STATS 369: Methods from Statistical Physics

Lecturer: Professor Sourav Chatterjee

Notes by: Andrew Lin

Autumn 2025

1 September 22, 2025

This will be a course about **lattice gauge theories**. Lecture notes will be posted on Canvas as the course goes on.

Lattice gauge theories are discrete versions of Euclidean Yang-Mills theories (or what the latter thing is really "supposed to be"), and the purpose of Euclidean Yang-Mills theories, briefly, is constructing quantum field theories. There's an approach where we construct first a probabilistic object (random field or random distribution) and then using tools from functional analysis construct a quantum field theory. (Just like any quantum system, this involves producing a Hilbert space and a time-evolution group or Hamiltonian, where the vectors are states and a group of unitary operators $U(t)$ describes how the states evolve over time.) In quantum field theory the Hilbert space is quite complicated, and there is a well-developed machinery of converting a Markov process to a quantum system: given a two-sided stationary Markov process $\{X_t\}_{t\in\mathbb{R}}$, we define the time-evolution semigroup

$$P_t f(x) = \mathbb{E}[f(X_t)|X_0 = x].$$

We can then typically find some Hamiltonian H where $P_t = e^{-tH}$, and once we've extracted H we can construct (with some machinery called Stone's theorem)

$$U_t = e^{-itH}.$$

So things connect to probability in a different way from the "probability in quantum mechanics," and lattice gauge theories are just about the rigorous discrete objects we can construct (and hopefully one day take a scaling limit to get Euclidean, then quantum Yang-Mills).

This is all rather abstract, so we'll do something a bit more concrete now:

Definition 1

Let G be a compact matrix Lie group contained in $U(n)$ for some n (such as $SU(n)$ or $SO(n)$), and let $\mathfrak{g}$ be its Lie algebra. A **$\mathfrak{g}$-valued 1-form** $A = \sum_{i=0}^{3} A_i dx_i$ on $\mathbb{R}^{1,3}$ is a 4-tuple $A = (A_0, A_1, A_2, A_3)$ of smooth $\mathfrak{g}$-valued functions. The **curvature form** F of A is the 2-form

$$F = \sum_{0\le i<j\le 3} F_{ij}\, dx_i \wedge dx_j$$

where $F_{ij} = \partial_i A_j - \partial_j A_i + [A_i, A_j]$.

For example, if the Lie group is $U(n)$, the Lie algebra is the space of skew-Hermitian matrices, and so to form a $\mathfrak{g}$-valued 1 form we have four such matrices at each point in $\mathbb{R}^4$, in a way smoothly varying in each coordinate. So derivatives make sense and we can get a curvature form which is a collection of six matrices.

Definition 2

The **Yang-Mills action** associated to a $\mathfrak{g}$-valued 1-form is

$$S(A) = -\frac{1}{2g_0^2}\int_{\mathbb{R}^4} \mathrm{Tr}\left[\sum_{0\le i<j\le 3} \eta_i\eta_j F_{ij}(x)^2\right] dx,$$

where $\eta_0 = 1$ and $\eta_1 = \eta_2 = \eta_3 = -1$ (this is usually called the Minkowski signature).

To put this in the framework of the program we just mentioned, we need to write down a Euclidean version of this action $S(A)$. For this, we will perform "Wick rotation," which is a heuristic process where we replace t by $-it$ everywhere (here $t = x_0$ is the first coordinate). If we make such a substitution, $\int dx = \int dx_0 dx_1 dx_2 dx_3$ gains a $-i$ factor, $\frac{\partial}{\partial x_0}$ gains a $\frac{1}{-i} = i$ factor, and $A = \sum_{j=0}^3 A_j dx_j$ stays the same if we replace A_0 by iA_0. Thus F_{jk} has no change if $1 \le j, k \le 3$, and for the others

$$F_{0j} = \partial_0 A_j - \partial_j A_0 - [A_0, A_j]$$

changes to $i\partial_0 A_j - i\partial_j A_0 - i[A_0, A_j] = iF_{0j}$. So actually $\eta_i\eta_j F_{ij}(x)^2$ has all of its negative factors canceling out nicely, and what we end up with is (we are also supposed to multiply what we end up with by i)

$$S_E(A) = -\frac{1}{2g_0^2}\int_{\mathbb{R}^4} \mathrm{Tr}\left[\sum_{0\le i<j\le 3} F_{ij}(x)^2\right] dx.$$

So what we'll want to do with the Euclidean theory is to construct some kind of probability measure on the space of all A – that is, on all $\mathfrak{g}$-valued 1-forms – with density proportional to $\exp(-S_E(A))$ "with respect to Lebesgue measure." That's what we'll focus on in this course through a discretization.

Much like the discrete version of Brownian motion is random walk, lattice gauge theories are certain lattice approximations that may allow us to understand the limiting object. We'll now be working in general d-dimensional space instead of $\mathbb{R}^4$ as we did above – we can similarly make the definitions for A, F, $S(A)$, and $S_E(A)$ in $\mathbb{R}^d$.

Definition 3

Let G be as before and let $\Lambda \subseteq \mathbb{Z}^d$ be a finite set. Suppose that for any two adjacent vertices $x, y \in \Lambda$, there is a matrix $U(x, y) \in G$, and we impose the constraint that $U(y, x) = U(x, y)^{-1}$. Such an assignment of matrices is called a **configuration** for the theory.

Definition 4

A **plaquette** is a square bounded by four edges (this is the definition for any $d \ge 2$). For each plaquette $p \in \Lambda$ (meaning all four edges lie in Λ) with vertices x_1, x_2, x_3, x_4 in anticlockwise order (let x_1 be lexicographically smallest and x_2 the next smallest), define

$$U_p = U(x_1, x_2)U(x_2, x_3)U(x_3, x_4)U(x_4, x_1).$$

The **Wilson action** on this configuration is then

$$S_\Lambda(U) = \sum_{p\in\Lambda} \mathrm{Re}(\mathrm{Tr}(I - U_p)).$$

(We're using free boundary conditions here.)

To see how this relates to Euclidean Yang-Mills theories, fix some $\beta > 0$ and let $\mu_{\Lambda,\beta}$ be the probability measure on the space of configurations $G(\Lambda)$ (that is, assignments of matrices to the edges of the lattice) defined as

$$\mu_{\Lambda,\beta} = \frac{1}{Z} e^{-\beta S_\Lambda(U)} d\sigma_\Lambda(U),$$

where $d\sigma_\Lambda(U)$ is the product normalized Haar measure on $G(\Lambda)$ (this makes sense since G is a compact Lie group and $G(\Lambda) = G^{E(\Lambda)}$ for $E(\Lambda)$ the set of positively oriented edges in Λ) and Z is the appropriate normalizing constant.

