Exponents, Universality, and the Conservation Law

1-D exponents by finite-size scaling (S3). Reading τ from the power-law slope of the largest lattice and the cutoff exponent D from the moment ratio $\langle x^2 \rangle /$ across $N = 64-1024$, the 1-D model gives $\tau_E = 1.03$ (the PDF close to the marginal $1/E$), energy cutoff $D_E \sim 2.0$ ($E_{\max} \sim N^2$, set by the largest avalanche spanning the lattice), and duration cutoff $D_T \sim 1.0$ ($T_{\max} \sim N$). The rescaled PDFs collapse (`sandpile_fss.png`). One structural fact matters for everything later: the avalanches are *quantized families*, not simple fractals. The energy-duration wedge (bounded by slopes +1 and +2) means E and T are not related by a single power law, so the usual scaling relation $\tau_T = 1 + (\tau_E - 1) D_E / D_T$ does not apply — it would predict $\tau_T \sim 1.06$, far from the measured ~ 0.6 . The 1-D model’s locality forces avalanches into discrete line/wedge families rather than one self-affine family, consistent with the chapter’s remark that E , P , and T are correlated “only in a statistical sense.” This is the seed of the duration anomaly that recurs through the whole arc.

The Grand Challenge: 2-D and universality (S4). Run from a pyramid over $L = 32-128$, the 2-D slope model reaches a stationary state (mean bond slope ~ 2.5 , about 50% below Z_c , because four neighbors give each site more relaxation channels) with power-law, collapsing avalanche PDFs (`sandpile_fss2d.png`). Two conclusions:

observable	1-D	2-D slope	2-D BTW
energy τ_E	1.03	0.87	—
size τ_S	—	0.89	1.14
energy cutoff D_E	2.0	~ 2.2	—

First, **1-D and 2-D slope sandpiles are not in the same universality class**: the energy exponent drops from 1.03 to 0.87 under the *identical* local rule, so the change is purely dimensional. Second, **the gradient rule is not in the BTW class**: at the same dimension and under the same log-binned-PDF / FSS pipeline, the slope model’s size exponent (0.89) differs clearly from BTW’s (1.14; `sandpile_btw.png`, consistent with the literature ~ 1.2). How a site decides to topple matters — toppling on a height *difference* exceeding a threshold (gradient) yields different critical behavior than toppling on absolute *height* (BTW). Both are SOC; they are distinct classes. The headline of the chapter falls out here: SOC as a *phenomenon* is robust, but the SOC *exponents* are not.

The honest caveats on S4 were the engine of the rest of the work. The exponents come from modest

L (32-128) with finite-size curvature, so the digits carry ± 0.05 -0.1; the *differences* are several times that, so the conclusions hold, but the numbers are not nailed down. And the slope-model size counts *bond* topplings while BTW counts *site* topplings, an $O(1)$ definitional difference that cannot change an exponent but invites an objection. Removing that objection, and replacing a single hedged exponent with the full scaling structure, is what Sections 4 onward do.

Bulk conservation is necessary (S5, S7). Charbonneau lists four ingredients that appear sufficient for SOC. Is the slope model’s bulk conservation actually necessary, or incidental? Breaking it with a tunable dissipation d (the lower node of each unstable bond receives only $(1-d)$ of the shed sand; the conservative path is kept bit-identical via a branch on $d = 0$), the decisive test is not whether avalanches still occur but whether their cutoff still *scales* with system size. It does not. In 1-D the conservative cutoff scales as $N^{2.00}$ (the criticality of S3); every $d > 0$ slashes the cutoff-scaling exponent below 1 and truncates the PDF at a dissipation-set characteristic size (`sandpile_dissipation.png`). The result transfers to 2-D (cutoff scaling $2.16 \rightarrow 1.46 \rightarrow 1.21$ at $d = 0, 0.05, 0.20$; `sandpile_dissipation2d.png`). Conservation is what makes criticality *exact* rather than approximate: with it, the only sink is the boundary, so the only length that can cut off an avalanche is the system size; break it, and a finite dissipation length becomes the characteristic avalanche size and destroys scale invariance. This is, in the simplest sandpile, the non-conservation sensitivity that makes the Olami-Feder-Christensen earthquake model perennially debated. (Honest limit: over the accessible L range the dissipative exponents land at ~ 0.5 - 0.9 rather than cleanly at 0, because for weak dissipation the correlation length is large; the dichotomy is unambiguous but the asymptotic zero would need larger lattices.)

