Topological Freezing with Q-Shift and Multilevel Sampling

in Lattice Gauge Theory

Ankur Singha, Jacob Kauffmann, Karl Jansen, Jacob Finkenrath,
Vipul Arora, Shinichi Nakajima

Bridging Machine Learning and Statistical Field Theory through Generative and Diffusion Processes
ECT*, Trento 7–11 September 2026









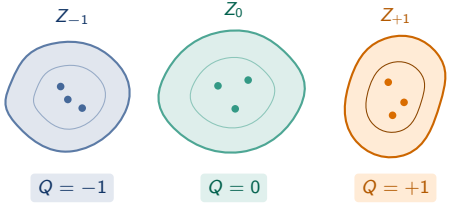


arXiv:2503.08918 arXiv:2604.10209

Part I — Introduction: sector-resolved topological sampling

- configuration space splits by the topological charge Q
- the path-integral weight is spread *across* sectors
- correct sampling needs both: fluctuations *within* a sector, and the weights *between* them
- freezing — from two different causes:
  1. *MCMC*: a high action barrier between sectors
  2. *generative models*: reverse-KL training mode-collapses — all modes are hard to learn

Lorenzo' talk



$$Z = \sum_Q Z_Q\,, \qquad P(Q) = \frac{Z_Q}{\sum_{Q'} Z_{Q'}}$$

We give a general construction that resolves the sectors explicitly in a generative-model setting.

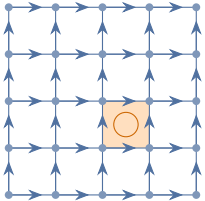
2D $U(1)$

2D $U(1)$ has nontrivial topology, and algorithms like HMC suffer severe topological freezing.

$$S[U] = -\beta \sum_p \cos \varphi_p$$

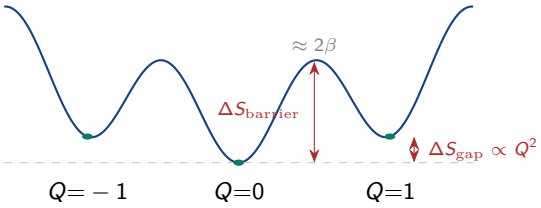
$$\boxed{Q[U] = \frac{1}{2\pi} \sum_p \varphi_p \ \in \ \mathbb{Z}}$$

compact links on a periodic $L \times L$ lattice;
plaquette angle $\varphi_p \in [-\pi, \pi)$



Freezing in HMC:Action barrier

Each Q is a separate sector. To change it you must force a plaquette across $\pm\pi$ — cross a barrier. The sector weights $P(Q)$ are physics ($\chi_{\text{top}} = \langle Q^2 \rangle / V$).



Barrier height increases with β — a canonical hard sampling problem

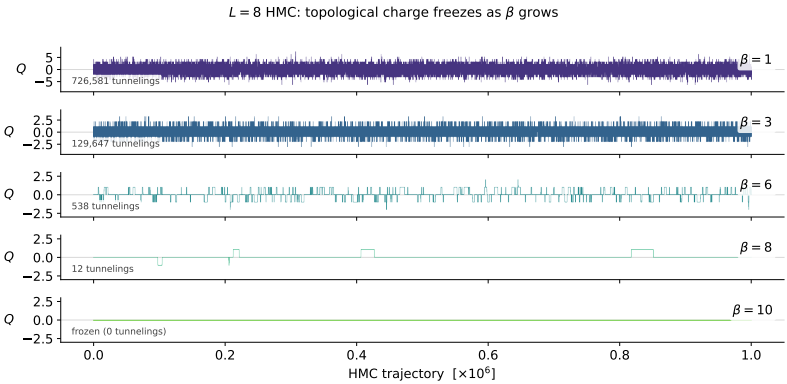
Talk by Lorenzo

Del Debbio et al., PLB 2004; Schaefer et al., NPB 2011.

Topological freezing in HMC: Toy simulation

Caused by $\Delta S_{\text{barrier}} \approx 2\beta$

local updates cannot cross it — the charge freezes



same lattice, growing β: tunnelings crash from $\sim 10^6$ to zero

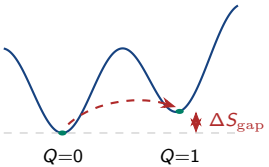
Lüscher & Schaefer, JHEP 2011.