Theorem 5

Let $A = \sum_{j=1}^{d} A_j dx_j$ be a compactly supported smooth $\mathfrak{g}$-valued 1-form on $\mathbb{R}^d$ (for any $d \geq 2$), and let $e_1, \cdots, e_d$ be standard basis vectors. We define a configuration on the lattice $\varepsilon\mathbb{Z}^d$ as follows: for any directed edge of the form $(x, x + \varepsilon_j)$, define

$$U(x, x + \varepsilon e_j) = e^{\varepsilon A_j(x)}$$

(and so $U(x+\varepsilon e_j, x) = U(x, x+\varepsilon e_j)^{-1}$). We can now define (because A is compactly supported the sum is finite)

$$S(U) = \sum_{p \in \varepsilon\mathbb{Z}^d} \mathrm{Re}(\mathrm{Tr}(I - U_p)).$$

Then

$$S_E(A) = \lim_{\varepsilon \to 0} g_0^{-2} \varepsilon^{d-4} S(U).$$

(Indeed, A_j is zero outside a compact set, so U_p are all the identity and the contribution to the total trace is zero outside of that compact set.) It's believed that we might actually need some $\log \varepsilon$ type factors in $d = 4$ (via renormalization group arguments), but it's not entirely clear how that actually works out.

We'll need a four-term version of the Baker-Campbell-Hausdorff formula:

Lemma 6

Let A and B be square (complex) matrices of order n with entries all of absolute value at most K. For all $t \in [0, 1]$, we have

$$e^{tA} e^{tB} = \exp\left(t(A+B) + \frac{t^2}{2}[A,B] + \frac{t^3}{12}([A,[A,B]] + [B,[B,A]]) - \frac{t^4}{24}([A,[B,[A,B]]]) + O(t^5) \right),$$

where $O(t^5)$ means a matrix whose entries have absolute value at most Ct^5, where C depends only on K and n.

(The actual Baker-Campbell-Hausdorff formula is an infinite series version of this which doesn't always converge, but our version here makes sense.)

Proof sketch. Call the quantity inside the exponential on the right-hand side $h(t)$, and let $g(t) = e^{h(t)}$. Since $h(t) = O(t)$, we have

$$g(t) = 1 + h(t) + \frac{h(t)^2}{2} + \frac{h(t)^3}{6} + \frac{h(t)^4}{24} + O(t^5),$$

and then plugging in $h(t)$ and keeping track of all coefficients verifies what we want. For example,

$$h(t)^2 = t^2(A+B)^2 + \frac{1}{2}t^3((A+B)[A,B] + [A,B](A+B)) + \cdots,$$

and if we collect all terms together we get the same expansion as if we write out $e^{tA}e^{tB}$ as a power series (for example, in both cases the coefficient of t^3 is $\frac{1}{6}(A^3 + 3A^2B + 3AB^2 + B^3)$). □

Lemma 7

Let $B_1, \cdots, B_m$ be skew-Hermitian matrices (we specialize like this because all of our matrices will be skew-Hermitian) with absolute values of entries at most K. Then for all $t \in [0,1]$,

$$e^{tB_1} \cdots e^{tB_m} = \exp\left(t \sum_{j=1}^{m} B_j + \frac{t^2}{2} \sum_{1 \le j < k \le m} [B_j, B_k] + R \right) + O(t^5),$$

where $R = O(t^3)$ and R is skew-Hermitian.

We'll in particular want to apply this to our plaquettes (so $m = 4$ when we multiply along the four edges).

Proof. We prove this by induction on m; for $m = 2$ this follows from the following lemma, since $[A, B]$ is skew-Hermitian if A, B are skew-Hermitian. And for the inductive step, we have by the inductive hypothesis

$$e^{tB_1} \cdots e^{tB_{m-1}} = \exp(tA) + O(t^5), \quad A = \sum_{j=1}^{m-1} B_j + \frac{t}{2} \sum_{1 \le j < k \le m-1} [B_j, B_k] + Q$$

for some skew-Hermitian Q with $Q = O(t^2)$. So now we can apply the previous lemma to $e^{tA} e^{tB_m}$ and get the result, since again commutators of A and B_m will all be skew-Hermitian. □

Corollary 8

Let $B_1, \cdots, B_m$ be as above. Then for all $t \in [0,1]$,

$$\mathrm{Re}(\mathrm{Tr}(I - e^{tB_1} \cdots e^{tB_m})) = -\frac{1}{2}\mathrm{Tr}\left[\left(t \sum_{i=1}^{m} B_i + \frac{t^2}{2} \sum_{1 \le i < j \le m} [B_i, B_j] \right)^2 \right] + \mathrm{Tr}\left[O(t^4) \sum_{i=1}^{m} B_i \right] + O(t^5).$$

Proof. From the previous lemma, we can again take a power series approach and write

$$e^{tB_1} \cdots e^{tB_m} = I + \sum_{k=1}^{4} \frac{1}{k!} (t \sum_i B_i + \frac{t^2}{2} \sum_{i<j} [B_i, B_j] + R)^k + O(t^5).$$

But the blue part is skew-Hermitian, hence has all imaginary eigenvalues; therefore the traces of its odd parts are pure imaginary and the traces of its even parts are real. So when we take the real part of the trace, only the $k = 2, 4$ terms will survive, and out of those the only parts that give us nontrivial contributions up to $O(t^4)$ are the ones listed above. (Indeed, the only cross-term involving R that may matter is between $t \sum B_i$ and R, which is indeed absorbed into the term $O(t^4) \sum B_i$.) □

Proof sketch of Theorem 5. Take some $x \in \varepsilon \mathbb{Z}^d$ and fix some $1 \le j < k \le d$. Consider the plaquette p being formed from $x_1 = x, x_2 = x + \varepsilon e_j, x_3 = x + \varepsilon e_j + \varepsilon e_k$, and $x_4 = x + \varepsilon e_k$. We have

$$\mathrm{Re}(\mathrm{Tr}(I - U_p)) = \mathrm{Re}(\mathrm{Tr}(I - U(x_1, x_2)U(x_2, x_3)U(x_3, x_4)U(x_4, x_1))),$$

and we defined $U(x, x + \varepsilon e_j) = e^{\varepsilon A_j(x)}$. So substituting these values in and applying our lemma, we can show that the approximation works – we'll do this next time. □

2 September 24, 2025

Last time, we set up the following framework: $A = \sum_{j=1}^{d} A_j dx_j$ is a smooth compactly-supported $\mathfrak{g}$-valued 1-form on $\mathbb{R}^d$, which we discretize by defining for any $x \in \varepsilon\mathbb{Z}^d$ the matrices

$$U(x, x + \varepsilon e_j) = e^{\varepsilon A_j(x)}$$

along the edges of our lattice. We were then trying to prove that $\lim_{\varepsilon\to 0} g_0^{-2}\varepsilon^{d-4}S(U)$ is the Euclidean Yang-Mills action $S_E(A)$, so that the Wilson action is a proper discretization.