Boundary falloff (S8, Exercise 2). Tracking the mass that actually leaves the open edge as its own series shows only $\sim 44\%$ of avalanches reach the boundary; falloff is itself scale-invariant but with a *shallower* exponent than bulk toppling (~ 0.64 versus ~ 1.0), so boundary evacuation is a distinct self-similar process; falloff energy correlates strongly with bulk energy (0.91-0.96); and the evacuated fraction is tiny and shrinks with N ($0.023 \rightarrow 0.005$), matching the chapter’s remark that even a large avalanche lowers the pile mass by only $\sim 0.2\%$ (`sandpile_falloff.png`). The interior is frozen at the angle of repose with all the action in a thin avalanching surface layer.

The fast engine and the duration cross-check (S9, S6 $\rightarrow$ S10). The reference engines rescan every bond each iteration ($O(N)$ in 1-D, $O(L^2)$ in 2-D) regardless of how little is moving, which is what left the duration cross-check inconclusive at $L \leq 128$ (S6: the slope model’s τ_T swung from 1.43 to 0.80 between sizes — a quantity that changes by a factor of two between adjacent sizes is not a measured exponent). Charbonneau’s Exercise 5 names the cure: an active-list engine that touches only the moving bonds. I built one (numba-compiled; `sandpile_fast.py`) and proved it is the *same* model, not merely a similar one — both reference engines gained an optional shared-forcing stream, and driven identically the fast and reference engines are bit-for-bit equal (avalanche series, counts, durations, falloff, dissipative paths; only the incrementally accumulated total mass differs at $\sim 1e-9$, and is used for nothing). The speedup is $\sim 600\times$ in 2-D. With it, the duration question is settled (S10, `sandpile_duration_fss.png`): at $L = 64$ - 512 with $\sim 2e5$ avalanches per size the slope model’s τ_T is *stable* at 0.56 versus BTW’s 1.22 — a gap of 0.67, more than twenty times the per-fit scatter. The convention-free duration measure now agrees with the size measure, removing the S4 bond-versus-site objection. The two models even differ in how the duration cutoff grows (slope $L^{1.07}$, a ballistically spreading front; BTW $L^{1.43}$), a second independent signature of the class difference. S6 is resolved not by overturning its honest “cannot tell at $L \leq 128$ ” but by removing the size limitation that caused it.

4 A Scaling Theory of the 2-D Avalanche

Sections 2-3 establish that the slope model is SOC, that its exponents differ from 1-D and from BTW, and that conservation is what makes it critical. But a single tau exponent is a blunt instrument, and when a distribution is multifractal it is ill-defined — which is exactly why S4 needed caveats. This section replaces the single exponent with the full scaling structure, and in doing so builds a mechanistic picture of what a 2-D slope avalanche *is*. The thread is one argument, not eight separate results.

4.1 The Moment-Spectrum Method, Validated on BTW (S11)

The field's decisive discriminator is the moment spectrum: how the family of moments $\langle x^q \rangle$ scales with system size. For a distribution obeying simple finite-size scaling (FSS), $P(x) = x^{-\tau} G(x/x_c)$ with $x_c \sim L^D$, the moment exponent $\sigma(q)$ defined by $\langle x^q \rangle \sim L^{\sigma(q)}$ is *linear* in q , so its local slope $D(q) = d\sigma / dq$ is *constant*. A *drifting* $D(q)$ means multifractal scaling. Before turning this on the slope model, I reproduced the established result for BTW (De Menech, Stella, Tebaldi 1998; Tebaldi, De Menech, Stella 1999): BTW is multifractal in toppling number but near-FSS in area. The machinery (`moments.py`, `moment_fss.py`) carries its own self-test — synthetic avalanches from an exact simple-FSS source must read out as a flat $D(q)$, and they do (drift < 0.01), so the method does not manufacture multifractality. Measured on BTW (`sandpile_moments_btw.png`):

BTW observable	$D(q)$ drift, q in $[1,4]$	read
toppling number S	0.302	multifractal
area A	0.034	near-FSS, $D_{\text{area}} \sim 2.05$

The toppling number is multifractal ($D(q)$ climbs $2.42 \rightarrow 2.72$, toward the literature avalanche dimension ~ 2.75); the area is flat at ~ 2.05 (compact, essentially 2-D footprints). The multifractality lives in the multiple topplings per site, not in the footprint. This is the trusted basis for the slope-model measurement, and it carries the key lesson: **area is the clean, well-behaved observable**.