Freezing: Overrelaxation with HMC

Caused by ΔS_{gap}

global moves *are* proposed — but accepted only with probability $\propto e^{-\Delta S_{\mathrm{gap}}}$

- wHMC — winding HMC
- Instanton update (HMC-IU)



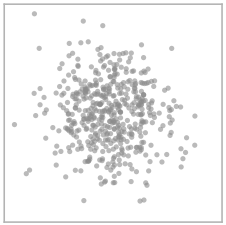
as the barrier grows toward the continuum the acceptance collapses exponentially — they freeze too

Albandea et al., EPJC 2021; Eichhorn & Hoelbling, PoS Lattice 2021.

Generative flows Talk by Gurtej

A normalizing flow gives independent samples and their exact density — but one flow must represent all sectors at once.

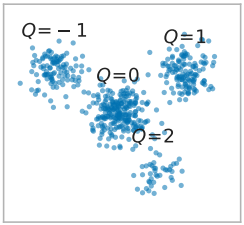
base noise $z \sim \mathcal{N}(0, \mathbb{1})$



f_ϕ

one shot

target p: several sectors

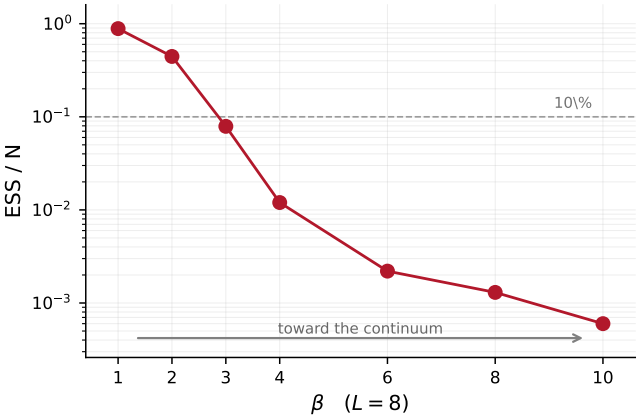


Lorenzo' talk
Elia' talk

mode-seeking training collapses to one sector

Kanwar et al., PRL 2020; Singha et al., PRD 2023; Caselle et al., JHEP 2022.

. . . and it shows up as soon as β grows



single-scale spline flow
trained at each β

$$\mathrm{ESS}/N$$

$$0.89 \longrightarrow 6 \times 10^{-4}$$

$\beta = 1 \rightarrow 10$ (fixed $L = 8$)

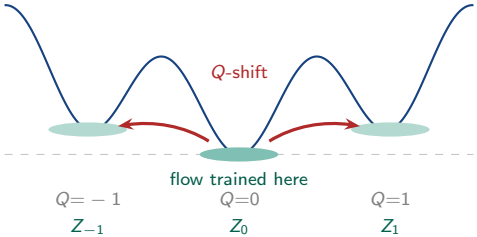
three orders of magnitude
on a *fixed* lattice

Talk by Fernando
Talk by Nils
Talk by Lena and Simran

volume scaling is a *separate* issue — the flow scales with neither β nor volume

The idea: resolve the sectors

Don't fight the barrier. Train a flow in one sector, then map it to the others and reweight.

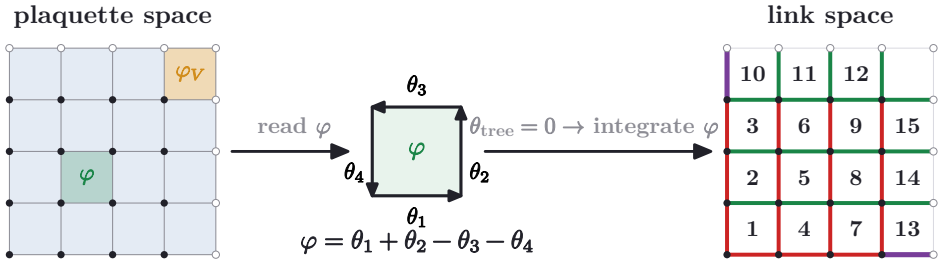


$$Z = \sum_Q Z_Q\,, \qquad P(Q) = \frac{Z_Q}{\sum_{Q'} Z_{Q'}}$$

the Q-shift recovers the Z_Q — never crossing a barrier

Unbiased observables from neural samplers: K. Nicoli, **S. Nakajima**, N. Strodthoff, W. Samek, **K.-R. Müller**, P. Kessel, PRE 2020.