Proof of Theorem 5. Writing out the proof sketch from last time in more detail, we again need to make use of Corollary 8. The Wilson action is the sum over all plaquettes of the lattice

$$S(U) = \sum_p \mathrm{Re}(\mathrm{Tr}(I - U_p))$$

(this sum is nonzero only finitely often, so this is well-defined). Like last time, we fix some $x \in \varepsilon\mathbb{Z}^d$ and consider p formed by $x_1 = x, x_2 = x + \varepsilon e_j, x_3 = x + \varepsilon e_j + \varepsilon e_k$, and $x_4 = x + \varepsilon e_k$. Our corollary then says that

$$\begin{aligned}\mathrm{Re}(\mathrm{Tr}(I - U_p)) &= \mathrm{Re}(\mathrm{Tr}(I - U(x_1, x_2)U(x_2, x_3)U(x_3, x_4)U(x_4, x_1))) \\ &= \mathrm{Re}\left(\mathrm{Tr}\left(I - e^{\varepsilon A_j(x_1)}e^{\varepsilon A_k(x_2)}e^{-\varepsilon A_j(x_4)}e^{-\varepsilon A_k(x_1)}\right)\right)\end{aligned}$$

where notice that we had to take the inverse for $U(x_3, x_4)$ and $U(x_4, x_1)$ because we're defining things in terms of the positively oriented edges. Applying Corollary 8 yields that

$$\begin{aligned}\mathrm{Re}(\mathrm{Tr}(I - U_p)) = -\frac{1}{2}\mathrm{Tr}\Bigg[\Bigg(&\varepsilon\Big(A_j(x) + A_k(x_2) - A_j(x_4) - A_k(x)\Big) \\ &+ \frac{\varepsilon^2}{2}\Big([A_j(x_1), A_k(x_2)] - [A_j(x_1), A_j(x_4)] - [A_j(x_1), A_k(x_1)] - [A_k(x_2), A_j(x_4)] - [A_k(x_2), A_k(x_1)] + [A_j(x_4), A_k(x_1)]\Big)\Bigg)^2\Bigg] \\ &+ \mathrm{Tr}\left[O(\varepsilon^4)(A_j(x) + A_k(x_2) - A_j(x_4) - A_k(x))\right] + O(\varepsilon^5).\end{aligned}$$

Now $A_j(x_1) = A_j(x)$, $A_k(x_1) = A_k(x)$, and then we linearize the other terms: $A_k(x_2) = A_k(x + \varepsilon e_j) = A_k(x) + \varepsilon\partial_j A_k(x) + O(\varepsilon^2)$, and similarly $A_j(x_4) = A_j(x) + \varepsilon\partial_k A_j(x) + O(\varepsilon^2)$. Thus the sum of all four of these terms is just $\varepsilon(\partial_j A_k(x) - \partial_k A_j(x)) + O(\varepsilon^2)$; in particular we see that we can absorb the $\mathrm{Tr}\left[O(\varepsilon^4)(A_j(x) + A_k(x_2) - A_j(x_4) - A_k(x))\right]$ term as part of the $O(\varepsilon^5)$ error. And similarly the commutators can be simplified: $[A_j(x_1), A_k(x_2)]$ is just $[A_j(x), A_k(x)] + O(\varepsilon)$, and if we do this with all six commutators the leading contributions add up to $2[A_j(x), A_k(x)]$. Thus plugging everything in,

$$\mathrm{Re}(\mathrm{Tr}(I - U_p)) = -\frac{1}{2}\mathrm{Tr}\left[\left(\varepsilon^2\Big(\partial_j A_k(x) - \partial_k A_j(x)\Big) + \varepsilon^2[A_j(x), A_k(x)] + O(\varepsilon^3)\right)^2\right] + O(\varepsilon^5).$$

So when we expand out the squares, all errors are of order at least ε^5, and we indeed get

$$\mathrm{Re}(\mathrm{Tr}(I - U_p)) = -\frac{\varepsilon^4}{2}\mathrm{Tr}\left[(\partial_j A_k(x) - \partial_k A_j(x) + [A_j(x), A_k(x)])^2\right] + O(\varepsilon^5) = -\frac{\varepsilon^4}{2}\mathrm{Tr}(F_{jk}(x)^2) + O(\varepsilon^5),$$

meaning that substituting back into the action,

$$\varepsilon^{d-4} S(U) = -\frac{\varepsilon^d}{2} \sum_{x \in \mathbb{Z}^d} \sum_{1 \le j < k \le d} \mathrm{Tr}(F_{jk}(x)^2) + O(\varepsilon^{d+1}).$$

But we're summing over all plaquettes in $\varepsilon\mathbb{Z}^d$ corresponding to a bounded region of nonzero $(I - U_p)$s, and the number of nonzero terms is $O(\varepsilon^{-d})$; hence the error term is $O(\varepsilon)$. And the main term is just a Riemann sum approximation converging to $\int \sum_{j,k} \mathrm{Tr}(F_{jk}(x)^2)$, which is the Yang-Mills action as desired if we add in the factor of g_0^{-2}. □

Remark 9. *This is a similar argument to what appears when going between Brownian motion and Gaussian random walk, but the trouble is that we can't make a "Cameron-Martin" type argument to go from smooth A to general A since things aren't Gaussian here.*

We'll turn now to gauge transforms:

Definition 10

Let $A = \sum_j A_j dx_j$ be a $\mathfrak{g}$-valued 1-form. A **smooth gauge transform** U is a smooth map from $\mathbb{R}^d$ into G. (In the language of differential geometry, consider the trivial principal bundle and take a section.)

Gauge transforms form a group under pointwise multiplication, and the group acts on smooth 1-forms in the following way: U acts on A as

$$A^U(x) = U(x)A(x)U(x)^{-1} - (dU(x))U(x)^{-1}.$$

More explicitly in coordinates, we have

$$A^U = \sum_{j=1}^{d} A_j^U(x) dx_j, \quad A_j^U(x) = U(x)A_j(x)U(x)^{-1} - (\partial_j U(x))U(x)^{-1}.$$

(That is, $\partial_j U(x) = \lim_{\varepsilon \to 0} \frac{1}{\varepsilon}(U(x + \varepsilon e_j) - U(x))$ as matrices.) We proved last year that this is actually an element of the Lie algebra for any smooth gauge transform, so A^U is another smooth 1-form and in fact (by a calculation) $(A^U)^V = A^{VU}$.

We'll now want to similarly get a gauge transform for the discrete setting and justify why it is a proper discretization. For any finite lattice $\Lambda \subseteq \mathbb{Z}^d$ and lattice gauge theory on Λ, the group of gauge transforms is just G^Λ (that is, no conditions on smoothness – we just assign a group element to each vertex). For a configuration $U \in G(\Lambda)$ and a gauge transform $g \in G^\Lambda$, we define

$$U^g(x, y) = g(x)U(x, y)g(y)^{-1}.$$

So this is a much more transparent description of the action of a gauge transform: the new matrix attached to the edge pre-multiplies by g at the left endpoint and post-multiplies by g at the right endpoint. By a direct calculation we have that because $U(y, x) = U(x, y)^{-1}$,

$$U^g(y, x) = g(y)U(y, x)g(x)^{-1} = (g(x)U(x, y)g(y)^{-1})^{-1} = U^g(x, y)^{-1},$$

and we can also similarly check that $(U^g)^h = U^{hg}$. The last thing we'll do for the continous-to-discrete description is show that these gauge transforms are the right discretization as well.

