Large deviations for the 3D dimer model

Catherine Wolfram
(joint work with Nishant Chandgotia and Scott Sheffield)

May 25, 2023

Massachusetts Institute of Technology

Plan for the talk

- **Part I**: set up.
  - what the dimer model is in 2D or 3D
  - what large deviation principle (LDP) for dimers would mean
- **Part II**: digression: LDP for dimers in 2D.
  - this is just for context, since our methods in 3D are different.
- **Part III**: statements of LDP theorems in 3D.
- **Part IV**: simulations!
- **Part V** (time permitting): a few main ingredients and ideas for the proofs in 3D.

Everything in this talk is joint work with Nishant Chandgotia and Scott Sheffield (preprint on arxiv with the same title as the talk) [1].

Part I: set up

What is the dimer model?

Dimers in 2D are *dominoes*, e.g. 1×2 or 2×1 blocks.

Dimers in 3D are *bricks*, e.g. $2 \times 1 \times 1$ or $1 \times 2 \times 1$ or $1 \times 1 \times 2$ blocks.

A dimer tiling of a region $R \subset \mathbb{Z}^2$ or $\mathbb{Z}^3$ is a collection of dimer tiles such that every square/cube is covered by exactly one tile.

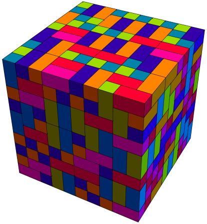

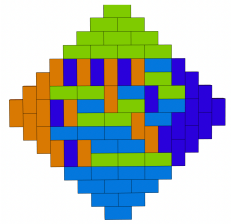

What do the colors of dimers mean?

What do the colors of dimers mean?

For any d, $\mathbb{Z}^d$ is a bipartite lattice, with underlying black and white checkerboard.

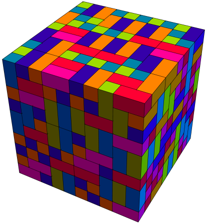

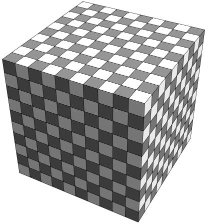

The colors of the dimers represent the cardinal direction of the dimer (north, south, east, west, up, down for $d = 3$), viewed as a vector from its white cube to its black cube.

There is a correspondence between 1) a dimer tiling τ of $\mathbb{Z}^d$ and 2) a *discrete vector field* v_τ defined by: for each edge e of $\mathbb{Z}^d$ oriented from white to black,

$$v_\tau(e) = \begin{cases} 1 & e \in \tau \\ 0 & e \notin \tau \end{cases}$$

Correspondence: dimer tilings and divergence free discrete vector fields

Observation: compute divergences of v_τ.

$$\operatorname{div} v_\tau(x) = \sum_{\substack{e \ni x \\ \text{oriented out of } x}} v_\tau(e) = \begin{cases} +1 & x \text{ is white} \\ -1 & x \text{ is black.} \end{cases}$$

Upshot: divergences depend only on the parity of x.

Subtracting a constant reference flow $r(e) = 1/(2d)$ for all $e \in \mathbb{Z}^d$, a dimer tiling τ corresponds to a divergence free discrete vector field f_τ which we call the ***tiling flow***.

When $d = 3$ this is

$$f_\tau(e) = \begin{cases} 1 - 1/6 = 5/6 & e \in \tau \\ -1/6 & e \notin \tau \end{cases}$$

The tiling flow construction works in **any dimension** and plays a role analogous to the height function in 2D (more on this later).

The main intuition throughout this talk is to think of a dimer tiling as a *flow*.

Set up for large deviations in 2D or 3D

Fix, for dimension $d = 2$ or 3:

- a "reasonable" compact region $R \subset \mathbb{R}^d$ and some boundary condition b (boundary condition can be specified e.g. using the flow)
- a sequence of grid regions $R_n \subset \frac{1}{n}\mathbb{Z}^d$ approximating R with the boundary conditions of R_n converging to b as $n \to \infty$.

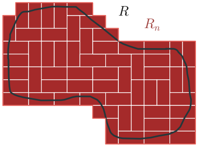

Question: what do uniform random dimer tilings of R_n look like in the fine-mesh limit as $n \to \infty$?

As a dimer tiling corresponds to a discrete divergence free flow on $\frac{1}{n}\mathbb{Z}^d$, the fine-mesh limit as $n \to \infty$ should be some *measurable divergence free vector field.*

Large deviations means quantifying: given a deterministic flow g, what is the probability that a tiling of R_n is close to g as $n \to \infty$? There is a *limit shape* if is there is one flow that random tilings concentrate on as $n \to \infty$.