4.2 The Slope Model Is Filamentary and Anomalous (S12)

The slope model has two genuine obstacles, both solved. Its size and energy have $\tau < 1$ (S3/S4), so every positive moment is set by the few largest avalanches and is noisy and correction-laden — the cure is avalanche AREA (distinct toppled bonds, the footprint), which is bounded by the lattice and not cutoff-dominated. And the slope pile equilibrates slowly and only from above: started over-steep it relaxes its mean down to repose, and the stationarity gauge must be the *mean slope*, not the noisy cutoff-dominated. With area added to the engine (validated bit-for-bit against a full-scan distinct-bond recount) and properly equilibrated runs over $L = 64\text{--}256$, the moment spectrum (`sandpile_moments_slope.png`) reads:

observable	$D(q=1)$	drift $[1,4]$	character
area A	1.01	0.275	filamentary, $D_{\text{area}} \sim 1.0\text{--}1.1$
toppling S	1.89	0.105	space-filling activity, $D \rightarrow 2$
energy E	1.90	0.107	space-filling activity, $D \rightarrow 2$

Three facts combine. The footprint is **filamentary** (area grows only $\sim$ linearly with L , a thin near-1-D front, against BTW’s compact $D_{\text{area}} \sim 2$). The activity is **space-filling** (size and energy scale as L^2 , the lattice area). Therefore each footprint bond topples about L times (size/area $\sim L^2/L^1 = L$) — a thin front swept many times. And no observable has a flat $D(q)$: the area drift is resolved at low-to-mid q (a rise of 0.15 against a noise of 0.003-0.02), so the model is **not** in the stochastic Manna single-fractal FSS class; it sits with the deterministic BTW/Zhang anomalous-scaling family. This is the caveat-free upgrade of S4: not “an exponent differs” but “the avalanche geometry and full moment structure differ” — BTW compact with multifractal toppling number, slope-model filamentary with space-filling activity along a thin front.

4.3 The Mechanism: A Ballistic Filamentary Front (S13)

S12 says *what* the avalanche looks like through moment spectra, which for this model are cutoff-dominated. The conditional exponent $\gamma(x|y)$, defined by $\langle x | y \rangle \sim y^{\{\gamma\}}$, measured at *fixed* y , never sees the system-size cutoff — a clean, caveat-free observable (`conditional.py`). Three earlier results combine into one prediction: a 2-D avalanche is a thin front of linear extent ℓ that propagates one node per iteration (so $T \sim \ell$), leaves a filamentary footprint ($A \sim \ell$), and sweeps each footprint bond $\sim \ell$ times ($S, E \sim \ell^2$). If a single scale ℓ governs the avalanche, the six conditional exponents are fixed and *over-determined* — single-scale scaling forces $\gamma(S|A) = \gamma(S|T)/\gamma(A|T)$, so the picture can fail a self-consistency test. Measured (pooled $L = 192+256$; `sandpile_conditional.png`):

relation	measured	predicted
$\langle E S \rangle$	1.00	1
$\langle S A \rangle$	1.93	2
$\langle E A \rangle$	1.94	2
$\langle A T \rangle$	0.97	1
$\langle S T \rangle$	1.73	2
$\langle E T \rangle$	1.74	2

Two conclusions, one clean and one subtle. **The filamentary-multiply-swept mechanism is confirmed, cutoff-free:** energy locks to size ($E \sim S$, each toppling sheds $\sim Z_c/4$ of mass), and a footprint of A bonds carries $\sim A^{1.93}$ topplings (each bond swept $\sim A$ times). The 2-D energy-duration plane hugs the slope-+2 “wedge” line, and the slope-+1 “line” branch that fills the lower 1-D wedge is essentially gone — the book’s upper wedge bound has become the generic 2-D law. **But single-scale scaling fails through duration:** the over-determination test does not close (direct $\gamma(S|A) = 1.93$ versus routed-through-duration $\gamma(S|T)/\gamma(A|T) = 1.79$, a gap of 0.14, about seven times the fit error), while the purely spatial identity $\gamma(E|S)$ closes exactly. Duration is a *loose* proxy for spatial extent — some long avalanches are long not because they reach farther but because they *linger* (re-topple in place, propagate intermittently). This is the 2-D residue of S3’s quantized line/wedge families. Every exponent drifts toward its single-scale value as L grows, so whether the gap heals or survives as a residual is the question flagged for later.