Recall that for any $U \in U(N)$, we can write $U = VDV^*$ for some unitary matrix $V \in U(n)$ and some diagonal matrix $D = \mathrm{diag}(e^{i\theta_1}, \cdots, e^{i\theta_n})$ with all $\theta_i \in (-\pi, \pi)$ (and this is unique up to permuting things) – this is the spectral decomposition. We can then define

$$\log U = V \mathrm{diag}(i\theta_1, \cdots, i\theta_n)V^*$$

uniquely because of the constraints above.

Theorem 11

Let A be a smooth compactly supported $\mathfrak{g}$-valued 1-form on $\mathbb{R}^d$, and let $g \in C^\infty(\mathbb{R}^d, G)$ be a smooth gauge transform that is the identity matrix I outside a compact region. Define the discretization on the lattice as before with $U(x, x+\varepsilon e_j) = e^{\varepsilon A_j(x)}$. Letting A^g, U^g be as before (so A^g is defined in the continuous way, and U^g is defined in the discrete way). If we let $B_j(x) = \frac{1}{\varepsilon} \log U^g(x, x+\varepsilon e_j)$, then

$$\sup_{x \in \varepsilon\mathbb{Z}^d, 1 \le j \le d} \left|\left| A_j^g(x) - B_j(x) \right|\right| = 0.$$

(Here, the norm denotes Euclidean norm on the entries of the matrix.)

(In fact, the norm will be $O(\varepsilon)$, as we'll see in the proof.) To prove this result, we'll need the following lemma:

Lemma 12

Let $U \in U(n)$ be such that $||U - I|| \le \frac{1}{2}$. Then $\log U = \sum_{k=1}^\infty \frac{(I-U)^k}{k}$, and the series converges absolutely with respect to the norm.

Proof. The absolute convergence is clear because $||UV|| \le ||U|| \cdot ||V||$ and thus the sum converges exponentially fast. To show that it's actually equal to $\log U$ as we defined before, write $U = VDV^*$ so that $\log U = V\Lambda V^*$ with $\Lambda = \operatorname{diag}(i\theta_1, \cdots, i\theta_n)$. Letting log now denote the analytic branch of the logarithm in $\mathbb{C} \setminus (-\infty, 0]$ agreeing with ln on $[0, \infty)$, we have a power series at 1 given by

$$\log z = -\sum_{k=1}^\infty \frac{(1-z)^k}{k}$$

with radius of convergence 1. So

$$\begin{aligned} ||U - I||^2 &= \operatorname{Tr}((U-I)^*(U-I)) \\ &= \sum_{j=1}^n |1 - e^{i\theta_j}|^2 \end{aligned}$$

by the spectral decomposition, and we know the left-hand side is at most $\frac{1}{4}$. Thus each term is at most $\frac{1}{4}$, meaning each $e^{i\theta_j}$ is within $\frac{1}{2}$ of 1 and in particular none of them are equal to $-\pi$. Thus by the choice of branch of log, we get $i\theta_j = \log e^{i\theta_j} = -\sum_{k=1}^\infty \frac{(1-e^{i\theta_j})^k}{k}$. So the rest of the proof just follows from the spectral decomposition and writing out each term $(I-U)^k$ in the sum in terms of V and our diagonal matrices. □

Proof of Theorem 11. Take any $x \in \varepsilon\mathbb{Z}^d$ and $1 \le j \le d$. By definition we have

$$\begin{aligned} U^g(x, x+\varepsilon e_j) &= g(x)U(x, x+\varepsilon e_j)g(x+\varepsilon e_j)^{-1} \\ &= g(x)e^{\varepsilon A_j(x)}g(x+\varepsilon e_j)^{-1} \\ &= g(x)(I + \varepsilon A_j(x) + O(\varepsilon^2))g(x+\varepsilon e_j)^{-1} \end{aligned}$$

(here again the constant in $O(\varepsilon^2)$ doesn't depend on x or ε because things are compactly supported); expanding this out yields

$$U^g(x, x+\varepsilon e_j) = g(x)g(x+\varepsilon e_j)^{-1} + \varepsilon g(x)A_j(x)g(x+\varepsilon e_j)^{-1} + O(\varepsilon^2).$$

We know that $g(x)g(x)^{-1} = I \implies \partial_j(g(x))g(x)^{-1} + g(x)\partial_j(g(x)^{-1}) = 0$ by taking a derivative, so $\partial_j(g(x)^{-1}) = -g(x)^{-1}(\partial_j g(x))g(x)^{-1}$. Thus linearizing, we get

$$g(x+\varepsilon e_j)^{-1} = g(x)^{-1} - \varepsilon g(x)^{-1}(\partial_j g(x))g(x)^{-1} + O(\varepsilon^2),$$

and we don't have to worry about singularities when taking inverses because everything is a unitary matrix here. Substituting this back in, we get

$$\begin{aligned} U^g(x, x+\varepsilon e_j) &= g(x)\left[g(x)^{-1} - \varepsilon g(x)^{-1}\partial_j(g(x))g(x)^{-1}\right] + \varepsilon g(x)A_j(x)g(x)^{-1} + O(\varepsilon^2) \\ &= I - \varepsilon\partial_j(g(x))g(x)^{-1} + \varepsilon g(x)A_j(x)g(x)^{-1} + O(\varepsilon^2) \\ &= I + \varepsilon A_j^g(x) + O(\varepsilon^2). \end{aligned}$$

But by definition $U^g(x, x+\varepsilon e_j)$ is a unitary matrix and it's very close to the identity, and the $O(\varepsilon^2)$ bounds are uniform in x. So for ε small enough we can apply our lemma above, and we get

$$\log U^g(x, x+\varepsilon e_j) = -\sum_{k=1}^{\infty} \frac{(-\varepsilon A_j^g(x) + O(\varepsilon^2))^k}{k} = \varepsilon A_j^g(x) + O(\varepsilon^2),$$

as desired. □

When people tried to get convergence to continuum limits in constructive field theory, they tried to choose gauge transforms to make all of the bond variables close to the identity. But this program has not been completed yet, and it's one of the first steps for the Yang-Mills existence problem. We'll see more of this later on.

Now that we've defined gauge transforms, we can understand gauge-invariant observables and gauge fixing. For $\Lambda \subseteq \mathbb{Z}^d$ and a function $f : G(\Lambda) \to \mathbb{C}$, we can define the **expected value** of f as

$$\langle f \rangle = \int_{G(\Lambda)} f(U) d\mu_{\Lambda,\beta}(x).$$

We say that an observable is **gauge-invariant** if $f(U^g) = f(U)$ for all U, g (that is, f is constant on orbits of the gauge transform). There's a physical reason to care only about these observables, but mathematically we also care because we can always produce gauge-invariant observables with the same expectation as a given observable with a certain prescription.