Ingredients of an LDP

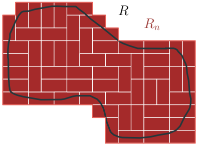

In general, a large deviation principle (LDP) needs:

1. A sequence of probability measures $(\rho_n)_{n\geq 1}$ that the large deviation principle is about.
2. A topology (to say what the fine-mesh limits are, and to compare things)
3. A rate function $I(\cdot)$, where I measures, for any fixed $\delta > 0$,

   " ρ_n(tiling flow f_τ is within δ of deterministic flow g) $\approx \exp(-n^d \cdot I(g))$"
4. When the rate function $I(\cdot)$ has a **unique minimizer** and $(\rho_n)_{n\geq 1}$ satisfy an LDP, then the ρ_n-probability that a random tiling is close to minimizer goes to 1 as $n \to \infty$. The minimizer is called the *limit shape*.
5. Main step for proving that $I(\cdot)$ has a unique minimizer is usually to prove that $I(\cdot)$ is *strictly convex*.

Part II: two dimensions (for context)

Context: large deviations in 2D

In 2000, Cohn, Kenyon and Propp [2] proved an LDP for the 2D dimer tilings of simply connected compact regions $R \subset \mathbb{R}^2$ and showed that there is a **limit shape**.

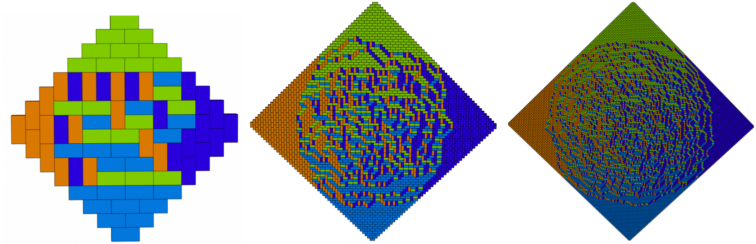

Example: finer and finer aztec diamonds. The interface between the frozen/rough regions is exactly a circle in the limit.

To compare with 3D, will explain the measures $(\rho_n)_{n \geq 1}$, topology, and how to understand the rate function $I(\cdot)$.

Dimers in 2D: LDP topology from height functions

In 2D, the divergence-free tiling flow f_τ can be upgraded, and we get a correspondence between 2D dimer tilings and Lipschitz functions called *height functions* (up to an additive constant).

Example of a height function of a tiling. From [10].

The LDP in 2D is in terms of height functions, i.e. two tilings are close if their height functions are close in the sup norm.

To get from tiling flow to height function: a div-free flow f in 2D is dual (by rotating by $\pi/2$) to a curl-free flow. The scalar potential of the curl-free dual of the tiling flow in 2D is the height function h, that is, ∇h is a rotation of f_τ.

In this sense the tiling flow f_τ is a natural generalization of ∇h.

The fine-mesh limits of height functions as $n \to \infty$ are called *asymptotic height functions*. These are 2-Lipschitz functions h that satisfy

$$\nabla h \in \{(x, y) : |x| + |y| \leq 2\} = \diamond.$$

To define measures, start with: simply connected compact region $R \subset \mathbb{R}^2$ and boundary value h_b on ∂R, which is an asymptotic height function restricted to ∂R.

Boundary conditions: suppose τ_1, τ_2 are dimer tilings with height functions h_1, h_2.

$$\tau_1, \tau_2 \text{ tilings of same region } R_n \subset \mathbb{Z}^2 \iff h_1, h_2 \text{ have } h_1 |_{\partial R_n} = h_2 |_{\partial R_n}.$$

I.e., choosing a sequence of regions $R_n \subset \frac{1}{n}\mathbb{Z}^2$ is equivalent to choosing a sequence of *boundary height functions* h_n.

We choose regions $R_n \subset \frac{1}{n}\mathbb{Z}^2$ so that the boundary conditions h_n converge to h_b as $n \to \infty$ and let ρ_n be the uniform measure on dimer tilings of R_n.

If h is an asymptotic height function with $h\,|_{\partial R}= h_b$, the measures ρ_n satisfy an LDP with *rate function* I of the form

$$I(h) = C - \mathrm{Ent}_2(\nabla h),$$

where

$$\mathrm{Ent}_2(h) = \frac{1}{\mathrm{area}(R)} \int_R \mathrm{ent}_2(\nabla h(x))\,\mathrm{d}x.$$

Note that I has a unique **minimizer** if Ent_2 has a unique **maximizer**.