Step 1 — Training in plaquette space: Guage fixing



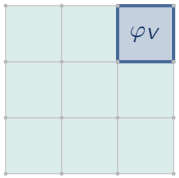
we sample the plaquette angles φ

gauge invariant, local in the action

the links follow by integration

tree gauge $\theta_{\text{tree}} = 0$; a bijection

Step 1 — collapsing to the reference sector



$\varphi_{p<V}$ $\sim$ flow

the flow generates the free plaquettes — φ_V is then fixed by Bianchi

$$\varphi_V = -\sum_{p<V} \varphi_p$$

so it reads off the winding of whatever was generated

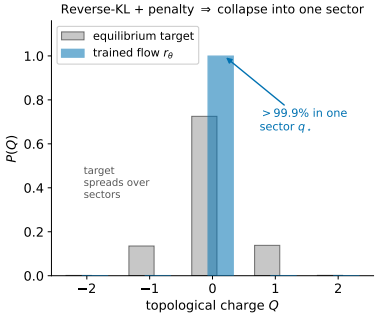
$$Q_{\text{soft}} = \frac{\varphi_V}{2\pi} \qquad \text{differentiable}$$

$$\mathcal{L} = \underbrace{\left\langle S + \log r_\theta \right\rangle_{r_\theta}}_{\text{reverse KL}} + \lambda \underbrace{\left\langle Q_{\text{soft}}^2 \right\rangle}_{\text{pins } Q=0}$$

one penalty on one plaquette — no barrier is crossed, the flow is simply built in the reference sector

Step 1 — in practice

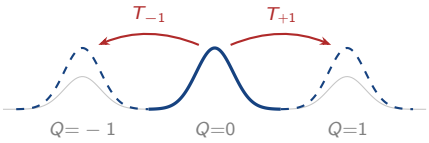
After training, $> 99.9\%$ of samples land in $Q = 0$. The rare $Q \neq 0$ leakage is discarded (*first* discard) and the density renormalized over $Q = 0$.

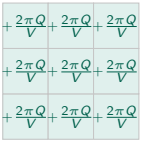


a clean sampler of the reference sector

Step 2 — the Q-shift: a bijective transformation

A deterministic, unit-Jacobian shift $T_{0\to Q} : \varphi_j \to \varphi_j + \frac{2\pi Q}{V}$ copies the collapsed $Q = 0$ flow into every sector.





$$\sum_j \frac{2\pi Q}{V} = 2\pi Q \;\Rightarrow\; \Delta Q = Q$$

unit Jacobian — the density transforms trivially

Step 3 — sector weights, and the mixture

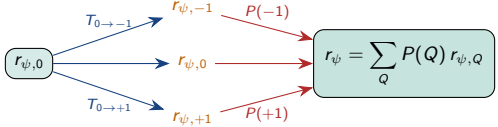
$r_{\psi,0}$: the trained flow — the reference sector

$T_{0\to Q}$: the shift — bijective, unit Jacobian

$r_{\psi,Q} = (T_{0\to Q})_{\#}\, r_{\psi,0}$: the density in every sector, known exactly

these three are all that is needed for the partition function

$$Z_Q = \left\langle e^{-S} / r_{\psi,Q} \right\rangle_{r_{\psi,Q}} \qquad P(Q) = \frac{Z_Q}{\sum_q Z_q} \qquad Z = \sum_Q Z_Q$$



unbiased across sectors by construction — not by tunnelling

Nicoli et al., PRL 2021.