Definition 13

A **loop** in $\mathbb{Z}^d$, usually denoted ℓ, is a sequence of points $x_0, x_1, \cdots, x_k$ such that x_j, x_{j+1} are neighbors for all $0 \le j \le k-1$ and x_0 is a neighbor of x_k. The **Wilson loop observable** for a loop entirely contained within Λ is defined as

$$W_\ell(U) = \mathrm{Tr}(U(x_0, x_1)U(x_1, x_2)\cdots U(x_{k-1}, x_k)U(x_k, x_0)).$$

There is a continuous analog of this definition (where we take the gauge connection A and perform parallel transport of A around a loop), though it is a bit more complicated to define since it's the Wick-ordered exponential of the connection rather than just the exponential. And this quantity is indeed gauge-invariant, because

$$W_\ell(U^g) = \mathrm{Tr}(g(x_0)U(x_0, x_1)g(x_1)^{-1}g(x_1)U(x_1, x_2)g(x_2)^{-1}\cdots g(x_k U(x_k, x_0)g(x_0)^{-1}),$$

and now all of the g factors cancel out except the first and last and trace is invariant under conjugation, meaning we end up with $W_\ell(U)$ as desired.

3 September 29, 2025

We started discussed gauge-invariant observables last time, and specifically we considered the Wilson loop observable $W_\ell(U)$ associated to a loop ℓ in $\mathbb{Z}^d$. For a plaquette p, note that $\mathrm{Tr}(U_p)$ is itself a Wilson loop observable, so the Wilson action $S_\Lambda(U) = \sum_p \mathrm{Re}(\mathrm{Tr}(I - U_p))$ is therefore gauge-invariant. This implies that the measure $\mu_{\Lambda,\beta}$ for the lattice gauge theory is also invariant under gauge transforms; that is,

$$\int f(U^g) d\mu_{\Lambda,\beta}(U) = \int f(U) d\mu_{\Lambda,\beta}(U).$$

Indeed, we can see this by letting σ_Λ denote the product (normalized) Haar measure on $G(\Lambda)$, so that we always have

$$\int f(U^g) d\sigma_\Lambda(U) = \int f(U) d\sigma_\Lambda(U)$$

since the gauge transform retains the independence and the distribution on each vertex of Λ. Therefore it is also true that

$$\begin{aligned}\int f(U^g) e^{-\beta S_\Lambda(U)} d\sigma_\Lambda(U) &= \int f(U^g) e^{-\beta S_\Lambda(U^g)} d\sigma_\Lambda(U) \\ &= \int f(U) e^{-\beta S_\Lambda(U)} d\sigma_\Lambda(U),\end{aligned}$$

by first using that S is gauge-invariant and then using our identity above. Dividing through by the normalizing constant Z then yields the result.

Remark 14. *Lattice models are "very real models" in a certain way: people use simulations of some other lattice models to estimate the radius of elementary particles, among other things. And these are the best available approximations to the true values verified by experiments. (For example, $SU(3)$ lattice gauge theory should describe the theory of hadrons or quarks; the pure gauge theory is just about the strong force, but if we add the particles and put in more things we do get things that aren't just toy models.) The point is also to construct a non-perturbative Yang-Mills theory because perturbation theory has its limits, and our best available understanding of that comes through these lattice gauge theories as discrete approximations.*

Example 15

We also mentioned last time that it "suffices mathematically to look at gauge-invariant functions," which we'll elaborate on now.

Let $f : G(\Lambda) \to \mathbb{C}$ be any function such that $\langle |f| \rangle < \infty$. We can then define a new function which is the "random gauge transform"

$$h(U) = \int f(U^g) \prod_{x \in \Lambda} d\sigma(g(x)).$$

(The notation $d\sigma(g(x))$ means that each coordinate of our transform is distributed according to the Haar measure.) We claim that h is gauge-invariant and also that $\langle h \rangle = \langle f \rangle$ under the lattice gauge theory. Indeed, we have

$$h(U^g) = f((U^g)^{g'}) \prod_{x \in \Lambda} d\sigma(g'(x)) = f(U^{g'g}) \prod_{x \in \Lambda} d\sigma(g'(x)),$$

but since the Haar measure is left- and right-invariant this is the same as $f(U^{g'}) \prod_{x \in \Lambda} d\sigma(g'(x)) = h(U)$. (Intuitively, applying a gauge transform and then applying a random one is just a random gauge transform.) And to show that h

has the same expectation as f, we compute

$$\langle h \rangle = \frac{1}{Z} \int h(U) e^{-\beta S_\Lambda(U)} d\sigma_\Lambda(U) = \frac{1}{Z} \int \left(\int f(U^g) \prod_{x \in \Lambda} d\sigma(g(x)) \right) e^{-\beta S_\Lambda(U)} d\sigma_\Lambda(U).$$

Since f has finite expectation under the lattice gauge theory, we can swap the order of integration and use Fubini's theorem to get

$$\frac{1}{Z} \int \left(\int f(U^g) e^{-\beta S_\Lambda(U)} d\sigma_\Lambda(U) \right) \prod_{x \in \Lambda} d\sigma(g(x))$$

But we just showed that the lattice gauge theory is invariant under gauge transforms, so this just simplifies to (replacing U^g with U)

$$\frac{1}{Z} \int \int f(U) e^{-\beta S_\Lambda(U)} d\sigma_\Lambda(U) \prod d\sigma(g(x)) = \langle f \rangle \int \prod d\sigma(g(x)),$$

which is just $\langle f \rangle$ since the Haar measure is normalized.

Last year, we talked about some gauge-fixing procedures in the continuum, and now we'll talk about a more unified way of doing so in the discrete setting. The following is the main result that will be helpful for us:

Proposition 16

Let Γ be a subset of the positively oriented edges $E(\Lambda)$ such that the undirected versions of the edges in Γ form no loops. Let $G(\Lambda, \Gamma)$ be the set of configurations $U \in G(\Lambda)$ where $U(x, y) = I$ for all $(x, y) \in \Gamma$. Then defining

$$d\mu_{\Lambda,\Gamma,\beta}(U) = \frac{1}{Z_\Gamma} \exp\left(-\beta S_\Lambda(U)\right) d\sigma_{\Lambda,\gamma}(U),$$

where $d\sigma_{\Lambda,\gamma}(U) = \prod_{(x,y) \in E(\Lambda) \setminus \Gamma} d\sigma(U(x, y))$ is product Haar measure on all edges except those in Γ. Then for any gauge-invariant f with $\langle |f| \rangle < \infty$, we have

$$\langle f \rangle = \int_{G(\Lambda,\Gamma)} f(U) d\mu_{\Lambda,\Gamma,\beta}(U)$$

In other words, we take our original measure and condition on all edges in Γ to be the identity, and that still gets us valid expectations while using a simpler theory. And this has important consequences: in lattice gauge theory we often care about mass gap and quark confinement, both of which hold for two-dimensional theories. And indeed any two-dimensional theory can be reduced to a collection of independent Markov chains, which is what we'll use.