The function ent_2 can be **computed explicitly**, and this is the main tool in 2D for showing strict convexity and proving that I has a unique minimizer given boundary conditions.

This uses linear algebra methods called *Kasteleyn theory*. The formula is

$$\mathrm{ent}_2(s_1, s_2) = \sum_{i=1}^{4} L(\pi p_i),$$

where p_i are determined by (s_1, s_2) with the equations $p_1 + p_2 + p_3 + p_4 = 1$, $s_1 = 2(p_1 - p_2), s_2 = 2(p_3 - p_4)$, and $\sin(\pi p_1)\sin(\pi p_2) = \sin(\pi p_3)\sin(\pi p_4)$ and $L(z) = \int_0^z \log|2\sin t|\,\mathrm{d}t$ is the *Lobachevsky function*.

Part III: moving to three dimensions

LDP and limit shape in 3D?

Yes! Simulation of random tiling:

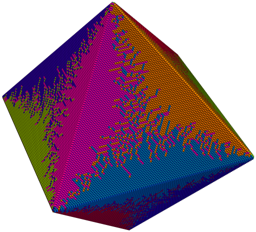

What has changed from 2D to 3D: 1) there is no height function correspondence, so we need a new topology and 2) there are no (known) exact formulas in 3D for the rate function, so need a new way to study this, for example to prove strict convexity.

- Measures ρ_n;
- Topology for comparing tilings;
- Rate function I.

We use *tiling flows* instead of height functions to define the topology for the LDP in 3D.

Recall that any dimension d, there is a correspondence

$$\left\{\text{dimer tilings } \tau \text{ of } \mathbb{Z}^d\right\} \iff \left\{\text{div free discrete flows } f_\tau\right\}.$$

The corresponding flow is called a *tiling flow*.

When $d = 3$, for each edge e of $\mathbb{Z}^3$ oriented from black to white,

$$f_\tau(e) = \begin{cases} 1 - 1/6 = 5/6 & e \in \tau \\ -1/6 & e \notin \tau \end{cases}$$

Topology on flows: induced by metric d_W on flows is a version of *Wasserstein distance*.

Two flows are close if we to transform one flow into another with low "cost" where "cost" is the minimum sum of 1) amount of flow moved times distance moved, 2) flow added, 3) flow deleted to transform one flow into the other.

The fine-mesh limits of tiling flows in this topology are measurable divergence-free vector fields g that we call *asymptotic flows* valued in the *mean-current octahedron*

$$\mathcal{O} = \{s = (s_1, s_2, s_3) : |s_1| + |s_2| + |s_3| \leq 1\}.$$

A element $s \in \mathcal{O}$ is called a *mean current*.

Motivating the name *mean current*: suppose μ is a measure on dimer tilings which is invariant under translations by even integers (these are the translations which preserve the colors of dimers/directions of flow) .

Then we can compute the *expected flow through the origin* (equivalently, expected flow through any white vertex) as

$$s_i = \mu(e_i \in \tau) - \mu(-e_i \in \tau) \qquad i = 1, 2, 3,$$

where e_1 (resp. $-e_1$) is the edge connecting the origin to $(1, 0, 0)$ (resp. $(-1, 0, 0)$).

Hard vs. soft boundary conditions: ways to define ρ_n

Like with height functions, τ_1, τ_2 are tilings of the same region $R_n \subset \mathbb{Z}^3$ if and only if their tiling flows have the same boundary values (i.e., same flow of vector field through boundary).

Fix a nice region $R \subset \mathbb{R}^3$ (compact, closure of a domain, ∂R piecewise smooth) and b a boundary value on ∂R.

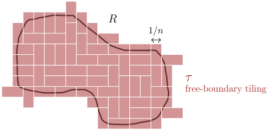

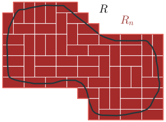

Hard boundary (HB): fix a sequence of regions $R_n \subset \frac{1}{n}\mathbb{Z}^3$ with boundary values b_n approximating b and let $\overline{\rho}_n$ be uniform measure on dimer tilings of R_n.

Soft boundary (SB): choose a sequence of "thresholds" $(\theta_n)_{n\geq 0}$ with $\theta_n \to 0$ slowly enough and let ρ_n be uniform measure on free-boundary tilings of $R \cap \frac{1}{n}\mathbb{Z}^3$ with boundary values within θ_n of b.