Benchmark: 2D $U(1)$ toward the continuum

Fix the physics ($\beta/L^2 = \text{const}$) and grow $L = 8 \to 32$

MCMC-based

local & topology-changing updates

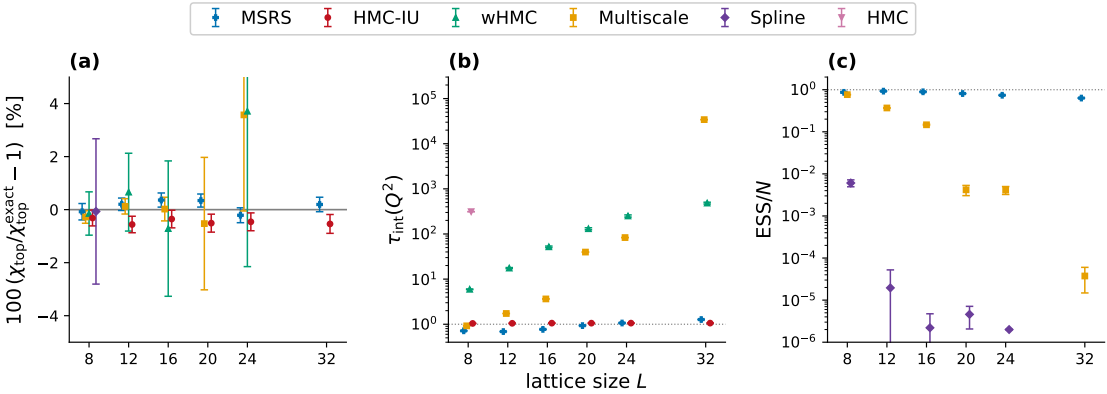
HMC · winding HMC · instanton update

Generative-model-based

normalizing flows trained on the theory

spline flow · multiscale flow

Results : Line of constant physics (LCP)



(a) unbiased $\chi_{\rm top}$

at every volume

(b) $\tau_{\rm int}(Q^2) = \mathcal{O}(1)$

baselines reach 10^4

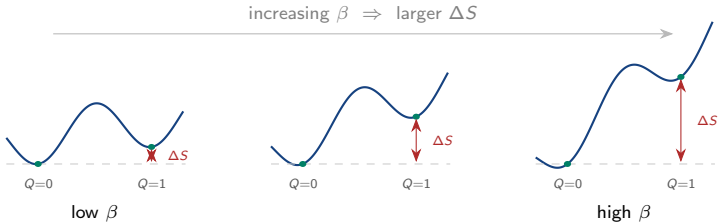
(c) ESS/N flat

spline flow at 10^{-6}

line of constant physics — every L is the same continuum theory

The instanton update: our best baseline

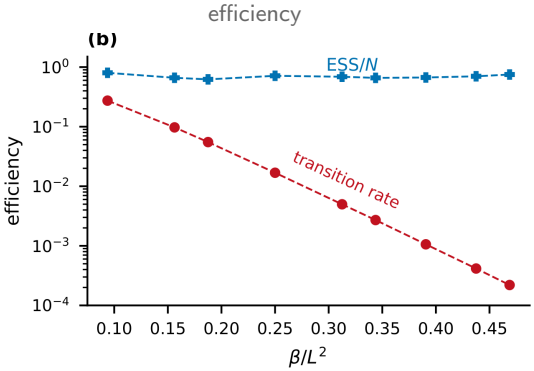
The instanton update proposes a global shift $\varphi_j \to \varphi_j \pm \frac{2\pi}{V}$
Accepted with Metropolis probability $\min(1, e^{-\Delta S})$



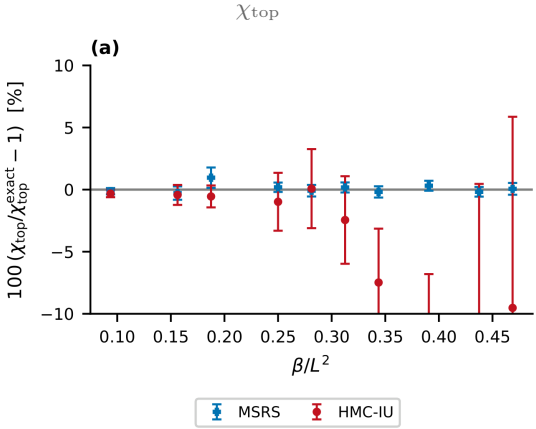
at low β: small ΔS, high acceptance — it tunnels freely and is unbiased

Eichhorn & Hoelbling, PoS Lattice 2021; Hasenbusch, PRD 2017.