4.4 The Geometry: A Ballistic Constant-Width Filament (S14)

S12 and S13 both conclude “filamentary, area $\sim L$ ” but each measures only a single area exponent — neither looks at the *shape*. The decisive observable is the per-avalanche mass-radius dimension D ($A \sim R g^D$, footprint cell count against radius of gyration), one number that places the avalanche

on an axis: compact BTW $D = 2$, the exactly-solvable directed (Dhar-Ramaswamy) sandpile $D = 3/2$, a constant-width thin filament $D = 1$. The fast engine gained a validated footprint dump (`dump_fp=True` records each avalanche’s distinct toppled-bond set, first-topple times, and launch site; a pure observer — dynamics bit-identical, sets matching a full-scan reference exactly), and `geometry2d.py` measures it. The estimator is unbiased: it reads synthetic lines, disks, and directed fronts back as $D = 1.000, 2.001, 1.499$.

quantity	slope model	reading
mass-radius D (pooled $L = 192+256$)	1.017 ± 0.003	a thin filament
longitudinal width $w_{\text{long}} \sim A^?$	$A^{0.99}$	the filament length = the area
transverse width (top decile A)	0.093 bonds	constant, ~ 0
one-bond-wide fraction (top decile)	0.98	98% are a single-bond line
anisotropy (top decile)	1.000	a perfect line
BTW mass-radius D (same pipeline)	2.00	compact

The footprint is a **constant-width thin filament** — thinner than the directed sandpile (which diffuses to transverse width $\sim \ell^{1/2}$), far from compact BTW. S12’s $D_{\text{area}} \sim 1$ is now a literal geometric fact, a straight thin line read off real footprints. And it is a **ballistic front**: first-topple time tracks radial distance from the seed one-to-one (slope 0.998 , correlation 0.990), the spatial face of S10’s duration cutoff and S13’s $\langle A|T \rangle$. On the geometric axis the slope model sits *below* the directed sandpile ($1 < 3/2 < 2$): the deterministic gradient rule, with no global drive, makes each avalanche a straight steepest-descent filament with none of the diffusive transverse wandering the stochastic directed model has — more ordered than the canonical directed sandpile, not less (`sandpile_geometry.png`).

4.5 The Causal Test: The Filament Is Made by the Deterministic Rule (S15)

The placement so far is *correlational*: I characterized a deterministic model and noted it differs from Manna. To make it causal, I added a tunable stochastic split to the engine (`psto`): a fraction `psto` of each bond’s downhill flux is diverted to a random transverse neighbor, conservatively. `psto = 0` is the deterministic rule and is bit-identical to S1-S14 (the engine self-test confirms no extra RNG draw, sand conserved to $1e-10$). Sweeping it (`sandpile_stochastic.png`):

<code>psto</code>	mass-radius D ($L=128$)	D ($L=192$)	one-bond-wide
0.00	1.005	1.015	0.99
0.05	1.270	1.284	0.23
0.10	1.538	1.562	0.01
0.20	1.728	1.755	0.00
0.50	1.869	1.862	0.00

Two conclusions. **The filament is a specific consequence of the deterministic rule**: turning on stochasticity drives the footprint from one-bond-wide ($D \sim 1$) toward compact ($D \sim 1.87$), monotonically, with the $L = 128$ and $L = 192$ curves lying on top of each other (intrinsic, not finite-size). Strikingly, near `psto` ~ 0.1 the model passes *through* the directed sandpile’s exact $D = 3/2$ — a small transverse leak reproduces the directed-sandpile geometry, tying the three reference sandpiles (filament 1, directed $3/2$, compact 2) onto one continuous knob. **But the model does**

not collapse onto the Manna class: the area moment drift does the opposite of flattening (0.23 $\rightarrow$ 0.38), and D_{mid} stays ~ 1.1 -1.2. The avalanches become compact in *shape* but *localized* (area $\sim L^{1.2}$, not the L^2 of compact spanning avalanches): the transverse leak breaks the coherent ballistic front into compact but sub-spanning splats without converting the global scaling to simple FSS. So “outside Manna” is not just reaffirmed but reinforced — the gradient rule’s anomalous scaling is robust, sitting in its own place rather than a hair’s breadth from Manna.