The soft boundary measures ρ_n will satisfy an LDP for a slightly larger class of regions that the hard boundary ones $\overline{\rho}_n$.

LDP rate function

For either $(\rho_n)_{n\geq 1}$ (soft boundary) or $(\overline{\rho}_n)$ (hard boundary), the *rate function* when an LDP holds is

$$I_b(g) = C - \mathrm{Ent}_3(g),$$

Like 2D, the entropy functional Ent_3 is an average of a local entropy function ent_3:

$$\mathrm{Ent}_3(g) = \frac{1}{\mathrm{Vol}(R)} \int_R \mathrm{ent}_3(g(x))\, dx.$$

The local entropy function $\mathrm{ent} = \mathrm{ent}_3 : \mathcal{O} \to [0, \infty)$ is defined more abstractly as:

$$\mathrm{ent}(s) = \max_{\mu \in \mathcal{P}^s} h(\mu),$$

where $\mathcal{P}^s$ are measures on tilings of $\mathbb{Z}^3$ invariant under even translations with mean current $s = (s_1, s_2, s_3) \in \mathcal{O}$, i.e. s is the expected flow through the origin in a tiling sampled from μ, and $h(\cdot)$ is *specific entropy*, defined as the limit of Shannon entropy per site.

$$h(\mu) := -\lim_{n\to\infty} |\Lambda_n|^{-1} \sum_{\tau \in \Omega(\Lambda_n)} \mu(\{\sigma \in \Omega : \sigma|_{\Lambda_n} = \tau\}) \log \mu(\{\sigma \in \Omega : \sigma|_{\Lambda_n} = \tau\}).$$

where $\Lambda_n = [-n, n]^3$, Ω is tilings of $\mathbb{Z}^3$, $\Omega(\Lambda)$ is tilings restricted to Λ.

Dictionary between 2D LDP and 3D soft boundary LDP set ups

	2D	3D
compact region R that is...	simply connected [2], multiply connected [7]	closure of connected domain, ∂R piecewise smooth
object associated to tiling τ	height function h	tiling flow f_τ
topology (to compare tilings)	sup norm on height functions	Wasserstein metric d_W on tiling flows
limits of discrete objects	asymptotic height functions: 2-Lipschitz functions	asymptotic flows: div-free meas. vector fields valued in $\mathcal{O}$
rate function	$C_2 - \mathrm{Ent}_2(\nabla h)$	$C_3 - \mathrm{Ent}_3(f_\tau)$

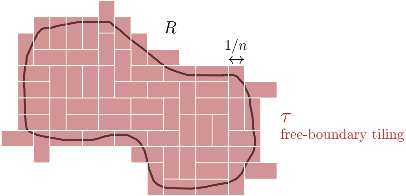

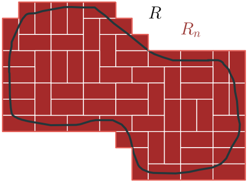

Main theorems in 3D summarized

Recall that $\mathcal{O}$ is the mean current octahedron, let $\mathcal{E}$ be the edges of $\partial\mathcal{O}$. We say a region/boundary value pair (R, b) is

- *flexible* if for every $x \in R$, there exists a neighborhood $U \ni x$ and an extension g of b such that $g(U) \subset \text{Int}(\mathcal{O})$.
- *semi-flexible* if for every $x \in R$, there exists a neighborhood $U \ni x$ and an extension g of b such that $g(U) \subset \mathcal{O} \setminus \mathcal{E}$.
- Otherwise (R, b) is *rigid*.

Theorems (Chandgotia, Sheffield, W.) Assume that $R \subset \mathbb{R}^3$ is the closure of a connected domain and ∂R is piecewise smooth.

(R, b)	SB LDP (ρ_n)	Ent maximizer/I_b minimizer unique	HB LDP $(\overline{\rho}_n)$
rigid	yes	not known	no
semi-flexible	yes	yes	no
flexible	yes	yes	yes

The hard boundary LDP is **provably not true** in full generality in 3D; there exists (R, b) semi-flexible where the HB LDP is false.

For (R, b) rigid, we can prove *weak uniqueness*. Namely, if f_1, f_2 are both Ent maximizers, then on the set A where they differ they are both valued in $\mathcal{E}$.