Compared with instanton update



ESS/N stays flat; the transition rate collapses



HMC-IU develops variance and bias; MSRS stays on the exact line

we propose in every sector by construction — never waiting for a transition

Summary



arXiv: XXXX.XXXXX

Part II
Sampling in plaquette space

1. Plaquette space — higher groups, higher dimensions
2. Multilevel sampling in plaquette space

2D $U(1)$	4D $U(1)$	2D $SU(2)$
volume	dimension	gauge group

talk by Elia
talk by Anna

Our earlier multilevel samplers before this:

spin systems — ICLR 2025, arXiv:2503.08918 · 2D ϕ^4 — arXiv:2604.10209

with Shin and Elia

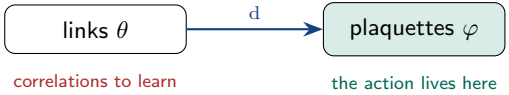
Sample the variables the action sees

$$S = -\beta \sum_p \cos \varphi_p$$

separable — one term per variable

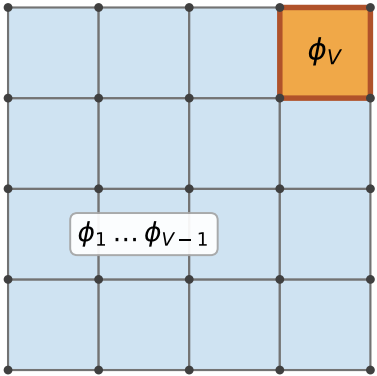
in links each term couples four variables

$$\varphi_p = \theta_\mu + \theta_\nu - \theta_\mu - \theta_\nu$$



any affine closure solves the same way $\Rightarrow$ in principle non-local actions

plaquette sampling : Constraints



one global constraint

$$\sum_{p=1}^{V} \phi_p \;=\; 2\pi Q$$

$$\phi_V \;=\; 2\pi Q \;-\sum_{p<V} c_p\,\phi_p\,, \qquad c_p \in \{-1, 0, +1\}$$

$$\underbrace{\#\{\, c_p \neq 0 \,\}}_{\text{support}} \;=\; V - 1$$

one angle correlated with *all* the others

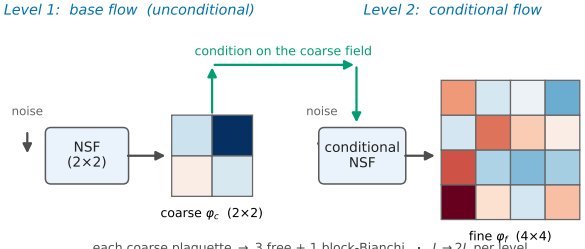
Multilevel: keep the solve *short*



■ generated by the flow ■ fixed by Φ_c

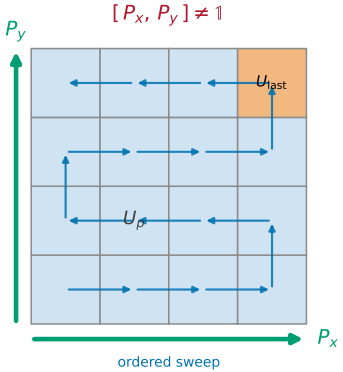
$$\phi_4 = \Phi_c - \phi_1 - \phi_2 - \phi_3$$

$$\underbrace{\#\{\, c_p \neq 0 \,\}}_{\text{support}} = 3 \quad \text{every level}$$



$$\sum_p \phi_p = \sum_c \Phi_c = \cdots = 2\pi Q$$ blocks telescope — global closure imposed once, at the coarsest level

2D $SU(2)$: the constraint changes



going round the square: $P_x \to P_y \to P_x^\dagger \to P_y^\dagger$

$$\prod_p^{\text{ordered}} \widehat{U}_p \;=\; [P_x, P_y] \;\equiv\; P_x P_y P_x^\dagger P_y^\dagger$$

$$\boxed{\widehat{U}_{\text{last}} = \Big(\prod_{p \neq \text{last}}^{\text{ordered}} \widehat{U}_p \Big)^{-1} [P_x, P_y]}$$

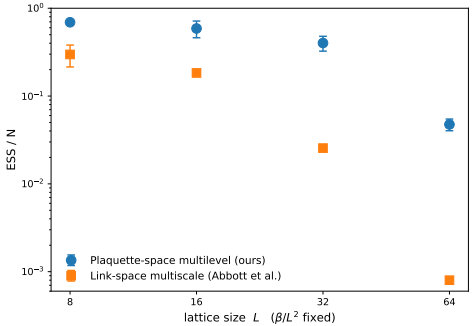
$U(1)$: $[P_x, P_y] = 0$ holonomy-free

$SU(2)$: $[P_x, P_y] \neq \mathbf{1}$ U_p and P_x, P_y coupled

2D results: $U(1)$ and $SU(2)$

$U(1)$ — across the volume

constant physics $\beta/L^2 = 0.188$



ESS/N 0.69 $\longrightarrow$ 0.05

$L = 8 \rightarrow 64$ (4096 d.o.f.)