4.6 The Dimensional Anchor: The 2-D Filament Is a 1-D Object (S16)

$D \sim 1$ was measured in 2-D, where it is non-trivial — a 2-D lattice *can* host a compact ($D = 2$) avalanche, and BTW does. But $D = 1$ is the *maximal* dimension a 1-D footprint can have, so the question is whether the 2-D filament actually matches a genuine 1-D slope avalanche under the same measurement. I added the S14 footprint dump to the 1-D engine (it had none — validated bit-for-bit against a full-scan brute reference) and measured the native 1-D avalanche with the same estimators (`geometry1d.py`).

quantity	1-D	2-D	reading
mass-radius D	0.999	1.02	same intrinsic dimension
solidity A/range	1.000	—	solid, gap-free interval
ballistic time vs distance	1.000 x dist, corr	0.998 x radius	one node per iteration
	1.000		
$\langle E S \rangle / \langle S A \rangle / \langle A T \rangle$	1.00 / 2.00 / 0.98	1.00 / 1.93 / 0.97	same conditional mechanism

The genuine 1-D avalanche is a **solid, gap-free, ballistic interval** of mass-radius dimension exactly 1 (a Cantor self-test confirms the estimator would catch fractal gaps as $D < 1$, so “solid” is a real measurement). This is exactly the 2-D footprint geometry. So **$D \sim 1$ is the intrinsic dimension of a slope avalanche, not a 2-D measurement artifact** — the directed sandpile (3/2) and BTW (2) are genuinely higher-dimensional objects; the gradient rule is not. The conditional mechanism is dimension-independent, and the S13 duration over-determination gap is ~ 0.11 in 1-D too (48% of avalanches are single-sweep “lines” — S3’s families seen directly), so the duration anomaly is intrinsic to the rule across dimensions, not a 2-D finite-size effect. The one genuine dimensional contrast is *directedness*: the 1-D front is downhill-directed (93-98% of the footprint downhill of the seed, because the 1-D model has a global drive), while the 2-D front radiates isotropically (no global drive). The 2-D avalanche has the *geometry* of a 1-D avalanche rather than being identical to one (`sandpile_geometry1d.png`).

This finding also delivered the enabler for the rest of the arc: a reusable `equilibrate(L)` (`equilibrate2d.py`) that warms a lattice in chunks and detects the stationary plateau by a windowed mean, returning a verified-stationary state. It establishes that $L = 512$ equilibrates to repose 2.745 (compute was never the barrier, a verifiable stopping rule was) and that the repose creep with L (2.42 $\rightarrow$ 2.745 over $L = 64$ -512) is finite-size, not under-equilibration — every L reaches an initial-condition-independent plateau, and $L = 256$ reproduces S12’s 2.64 (`sandpile_equilibrate.png`).

4.7 The Duration Closure: Exact in Space, Residual in Time (S17)

With the enabler, S13’s open question can be answered: does the single-scale over-determination gap *heal* as L grows, or survive as a true residual? Measuring the conditional exponents at verified-stationary states up to $L = 512$ (`duration_closure.py`, `sandpile_duration_closure.png`):

L	gap	gamma(S A)	gamma(A T)	gamma(S T)	gamma(E S)
96	0.066	1.891	0.912	1.664	1.003
256	0.118	1.917	0.964	1.734	1.002
512	0.183	1.947	0.992	1.750	1.002

Reading the directly-measured components rather than the derived gap (the discipline of not trusting an auto-verdict over a ratio of fits): **the spatial sector heals to exact single-scale** — $\gamma(S|A) \rightarrow 2$ and $\gamma(A|T) \rightarrow 1$ as L grows — **but $\gamma(S|T)$ saturates at ~ 1.75** , not 2 (the last step, $1.751 \rightarrow 1.750$, is flat). So the gap does not heal; it *grows* with L ($0.066 \rightarrow 0.183$), toward a residual of ~ 0.2 (a $1/L$ fit gives 0.19 and underestimates, since the gap is convex; the $S|T$ plateau against $\gamma(S|A) \rightarrow 2$, $\gamma(A|T) \rightarrow 1$ implies ~ 0.25). The spatial lock $\gamma(E|S) = 1.00$ at every L confirms the breakdown is duration-specific, not a fit artifact.

The answer to S13: **single-scale scaling in the 2-D slope sandpile is exact in space and residual in time**. The spatial observables (area, size, energy) form a single-scale family that becomes *exact* in the thermodynamic limit ($S, E \sim A^2$ with the exponent converging to 2, $A \sim T$ ballistic with the exponent converging to 1). But at fixed duration the size grows only as $T^{\sim 1.75}$, not the T^2 the spatial chain predicts, and that shortfall does not vanish with system size — it is a true residual. Some long-duration avalanches are long because they *linger*, so duration carries information the spatial extent does not. This is S3’s 1-D family structure surviving to the thermodynamic limit rather than washing out, and it is *stronger* in 2-D (residual ~ 0.2) than in the 1-D anchor (0.11), consistent with a 2-D front having more ways to linger. AREA, not duration, is the model’s clean scaling variable.