Part IV: simulations

Simulations: aztechedron and slices

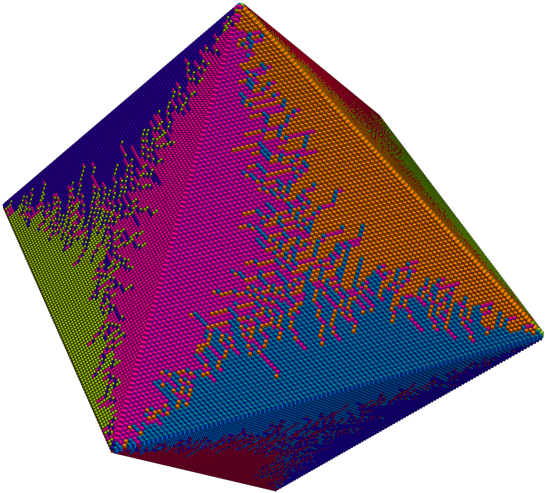

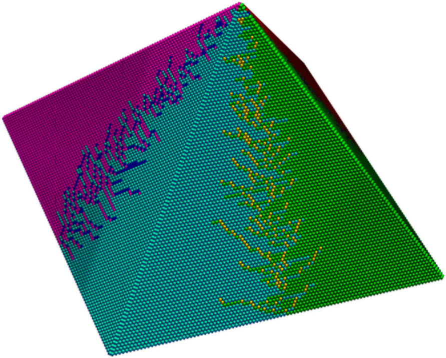

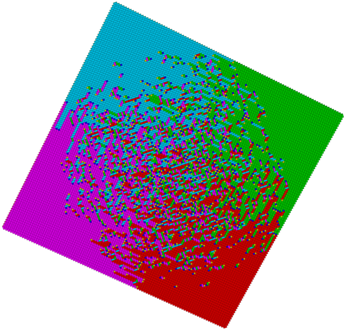

Part V: proofs?

Proofs?

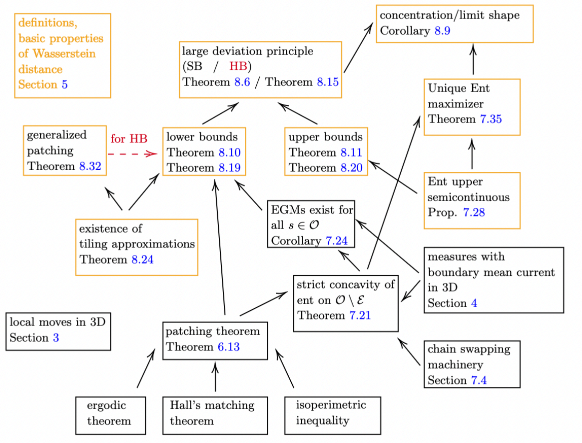

Will say a bit about 1) patching, 2) how to understand the rate function, in particular strict concavity of ent, and 3) measures with boundary mean current (since this is something where 3D is different from 2D).

Patching result: if two tilings τ_1, τ_2 have "asymptotically the same" mean current $s \in \text{Int}(\mathcal{O})$, then you can cut out a piece of one tiling and "patch it in" to the other.

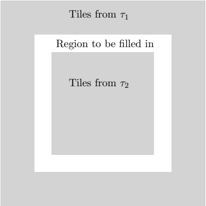

This is some sort of "locality" property that says that the mean current s is sufficient information to be able to stitch two tilings together, as long as $s \in \text{Int}(\mathcal{O})$.

In 2D [2], patching can be proved using Lipschitz extension theory, but those techniques don't work for us.

Note: the patching theorem is *false* if the tilings approximate a constant flow with a value $s \in \partial\mathcal{O}$ instead of $s \in \text{Int}(\mathcal{O})$.

Let $B_n = [-n, n]^3$ and fix $\delta > 0$. Given two tilings τ_1, τ_2 of $\mathbb{Z}^3$, when can we "patch together" τ_1 outside B_n to τ_2 inside $B_{(1-\delta)n}$ by tiling the annulus $A_n = B_n \setminus B_{(1-\delta)n}$ between them?

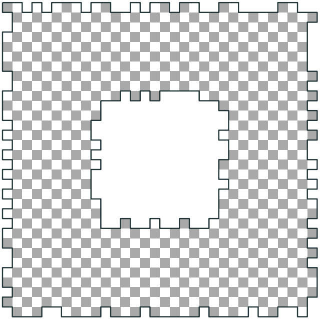

In other words, under what conditions is an annular region like the one above exactly tileable by dimers?