Proof. Let $V(x, y)$ be iid (random matrices) from σ for all $(x, y) \in E(\Lambda) \setminus \Gamma$. Extend V to an element of $G(\Lambda, \Gamma)$ by setting all $V(x, y) = I$ for $(x, y) \in \Gamma$. Let $g(x)$ be iid from σ for all $x \in \Lambda$ (so that we have a random gauge transform), and let $W = V^g$ be that gauge transform applied to V. We claim that the matrices $W(x, y)$ are now all iid from σ for all $(x, y) \in E(\Lambda)$.

To prove this, we first show that for $(x, y) \in \Gamma$, the $W(x, y)$ are iid from σ. Indeed,

$$W(x, y) = g(x) V(x, y) g(y)^{-1} = g(x) g(y)^{-1},$$

and we claim we can order the elements of Γ as $(x_1, y_1), (x_2, y_2), \cdots$ such that for all n, at most one of x_n and y_n is among the previous vertices $x_1, y_1, x_2, y_2, \cdots, x_{n-1}, y_{n-1}$. Indeed, if the undirected graph induced by Γ is connected (meaning it is a tree), then we first let (x_1, y_1) be any element of Γ, and then once we've chosen the first $(n-1)$ elements in the ordering, we choose (x_n, y_n) from among the remaining edges (if any) such that **at least** one of x_n, y_n is in the set of previously observed vertices. By connectedness we will always have such an edge available until Γ is exhausted, since the undirected graph formed by the first n edges is always connected. And we can't have both x_n, y_n

in the set of previously observed vertices, because otherwise we would form a cycle (by first taking the previous path between x_n and y_n in that connected graph so far, and then connecting that with our edge (x_n, y_n)). So we can indeed do the ordering we requested for a single connected component. And finally, if the undirected graph is not connected, we just separate the edges into connected components and do them one after another.

The point of this ordering is to show that each $W(x_n, y_n)$ is independent of all previous ones, and that each $W(x, y)$ is distributed iid from σ, so $(W(x, y))_{(x,y)\in\Gamma}$ is indeed overall iid from σ. So our goal is to prove that conditionally given $W(x_1, y_1), \cdots, W(x_{n-1}, y_{n-1})$, the matrix $W(x_n, y_n)$ is always distributed according to σ. For this, it suffices to prove the stronger statement that given $g(x_1), g(y_1), \cdots, g(x_{n-1}), g(y_{n-1})$ (which together determine all previous $W(x_i, y_i)$s), we have $W(x_n, y_n) \sim \sigma$. But now $W(x_n, y_n) = g(x_n)g(y_n)^{-1}$; if both x_n and y_n are both not in the previous set of vertices then $g(x_n)g(y_n)^{-1}$ is Haar distributed and independent of the rest, and even if one of x_n and y_n are in the set but the other is not, then $g(x_n)g(y_n)^{-1}$ still conditionally follows σ by the invariance of the Haar measure. This proves our claim that $W(x, y)$ are iid from σ for all $(x, y) \in \Gamma$.

For the remaining edges, note that for $(x, y) \in E(\Lambda) \setminus \Gamma$, we have

$$W(x, y) = g(x)V(x, y)g(y)^{-1}$$

and the $V(x, y)$s are iid from the Haar measure. So conditional on the gauge transform g, $\{W(x, y)\}_{(x,y)\in E(\Lambda)\setminus\Gamma}$ are also iid from σ, again by invariance. So in fact these are independent of g (while the $W(x, y)$s in Γ are deterministic functions of g), and thus all W are indeed iid from γ.

To complete the proof of the theorem, let $h : G(\Lambda) \to \mathbb{C}$ be any gauge-invariant function. We then have

$$\int_{G(\Lambda,\Gamma)} h(U)d\sigma_{\Lambda,\Gamma(U)} = \mathbb{E}[h(V)] = \mathbb{E}[h(V^g)]$$

by gauge-invariance of h, and this last quantity is exactly $\mathbb{E}[h(W)]$. But W is iid Haar, so this last quantity is exactly $\int_{G(\Lambda)} h(U)d\sigma_\Lambda(U)$.

In particular, if we take any gauge-invariant $f : G(\Lambda) \to \mathbb{C}$ and let $h(U) = f(U)e^{-\beta S_\Lambda(U)}$, then h is also gauge-invariant and thus $\int_{G(\Lambda,\Gamma)} f(U)e^{-\beta S_\Lambda(U)}d\sigma_{\Lambda,\Gamma}(U) = \int_{G(\Lambda)} f(U)e^{-\beta S_\Lambda(U)}d\sigma_\Lambda(U)$. And there's one last step to complete the proof: taking f to be identically 1 yields $Z_\Gamma = Z$, so that dividing by the appropriate normalizing constants on each side proves the desired result. □

We'll now turn to mass gap and quark confinement, and we won't go into the physical details of what exactly these mean for now, only the mathematical formulations.

Fact 17

Mass gap is equivalent to **exponential decay of correlations** – recall from last year that to construct a quantum field theory, we need a self-adjoint operator H. We then say that we have a mass gap if there is a positive spectral gap between 0 and the next smallest eigenvalue. We then get a quantum operator e^{-itH} with a corresponding Markov process with semigroup e^{-tH}, and we get exponential mixing if the gap is of positive size. And in fact the connection between "exponential mixing" and "exponential decay" comes in gauge-fixing – if we fix everything going in a particular coordinate direction (which we can call "time"), the (time-)slices evolving over that coordinate yield a Markov process.

On the other hand, we can also have theories with massless particles which will then not have a mass gap. So we'll discuss all that later.

Definition 18

Let $\Lambda_N \uparrow \mathbb{Z}^d$, and consider the lattice gauge theory on Λ_N. Let $f : G(\Lambda_N) \to \mathbb{C}$ be a gauge-invariant observable; we say that f is **supported on a set of directed edges E** if $f(U)$ has no dependence on $U(x,y)$ for any $(x,y) \notin E$ (so for example a plaquette variable is supported only on its four bordering edges). We then say the theory (that is, the sequence of theories on $\Lambda_1, \Lambda_2, \cdots$) has **exponential decay of correlations** if for all gauge-invariant bounded functions f_1, f_2 supported on finite $S_1, S_2 \subseteq \mathbb{Z}^d$ and all N large enough so that $S_1, S_2 \subseteq \Lambda_N$, we have

$$|\langle f_1 f_2 \rangle - \langle f_1 \rangle \langle f_2 \rangle| \leq C_1 e^{-C_2 \text{dist}(S_1, S_2)},$$

where C_1 may depend on $\beta, d, G, f_1, f_2, |S_1|, |S_2|$ but not N, while C_2 may depend on β, d, G but not f_1, f_2, or N. Here $\text{dist}(S_1, S_2)$ is the minimum graph distance between points in the two sets.