4.8 The Multifractality Closure: Asymptotic, Not a Corona (S18)

One question remained from S12: is the avalanche-area multifractality (the $D(q)$ drift ~ 0.27) a *true asymptotic* property, or a finite-size corona that heals to simple FSS? S17 sharpened the stakes: it proved the conditional mean $\langle S|A \rangle \rightarrow A^2$ exactly, with the activity space-filling — read naively, as if the mean were the whole distribution, that predicts a mono-fractal area with a flat $D(q)$, so the multifractality *should* heal. I tested the prediction with the unconditional area moments, localizing the signature by a *sliding* three-lattice window (`area_multifractality.py`, `sandpile_area_multifractal.png`):

L-window center	area $D(q)$ drift over [1,4]
92	0.219 +- 0.040
133	0.272 +- 0.023
185	0.243 +- 0.042
266	0.145 +- 0.062
369	0.194 +- 0.031

The drift does not heal. Across a 4x range in window-center L it stays at about 0.2 and extrapolates to a residual of 0.169 ± 0.048 (about 3.5 sigma from zero). A finite-size corona would have fallen $\sim 4x$ over this range; it does not fall at all. The self-test reads exact simple-FSS data as drift < 0.01 at every window, so a nonzero intercept is real. So S12’s area multifractality is asymptotic.

The reconciliation with S17 is the content of the finding. S17’s $\gamma(S|A) \rightarrow 2$ is a statement about a conditional *mean* — the first moment of S at fixed A . Single-scale scaling of that mean

does *not* imply the area distribution is single-fractal: the higher moments, which weight the rare largest avalanches and the fluctuations about the mean, can scale with their own larger exponents, and that is exactly what multifractality is. The conditional method reads the single-scale “core” (the typical avalanche); the moment method weights the tail and reads the multifractal fluctuations (the largest avalanches). The q -resolved data show it directly — within the largest window $D(q)$ rises from 1.14 at $q \sim 1$ (mean-weighted, near the typical filament) to 1.34 at $q \sim 4$ (tail-weighted, the avalanches that depart from the one-bond filament). The slope sandpile has both: a conditional mean that becomes exactly single-scale in space, and an area distribution whose higher moments stay multifractal asymptotically. (S18 also flagged a tentative second reading — that the typical area dimension is actively *rising* at the largest L , $D_{\text{mid}} 1.07 \rightarrow 1.23$ — but rested it on the fewest-avalanche points. A follow-up self-test (S19, `filament_fattening.py`) settled it: at double the seeds the typical avalanche is L -stable at the filament values (median area $\sim L^{0.90}$, typical mass-radius $D = 1.00 \rightarrow 1.01$, $\sim$ one bond wide to $L = 512$) and the mean-area growth slope that read 1.22 at five seeds averages down to 0.99 ± 0.03 at ten. The footprint does *not* fatten; the apparent fattening was tail-undersampling. S14/S16’s one-bond filament is the asymptotic geometry of the typical avalanche, and the L -growth that S18 saw belongs entirely to the multifractal tail — sharpening S18 rather than overturning it.)

5 Synthesis: One Avalanche, One Theory

Eight findings (S11-S18) describe a single object. The 2-D slope avalanche is a **ballistic, filamentary, multiply-swept front**. *Ballistic*: it propagates one cell per time step, first-topple time tracking radial distance one-to-one (S14), so duration scales with linear extent. *Filamentary*: its footprint is a constant-width thin line of mass-radius dimension 1 (S14), thinner than the directed sandpile and far from BTW, and that dimension is the intrinsic dimension of a slope avalanche, the same in 1-D where it is maximal as in 2-D where it is one of three possibilities (S16). *Multiply-swept*: the activity is space-filling while the footprint is filamentary, so each footprint bond topples about as many times as the footprint is long (S12, S13).