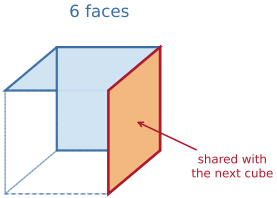
$SU(2)$ — across β, at 16×16

ESS/N [%]

		published		*this work*	
β	ξ/a	Gerdes	Boyda	single	multilevel
2.2	1.14	**87**	80	83.1	—
2.7	1.20	68	56	**81.4**	—
8	2.24	—	—	75.3	—
12	2.77	—	—	28.1	51.5
20	3.61	—	—	31.3	55.7
40	5.13	—	—	3.4	58.8
60	6.30	—	—	4.3	65.0
65.9	6.60	—	—	0.14	65.9

Gerdes '24, Boyda '21 · $N = 5 \times 10^5$: single level = median of 5 seeds; multilevel = single draw

4D: the constraint becomes *local*



$$(\mathrm{d}\phi)_{\mu\nu\rho}(x) = \sum_{p \in \partial\,\text{cube}} \pm\phi_p \;=\; 0 \pmod{2\pi}$$

one elementary 3-cube · 4 triples $(\mu\nu\rho)$ per site $\Rightarrow$ $4V$ equations

every face is shared by two cubes

$\Rightarrow$ they are not independent

$$\#\,\text{independent} \;=\; 3V + 3 \;=\; 771$$

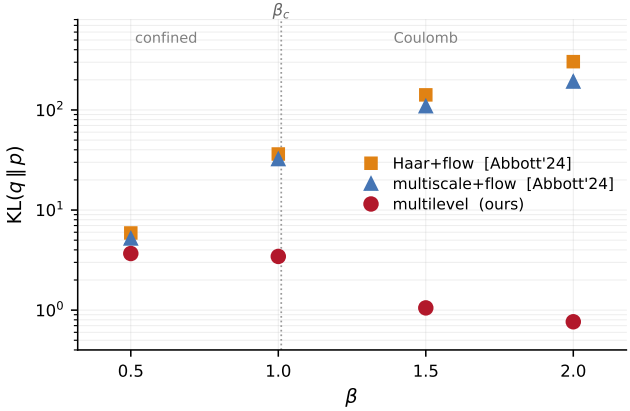
$6V = 1536$ plaquettes at 4^4

771 determined (50.2%)	**765** sampled by the flow (49.8%)

one relation per cube · half the plaquettes follow from the other half

gauge fixing (maximal tree) makes the map one-to-one: $4V = 1024$ links $\rightarrow$ $3V{+}1 = 769$ $=$ $\underbrace{765}_{\text{free }\phi} + \underbrace{4}_{\text{Polyakov}}$

4D $U(1)$: results



same benchmark, same KL

β	KL	ESS/N
0.5	3.68	0.16%
1.0	3.44	0.02%
1.5	1.05	0.60%
2.0	0.77	11.9%

both: median of 12 draws, $N = 262{,}144$

reweighted $\langle \cos\phi \rangle$ agrees with the exact reference at $\beta = 0.5, 1.5, 2.0$

0.5σ, 0.8σ, 2.6σ — not at β_c

baseline: Abbott et al., arXiv:2404.10819, Fig. 4

Two problems, two constructions

topological freezing

the sectors are disconnected

Q-shift

train $Q{=}0$ $\rightarrow$ $T_{0\to Q}$ $\rightarrow$ reweight

$\tau_{\mathrm{int}}(Q) = 1.1$ vs 1439

χ_{top} unbiased against exact

scaling with volume and dimension

the closure couples the whole lattice

multilevel in plaquette space

coarse $\rightarrow$ refine $\rightarrow$ refine

2D $U(1)$: ESS $0.69 \rightarrow 0.05$ to $L = 64$

4D $U(1)$: KL 0.77 vs 192

sample the variables the action sees · solve the constraint, don't penalise it

With thanks to my collaborators


Shinichi Nakajima


Karl Jansen


Jacob Kauffmann


Jacob Finkenrath


Vipul Arora


Stefan Kühn

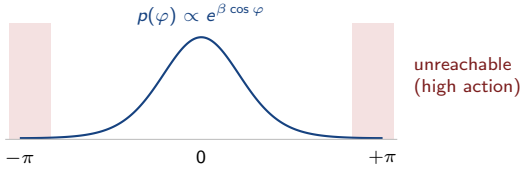
DESY · TU Berlin · BIFOLD · KU Leuven · IIT Kanpur · Wuppertal

Thank you for your attention!