4 October 1, 2025

Last time, we began discussing exponential decay of correlations (also called mass gap). We often discuss mass gap in terms of infinite-volume Gibbs measures, but we stated our version in Definition 18 to avoid some of those technical details. (From this version, we can essentially prove that the infinite-volume Gibbs measure is unique, in that it uniquely determines the gauge-invariant expectations, since we can think about how expectations depend on boundary values on the box and show that the contributions of those effects are small. But we're trying to not introduce notation for boundary conditions here for simplicity.)

The Yang-Mills mass gap problem requires us to prove that the continuous-time Hamiltonian has a gap in the spectrum, which is more complicated than this discrete statement. But even for something like $SU(3)$ lattice gauge theory on $\mathbb{Z}^4$ (or any nonabelian gauge group – it's expected not to be true for abelian gauge groups), it would be a big deal and a big step towards proving that conjecture. It's open to try to prove this for large enough β, and we'll see how to do so for small β soon.

Remark 19. *The reason that decay of correlations are called "mass" is that if we consider the massive scalar free field (which is the quantum field theory which describes free bosons with mass m), the mass is exactly the gap in the spectrum in that model. And in systems of particles in general, that generalizes in some way. We can add mass terms to our lattice gauge theories too, but it breaks the gauge symmetry.*

We'll prove this in two dimensions shortly, where things are much easier to do. First, though, there's also a related problem of quark confinement, which we'll state as a math problem:

Definition 20

Let $\Lambda_N \uparrow \mathbb{Z}^d$, and consider the lattice gauge theory on Λ_N. We say that the theory (that is, the sequence of theories) **confines quarks** if there exists a function $V : (0, \infty) \to \mathbb{R}$ with $V(x) \to \infty$ as $x \to \infty$, such that for any rectangular loop $\ell \subseteq \Lambda_N$ with side lengths $R \leq T$, we have that

$$|\langle W_\ell \rangle| \leq C e^{-V(R)T},$$

with C depending only on β, d, G. In the special case where $V(R) \propto R$, this identity is called **Wilson's area law** (since we then have exponential decay of expectations dictated by the area enclosed by ℓ).

The idea is that in physics, we don't understand why quarks and antiquarks are always bound together rather than freely existing in nature. The hypothesis is that the potential energy between those objects must grow to infinity as $R \to \infty$ so that conservation of energy does not allow them to move apart too much, and Wilson showed in quantum field theory calculations that if V is a potential of a quark-antiquark pair separated by R, then this Wilson loop expectation should behave as written as T grows. (This comes from taking the quantum Yang-Mills model, introducing the quark-antiquark pair as fermions via Grassmann variables, and then taking the masses to infinity.)

This is also an open problem – both this and mass gap do not have convincing physics arguments beyond numerical simulations in four dimensoins.

Fact 21

Remember that the product of matrices along a Wilson loop is always unitary, so the product itself is not actually getting smaller and smaller. But the point is that when we take the expectation of the trace, it will indeed get quite small – we can prove that under fairly general conditions we have $|\langle W_\ell \rangle| \leq C_1 e^{-c_2(R+T)}$. This is because by conditioning on all matrices outside the loop, the remaining matrices are almost independent and have some uniformly bounded densities with respect to the Haar measure. Of course, this "perimeter law" is weaker than the conjectured "area law."

We'll prove that the area law always holds for $d = 2$, and we'll also be able to prove that it holds in general dimension for small enough β. But the important open question is again for $SU(3)$ theory on $\mathbb{Z}^4$ and large enough β, or even any non-abelian theory in place of $SU(3)$. And real-world models are supposed to correspond to very large β in some sense (in fact $\beta \to \infty$ in ideal models).

Fact 22

Professor Chatterjee has a theorem that a stronger version of mass gap implies confinement, under the assumption that G has nontrivial center (which is a necessary condition). But both mass gap and quark confinement seem to be very hard, and it's not clear which one is easier to prove.

We'll now turn to the proofs for $d = 2$ – we'll need some notation related to Markov chains.

Example 23

Let $(S, \mathcal{S})$ be a measurable space, and let μ be a (not necessarily finite) measure on this space. Let $X_0, X_1, \cdots$ be a Markov chain on S with transition kernel p with respect to μ, meaning that

$$\mathbb{P}(X_{n+1} \in A | X_n = x) = \int_A p(x, y) d\mu(y).$$

(If we're concerned about conditioning on x here, we can in fact make sense of that via regular conditional probabilities on Polish spaces, which we're always working on.) Let $p^n(x, y)$ be the n-step transition density (that is, the probability density function with respect to μ of X_n given $X_0 = x$), so that (by Chapman-Kolmogorov)

$$p^n(x, y) = \int_S \cdots \int_S p(x, x_1) p(x_1, x_2) \cdots p(x_{n-2}, x_{n-1}) p(x_{n-1}, y) d\mu(x_1) \cdots d\mu(x_{n-1}).$$

Let ν be an invariant probability measure for this chain, and assume that ν has some density f with respect to μ, meaning that $\int_S f(x) p(x, y) d\mu(x) = f(y)$.

The following result gives us a way to show exponentially fast mixing of such a Markov chain:

Theorem 24 (Doeblin condition, special case)
Take the notation above. Suppose there is some $\varepsilon > 0$ such that the transition density satisfies $p(x, y) \geq \varepsilon f(y)$ for all x, y (this is true for example for random walk on a compact group). Then for all $x \in S$, we have

$$\int_S |p^n(x, y) - f(y)| d\mu(y) \leq 2(1 - \varepsilon)^n.$$

Continuous state spaces can be difficult to analyze, so this is a particularly useful result in that setting. And we can see from the bound that the proof is not too complicated:

Proof. Let $q(x, y) = \frac{p(x,y) - \varepsilon f(y)}{1-\varepsilon}$. Then $q(x, y)$ is nonnegative and

$$\int q(x, y) d\mu(y) = \frac{1}{1-\varepsilon} \left(\int p(x, y) d\mu(y) - \varepsilon \int f(y) d\mu(y) \right) = \frac{1}{1-\varepsilon}(1 - \varepsilon) = 1,$$

so q is also a transition density (though we won't use this). Defining $q^n(x, y)$ in the same way that we defined $p^n(x, y)$ as an n-step transition, we can similarly show that $\int_S q^n(x, y) d\mu(y) = 1$ for all n and x.

We claim that f is also an invariant density for q; that is, for all y we have $\int_S f(x) q(x, y) d\mu(y) = f(y)$. Indeed,

$$\begin{aligned} \int_S f(x) q(x, y) d\mu(x) &= \frac{1}{1-\varepsilon} \left(\int_S f(x) p(x, y) d\mu(x) - \varepsilon \int_S f(x) f(y) d\mu(x) \right) \\ &= \frac{1}{1-\varepsilon}(f(y) - \varepsilon f(y)) \\ &= f(y). \end{aligned}$$

Thus, we claim that we can actually write

$$p^n(x, y) = (1 - \varepsilon)^n q^n(x, y) + (1 - (1 - \varepsilon)^n) f(y).$$

We can prove this by induction: by definition it's true for $n = 1$, and now if this holds for $n = k - 1$, we have

$$\begin{aligned} p^k(x, y) &= \int_S p^{k-1}(x, z) p(z, y) d\mu(z) \\ &= \int_S \left((1 - \varepsilon)^{k-1} q^{k-1}(x, y) + (1 - (1 - \varepsilon)^{k-1}) f(z) \right) \left((1 - \varepsilon) q(z, y) + \varepsilon f(y) \right) d\mu(z). \end{aligned}$$

Expanding out the cross-terms, the product of qs gets us $(1 - \varepsilon)^k q^k(x, y)$. But then all other terms, by the fact that f is an invariant measure for q, just yield constant multiples of $f(y)$, so that summing up the coefficients yields the inductive hypothesis.