From that single picture a single-scale theory predicts the entire scaling structure, and the predictions are over-determined, so the theory is falsifiable rather than fitted. It passes in space and fails, in a precise and quantified way, in time. The spatial exponents (size $\sim \text{area}^2$, area $\sim \text{duration}$, energy $\sim \text{size}$) reach their single-scale values exactly in the thermodynamic limit (S17). Duration is the one loose variable — at fixed duration the avalanche keeps a lingering degree of freedom that does not wash out, a residual of order 0.2 that is the 2-D survival of the 1-D line/wedge families the chapter started from (S3, S16, S17). And the *distribution*, as opposed to the conditional mean, is multifractal in its tails asymptotically (S18), so the single-scale-in-space statement is precisely a statement about the *typical* avalanche, with the largest avalanches scaling anomalously.

This places the model concretely in the SOC landscape, which is the question the whole universality thread (S4 onward) was reaching for. Against the three reference sandpiles:

- **versus BTW** (compact, area $D = 2$, multifractal in toppling number): the slope model is filamentary (area $D = 1$) with multifractality in the *area* tail instead — a different geometry and a different locus for the anomaly.
- **versus the directed (Dhar-Ramaswamy) sandpile** (exactly single-scale, area $D = 3/2$): the slope model is *thinner* ($D = 1$, no diffusive transverse spread) and is single-scale only in its spatial means, not exactly — it has a duration residual the directed model does not, and a multifractal area tail.

- **versus Manna** (stochastic, single-fractal FSS): the slope model is robustly *outside* Manna — a tunable stochastic split moves its geometry toward compact but never flattens its moment spectrum to simple FSS (S15).

The single most useful sentence the arc produces is that the slope sandpile is **single-scale in its means and multifractal in its tails, simultaneously and asymptotically**. That is a sharper placement than any single exponent gives, and it resolves what looked like a contradiction between the chapter’s two main tools (conditional exponents say “single-scale,” moment spectra say “multifractal”) by locating each precisely.

The headline of Sections 2-3 survives intact and is now mechanistic. SOC as a *phenomenon* — scale-free avalanches as a dynamical attractor with no parameter tuning — is robust to dimension and to the toppling rule. The SOC *exponents*, and now the whole scaling structure, are not: they are set by the specific geometry the rule produces, and for the deterministic gradient rule that geometry is a one-bond-wide ballistic front.

6 Method, Self-Tests, and Honest Limits

A recurring discipline runs through this work and is worth stating on its own, because it is the part most easily lost in a results summary. Every analysis tool was validated on a known result or a synthetic source before being trusted on the data, and several planned checks were reported as failures rather than read off a bad fit.

- The moment-spectrum and conditional-exponent machinery each carry a self-test that recovers the right answer on an exact synthetic source (flat $D(q)$ on a simple-FSS source; 1, 2, 2, 2, 2, 1 on a single-scale conditional source; a measured drift on a bifractal mixture) before touching the model (S11, S13, S17, S18).
- The geometry estimators read synthetic lines, disks, directed fronts, and a Cantor set back at their known dimensions, so a measured $D = 1$ is a real “solid filament” statement, not a tautology (S14, S16).
- Every engine extension (area tracking, the footprint dump in 2-D and 1-D, the stochastic split, dissipation) was validated bit-for-bit against the reference implementation, with the default path proven identical to the original dynamics, so no result rests on an unverified rewrite (S9, S12, S14, S15, S16).
- The duration cross-check was reported as *inconclusive* at $L \leq 128$ rather than forced to a number (S6), and resolved only once the fast engine removed the size limit (S10). The auto-generated “verdict” in that script, which averages an unstable fit, is exactly the trap the honest record avoids — and the same discipline (read the directly-measured components, not the derived auto-verdict) is what makes the S17 and S18 closures trustworthy.

The honest limits are equally explicit and are owned in `findings_sandpile.md` finding by finding. Exponents come from log-binned fits over modest lattices, so individual digits carry ± 0.05 - 0.1 ; the conclusions rest on *differences* and *trends* that are several times that uncertainty, not on any single number. The repose mean slope creeps upward with L ($2.42 \rightarrow 2.745$ over $L = 64$ - 512); this is confirmed finite-size rather than under-equilibration, but whether it saturates or grows slowly is not settled at $L \leq 512$. The asymptotic magnitudes of the duration residual (~ 0.2) and the multifractal residual (~ 0.17) are bounded but not pinned to a single digit. And the most interesting hint — that the typical filament may fatten at the largest lattices (S18) — is the least certain, resting on the fewest-avalanche points, and is left as the natural follow-up. None of these limits

touches the qualitative spine of the arc, which rests on cross-checked trends and dimension-matched comparisons.