Hall's matching theorem

Necessary condition: a dimer contains 1 black cube and 1 white cube, so the region R needs to have white$(R) =$ black(R). We call this *balanced.*

Hall's matching theorem [5] gives a necessary and sufficient condition:

Theorem. A balanced region $R \subset \mathbb{Z}^3$ is tileable by dimers if and only if there is no *counterexample set* $U \subset R$, i.e. no set of cubes which has white$(U) >$ black(U), despite having only black cubes along its boundary within R.

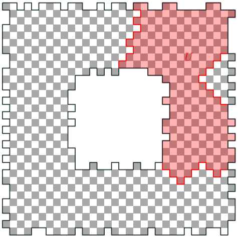

To prove the patching theorem, we show that there are no counterexamples to tileability of A_n when n is large enough and apply Hall's matching theorem.

To prove the HB LDP, we use a *generalized patching* theorem, that says we can patch two tilings close to the same asymptotic flow g extending b.

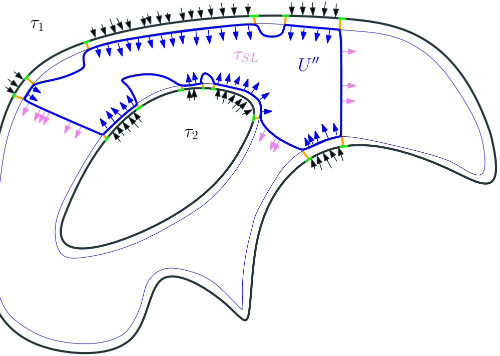

Analogous to the condition that $s \in \mathrm{Int}(\mathcal{O})$ for the regular patching theorem, for generalized patching we need g to be valued in $\mathrm{Int}(\mathcal{O})$ (at least on any compact set $D \subset R$).

This is why we need (R, b) to be *flexible* for HB LDP to hold.

Unlike in 2D, we do **not** have a formula for $\text{ent}(s)$ in $\text{Int}(\mathcal{O})$. (We do have a formula on $\partial\mathcal{O}$ though—I will get to this in a minute time permitting!)

To prove that the rate function

$$I_b(g) = C - \text{Ent}(g) = C - \frac{1}{\text{Vol}(R)} \int_R \text{ent}(g(x))\,\mathrm{d}x$$

has a unique minimizer, one of the important steps is to show that $\text{ent}(s)$ is *strictly concave* on $\mathcal{O} \setminus \mathcal{E}$, where $\mathcal{E}$ is the edges of $\mathcal{O}$.

Like in 2D, $\text{ent}\,|_{\mathcal{E}} \equiv 0$, so it is not strictly concave on $\mathcal{E}$. (This is why we need the *semi-flexible* condition to prove that the Ent maximizer is unique.)

Without a formula for ent we need "soft arguments" for strict concavity. The main idea, for $s \in \text{Int}(\mathcal{O})$, is a method called *chain swapping*.

Before explaining chain swapping, want to explain some background about *Gibbs measures*.

A measure μ is *Gibbs* if for any finite region B, μ conditional on a tiling σ of $\mathbb{Z}^3 \setminus B$ is uniform on tilings τ of B extending σ.

Gibbs measures have a special relationship with entropy.

- Classical result [8]: specific entropy $h(\cdot)$ is maximized by Gibbs measure.
- Straightforward to extend this to say that

$$\text{ent}(s) = \max_{\mu \in \mathcal{P}^s} h(\mu),$$

  is realized by a Gibbs measure of mean current s.
- As a corollary of the **patching theorem**: if μ_1, μ_2 are *ergodic Gibbs measures* (EGMs) of the same mean current $s \in \text{Int}(\mathcal{O})$, then $h(\mu_1) = h(\mu_2)$.

Idea with chain swapping: uses two measures μ_1, μ_2 of mean currents s_1, s_2 to construct new two measures with mean currents $(s_1 + s_2)/2$ and the same total entropy, but then show that this breaks the Gibbs property.

Let $\mu = (\mu_1, \mu_2)$ be a measure on pairs of dimer tilings which is invariant under even translations, sample (τ_1, τ_2) from μ. The union $\tau_1 \cup \tau_2$ is a collection of double tiles, finite loops, and *infinite paths*.