Backup

Step 2 — the Q-shift: the second discard

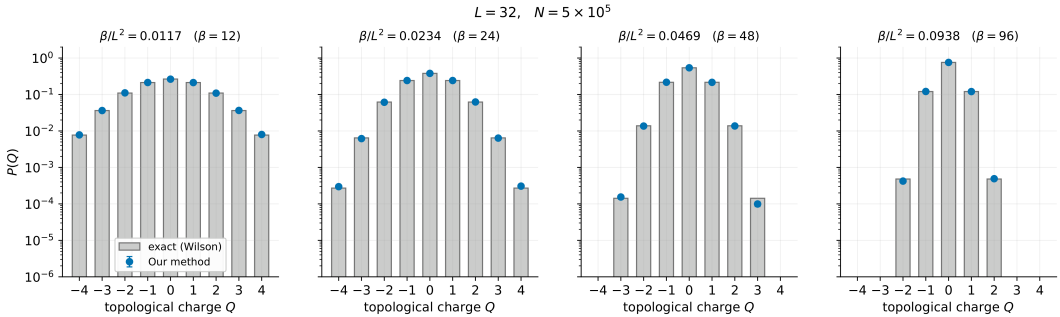
- First discard — the $Q \neq 0$ *leakage* at collapse.
- Second discard — after the shift a plaquette can *cross* $\pm\pi$ into the wrong sector ($< 0.04\%$); this leaves a thin region unreachable.



exponentially suppressed: $P(\Omega_Q) \approx P(\textit{reachable})$

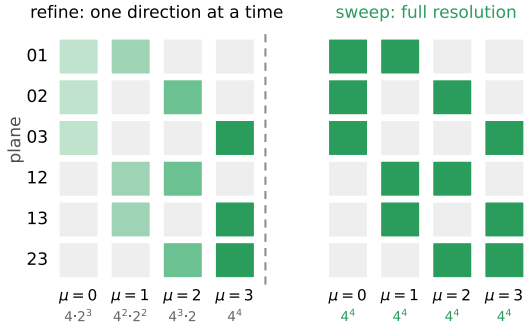
Result: the full charge distribution

The topological-charge distribution $P(Q)$ matches the exact Wilson result across volumes and sectors.



Backup — 4D: full-resolution sweeps

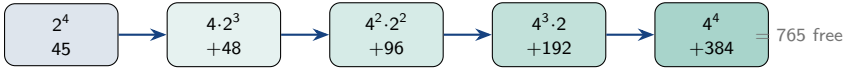
$$\phi_{\text{det}} = -\big(\underbrace{C\,\phi_{\text{free}}}_{\text{this level}} + \underbrace{D\,\Phi_{\text{c}}}_{\text{the coarse level}} \big), \qquad \text{same-level support} \le 4$$



refining along μ frees only the 3 planes containing μ — the sweeps give every plane a turn at 4^4

4D: multilevel keeps the solve *short*

$$\phi_{\text{det}} \;=\; -\Big(\; \underbrace{C\,\phi_{\text{free}}}_{\text{drawn at this level}} \;+\; \underbrace{D\,\Phi_{\text{c}}}_{\text{already generated}} \;\Big) \qquad C, D \text{ integer} \in \{-1, 0, 1\}$$



upsample one direction at a time — the flow draws only the *new* angles

one refinement ($4^3{\cdot}2 \to 4^4$): every determined plaquette has one of *three* roles

	rows	free angles entering
copied from the coarse field	384	0
block partner $\Phi_{\text{c}} - \phi$	384	1
closed by one cube	384	4
direct solve on 4^4	771	5–47, *median* 13

the constraint is local, so the solve stays local — the flow never has to correlate an angle with the whole lattice