7 Conclusions

1. The continuous-height slope sandpile reproduces the chapter’s SOC signatures exactly (conservation to $3e-11$, N -independent power-law avalanches, initial-condition-independent attractor), and clarifies one textbook conflation: the angle of repose sits $\sim 16\%$ below the threshold for any grain size (the *mean* slope), while only the *peak* slope approaches the threshold as forcing becomes fine.
2. SOC as a phenomenon is robust, but its exponents are not. The slope model’s energy exponent changes with dimension under the identical local rule ($1.03 \rightarrow 0.87$), and the gradient rule is in a different universality class from BTW (size exponent 0.89 versus 1.14 , confirmed convention-free by the duration exponent 0.56 versus 1.22).
3. Bulk conservation is necessary for true scale invariance, in 1-D and 2-D: any dissipation imposes a finite correlation length that becomes the characteristic avalanche size and stops the cutoff from scaling with system size. This is the OFC non-conservation sensitivity in the simplest possible setting.
4. The 2-D slope avalanche is a ballistic, filamentary, multiply-swept front: a constant-width one-bond-wide footprint (mass-radius dimension 1, thinner than the directed sandpile, far from compact BTW) that propagates one cell per step and sweeps each footprint bond about as many times as the footprint is long. The filament is *caused* by the deterministic rule (stochasticity drives its dimension to 1.87) and is the intrinsic geometry of a slope avalanche, the same as in the model’s native 1-D.
5. A single-scale theory of that front predicts the entire spatial exponent set exactly in the thermodynamic limit, with one quantified exception: duration is a loose proxy for spatial extent, and the looseness does not heal but grows to a residual of order 0.2 — the 1-D family structure surviving to the thermodynamic limit.
6. The avalanche-area distribution is multifractal asymptotically, not as a finite-size artifact. This reconciles with the single-scale result by distinguishing the conditional mean (single-scale, the typical avalanche) from the distribution’s higher moments (multifractal, the largest avalanches). The model is single-scale in its means and multifractal in its tails, simultaneously and asymptotically — a more precise placement against BTW, Manna, and the directed sandpile than any single exponent provides.

8 References

- P. Charbonneau, *Natural Complexity: A Modeling Handbook* (Princeton University Press, 2017), Chapter 5.
- P. Bak, C. Tang, K. Wiesenfeld, “Self-organized criticality: An explanation of $1/f$ noise,” *Phys. Rev. Lett.* **59**, 381 (1987).
- M. De Menech, A. L. Stella, C. Tebaldi, “Rare events and breakdown of simple scaling in the Abelian sandpile model,” *Phys. Rev. E* **58**, R2677 (1998).

- C. Tebaldi, M. De Menech, A. L. Stella, “Multifractal scaling in the Bak-Tang-Wiesenfeld sandpile and edge events,” *Phys. Rev. Lett.* **83**, 3952 (1999).
- D. Dhar, R. Ramaswamy, “Exactly solved model of self-organized critical phenomena,” *Phys. Rev. Lett.* **63**, 1659 (1989).
- S. S. Manna, “Two-state model of self-organized criticality,” *J. Phys. A* **24**, L363 (1991).
- Z. Olami, H. J. S. Feder, K. Christensen, “Self-organized criticality in a continuous, nonconservative cellular automaton modeling earthquakes,” *Phys. Rev. Lett.* **68**, 1244 (1992).

9 Appendix: Code and Reproducibility

All code is in `sandpile/`, self-contained, vectorized, ASCII-only, and `__main__`-guarded with self-tests. The reference engines are `sandpile1d.py` and `sandpile2d.py`; the validated active-list engine is `sandpile_fast.py`. Analysis: `validate1d.py` (S1), `repose_peak.py` (S2), `fss1d.py` (S3), `fss2d.py` / `btw_compare.py` (S4), `dissipation.py` / `dissipation2d.py` (S5, S7), `falloff.py` (S8), `duration_compare.py` / `duration_fss2d.py` (S6, S10), `moments.py` / `moment_fss.py` (S11), `moment_slope.py` (S12), `conditional.py` (S13), `geometry2d.py` (S14), `stochastic_split.py` (S15), `geometry1d.py` / `equilibrate2d.py` (S16), `duration_closure.py` (S17), `area_multifractality.py` (S18). Every engine extension is validated bit-for-bit against the reference implementation under a shared forcing stream, with the default path proven identical to the original dynamics. The full per-finding record, including all caveats, is in `findings_sandpile.md`.