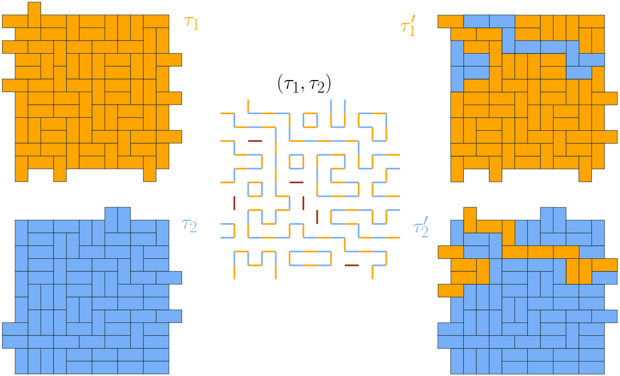

Chain swapping: for each infinite path "of nonzero slope" $\ell \subset (\tau_1, \tau_2)$, with independent probability 1/2 we swap the tiles from τ_1, τ_2 to construct a new pair of tilings (τ_1', τ_2'). This defines a new *swapped measure* $\mu' = (\mu_1', \mu_2')$.

Chain swapping to prove strict concavity on Int($\mathcal{O}$)

Suppose μ is an erogdic coupling of ergodic measures μ_1, μ_2 on dimer tilings, with mean currents $s(\mu_1) \neq s(\mu_2)$. Let μ' be the swapped measure, with marginals μ_1', μ_2'. Chain swapping...

- **Preserves ergodicity**: μ' and hence μ_1', μ_2' are ergodic.
- **Preserves total entropy**: $h(\mu_1) + h(\mu_2) = h(\mu) = h(\mu') = h(\mu_1') + h(\mu_2')$.
- **Preserves but redistributes mean current**: for $i = 1, 2$,

$$s(\mu_i') = \frac{s(\mu_1) + s(\mu_2)}{2}.$$

- **BREAKS the Gibbs property**: if μ_1, μ_2 are Gibbs, then μ_1', μ_2' are **not** Gibbs.

proof* of strict concavity for $s \in$ **Int($\mathcal{O}$)**: Given mean currents $s_1 \neq s_2$, $\frac{s_1+s_2}{2} \in \text{Int}(\mathcal{O})$, let μ_1, μ_2 Gibbs be such that $\text{ent}(s_i) = h(\mu_i)$. If μ_1, μ_2 are *ergodic* Gibbs measures, then by chain swapping

$$2\text{ent}\left(\frac{s_1 + s_2}{2}\right) > h(\mu_1') + h(\mu_2') = h(\mu_1) + h(\mu_2) = \text{ent}(s_1) + \text{ent}(s_2).$$

*full proof uses case work based on ergodic decompositions (we don't yet know that EGMs of every mean current exist), but this is the main idea.

Both patching and the “proof” of strict concavity need $s \in \text{Int}(\mathcal{O})$. What is happening for $s \in \partial\mathcal{O}$?

Moral: measures with boundary mean current in dimension d have a constraint, so they behave like measures of arbitrary mean current on a $(d-1)$ dimensional lattice.

In 2D, we see something trivial on $\partial\mathcal{O}_2$ since it reduces to a 1-dimensional problem, and indeed $\text{ent}_2 \,|_{\partial\mathcal{O}_2} \equiv 0$.

In 3D, it reduces to a non-trivial *2-dimensional problem*.

If μ is a measure of mean current $s \in \partial\mathcal{O}$, then μ a.s. samples only three orthogonal kinds of tiles (wlog north, east, and up).

A tiling with only these three types of tiles in 3D corresponds to a sequence of tilings on 2D *slabs*,

$$L_c = \{(x_1, x_2, x_3) : x_1 + x_2 + x_3 = 2c \text{ or } 2c + 1\}.$$

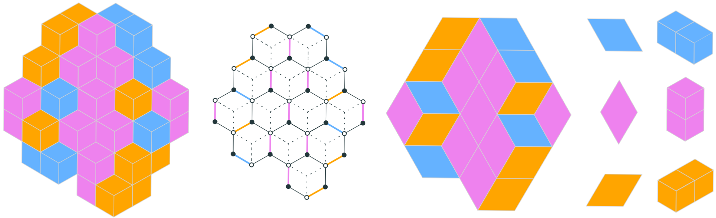

Theorem. Each slab L_c is a copy of the hexagonal lattice. If τ is a tiling of $\mathbb{Z}^3$ with mean current $(s_1, s_2, s_3) \in \partial\mathcal{O}$, then $\tau\,|_{L_c}$ is a lozenge tiling with tile densities (s_1, s_2, s_3).

Measures with boundary mean current (i.e. $s \in \partial\mathcal{O}$)

An EGM μ with boundary mean current s corresponds to a sequence of marginal measures $(\rho_c)_{c\in\mathbb{Z}}$ on each slab L_c, where $s(\rho_c)$ averages to s. Maximum entropy sequence: $(\rho_c)_{c\in\mathbb{Z}}$ i.i.d. copies of the unique EGM on lozenge tilings of slope s.

Theorem. For any face $\mathcal{F} \subset \partial\mathcal{O}$, for $s = (s_1, s_2, s_3) \in \mathcal{F}$,

$$\mathrm{ent}_3(s_1, s_2, s_3) = \mathrm{ent}_{\mathrm{loz}}(|s_1|, |s_2|, |s_3|) = L(\pi|s_1|) + L(\pi|s_2|) + L(\pi|s_3|),$$

where $L(z) = \int_0^z \log|2\sin t|\,\mathrm{d}t$.

In 2D, $\mathcal{O}_2 = \diamond$ and $\mathrm{ent}_2(s) = 0$ for all $s \in \partial\mathcal{O}_2$ because this becomes the *1-dimensional problem* of tiling strips.

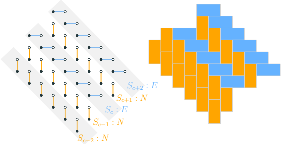

Thank you for listening!!

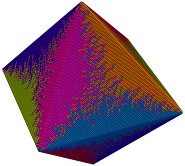

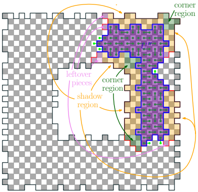

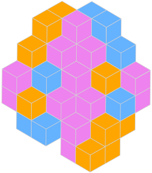

Various open questions...

- Is there is a unique EGM of mean current s for all $s \in \text{Int}(\mathcal{O})$?
- What can be said about the interfaces between frozen and liquid regions in the limit shapes? How big should the fluctuations be?
- Do there exist regions $R \subset \mathbb{R}^3$ (with ∂R piecewise smooth) and boundary conditions b where (R, b) has more than one Ent maximizer?
- Now we know a limit shape *exists*. Are there soft arguments, for example, for the existence of frozen regions in the limit shape?
- and more...

Open question: local moves and mixing times?

Any two dimer tilings of a simply connected region $D \subset \mathbb{R}^2$ are connected by a finite sequence of *flips*.

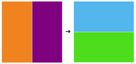

Random tilings in 2D can be simulated by Glauber dynamics with flips. Flip connectedness *fails* in 3D. Lots of interesting work on this, e.g. [9, 3, 4, 6].

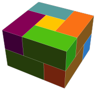

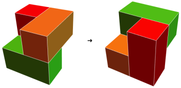

Open question: are flips and trits enough to connect any two tilings of a $M \times N \times L$ box, for $M, N, L > 2$?

Our simulations are generated a Markov chain that constructs *random loops* in a tiling τ: 1) choose a random point, 2) follow alternating sequence of tiles/random directions until it forms a lasso, 3) "shift" tiles on loop part.

Open question: what is the mixing time of this *loop shift Markov chain*?

N. Chandgotia, S. Sheffield, and C. Wolfram.
Large deviations for the 3d dimer model.
arXiv preprint arXiv:math/2304.08468, 2023.

H. Cohn, R. Kenyon, and J. Propp.
A variational principle for domino tilings.
Journal of the American Mathematical Society, 14(2):297–346, 2001.

M. Freedman, M. B. Hastings, C. Nayak, and X.-L. Qi.
Weakly coupled non-Abelian anyons in three dimensions.
Physical Review B, 84(24):245119, 2011.

J. Freire, C. J. Klivans, P. H. Milet, and N. C. Saldanha.
On the connectivity of spaces of three-dimensional tilings.
Transactions of the American Mathematical Society, 375:1579–1605, 2022.

P. Hall.
On representatives of subsets.
Journal of the London Mathematical Society, s1-10(1):26–30, 1935.

I. Hartarsky, L. Lichev, and F. Toninelli.
Local dimer dynamics in higher dimensions.
arXiv preprint arXiv:2304.10930, 2023.

N. Kuchumov.
A variational principle for domino tilings of multiply connected domains.
arXiv preprint arXiv:2110.06896, 2022.

O. E. Lanford, III and D. Ruelle.
Observables at infinity and states with short range correlations in statistical mechanics.
Comm. Math. Phys., 13:194–215, 1969.

N. C. Saldanha.
Domino tilings of cylinders: the domino group and connected components under flips.
Indiana Univ. Math. J., 71(3):965–1002, 2022.

W. P. Thurston.
Conway's tiling groups.
The American Mathematical Monthly, 97(8):757–773, 1990.