And now we're done, since plugging in this expression for p^n into the left-hand side of the theorem yields

$$\int_S |p^n(x, y) - f(y)| d\mu(y) = \int_S (1 - \varepsilon)^n (q^n(x, y) - f(y)) d\mu(y),$$

and by the triangle inequality this is bounded by

$$(1 - \varepsilon^n) \int |(q^n(x, y)| + |f(y)|) d\mu(y) = (1 - \varepsilon)^n (1 + 1)$$

because q^n is nonnegative (this is the only place where we use this fact), which is the desired claim. □

In the language of total variation distance, this result is saying that $(1 - \varepsilon)^n$ is a bound for TV between p^n and the stationary distribution. And probabilistically, the idea with this proof is that at each step we either toss an ε-coin to

end up at the stationary distribution or take a step from q; once we end up heads we are automatically at stationarity.

Corollary 25

With the notation and condition in the theorem above, let g_1 be the joint density of $(X_0, X_1, \cdots, X_i)$ and g_2 the joint density of $(X_j, \cdots, X_k)$, where $i < j < k$. Also let g be the joint density of $X_0, \cdots, X_i, X_j, \cdots, X_k)$. (All densities are with respect to the product μ measure.) Assume that X_0 has some density h. Then if i and j are far apart, g is "approximately" $g_1 \times g_2$ in the following total variation sense:

$$\int_S \cdots \int_S |g(x_0, \cdots, x_i, x_j, \cdots, x_k) - g_1(x_0, \cdots, x_i) g_2(x_j, \cdots, x_k)| d\mu(x_0) \cdots d\mu(x_i) d\mu(x_j) \cdots d\mu(x_k) \leq 4(1-\varepsilon)^{j-i}.$$

Proof. We have $g_1(x_0, \cdots, x_i) = h(x_0)p(x_0, x_1) \cdots p(x_{i-1}, x_i)$, and we also have

$$g(x_0, \cdots, x_i, x_j, \cdots, x_k) = h(x_0)p(x_0, x_1) \cdots p(x_{i-1}, x_i)p^{j-i}(x_i, x_j)p(x_j, x_{j+1}) \cdots p(x_{k-1}, x_k)$$

(note the big jump in the middle). The joint density of the latter random variables is then $g_2(x_j, \cdots, x_k) = w(x_j)p(x_j, x_{j+1}) \cdots p(x_{k-1}, x_k)$, where $w = \int_S h(X_0)p^j(x_0, x)d\mu(x_0)$ is the density of X_j. Thus we only have one term that's different when we substitute things in:

$$\begin{aligned} &|g(x_0, \cdots, x_i, x_j, \cdots, x_k) - g_1(x_0, \cdots, x_i) g_2(x_j, \cdots, x_k)| \\ &= \left| h(x_0)p(x_0, x_1) \cdots p(x_{i-1}, x_i)(p^{j-i}(x_i, x_j) - w(x))p(x_j, x_{j+1}) \cdots p(x_{k-1}, x_k) \right|, \end{aligned}$$

so that if we integrate both sides what we need to bound is

$$\int_S \cdots \int_S h(x_0)p(x_0, x_1) \cdots p(x_{i-1}, x_i)|p^{j-i}(x_i, x_j) - w(x_j)|p(x_j, x_{j+1}) \cdots p(x_{k-1}, x_k) d\mu(x_0) \cdots d\mu(x_i) d\mu(x_j) \cdots d\mu(x_k).$$

If we integrate out x_k, we just have a probability kernel and so that variable disappears. Then we can successively integrate out all other variables up to x_{j+1}, so that all we have left is

$$\int_S \cdots \int_S h(x_0)p(x_0, x_1) \cdots p(x_{i-1}, x_i)|p^{j-i}(x_i, x_j) - w(x_j)| d\mu(x_0) \cdots d\mu(x_i) d\mu(x_j).$$

It suffices to show now that for all x_i,

$$\int |p^{j-i}(x_i, x_j) - w(x_j)| d\mu(x_j) \leq 4(1-\varepsilon)^{j-i},$$

because once we substitute that in we can integrate out all of the remaining variables as well. We already know that $\int |p^{j-i}(x_i, x_j) - f(x_j)| d\mu(x_j) \leq 2(1-\varepsilon)^{j-i}$, so it suffices to prove that $\int |w(x_j) - f(x_j)| d\mu(x_j) \leq 2(1-\varepsilon)^{j-i}$; in fact we can show that $\int |w(x_j) - f(x_j)| d\mu(x_j) \leq 2(1-\varepsilon)^j$. This is because

$$w(x_j) = \int h(x_0)p^j(x_0, x) d\mu(x_0),$$

so substituting this in and applying Jensen's inequality to bring the absolute value inside, we have

$$\int |w(x_j) - f(x_j)| d\mu(x_j) \leq \iint h(x_0)|p^j(x_0, x_j) - f(x_j)| d\mu(x_j) d\mu(x_0).$$

Integrating over j first and using Doeblin's theorem shows that this is bounded by $\int h(x_0) 2(1-\varepsilon)^j d\mu(x_0) = 2(1-\varepsilon)^j$, which completes the proof. □

So the point here is that if we have a function of a bunch of variables in two very separated time intervals and want to understand the correlations, we only care about the distance between the closest points, and that's very

related to exponential decay of correlations. More precisely, for any function F of $(X_0, \cdots, X_i)$ and any function G of $(X_j, \cdots, X_k)$ where $|F|, |G| \leq 1$, we have

$$\begin{aligned}&\mathbb{E}[FG] - \mathbb{E}[F]\mathbb{E}[G]\\ &= \int F(x_0, \cdots, x_i)G(x_j, \cdots, x_k)(g(x_0, \cdots, x_i, x_j, \cdots, x_k) - g_1(x_0, \cdots, x_i)g_2(x_j, \cdots, x_k))d\mu(x_0)\cdots d\mu(x_i)d\mu(x_j)\cdots d\mu(x_k).\end{aligned}$$

But then taking absolute values and then using our bounds from the previous result shows that

$$\mathrm{Cov}(F, G) = |\mathbb{E}[FG] - \mathbb{E}[F]\mathbb{E}[G]| \leq 4(1-\varepsilon)^{j-i}.$$

So we'll see next time how to use gauge fixing to reduce a two-dimensional system to such a system of Markov chains and prove our desired results for $d = 2$.

5 October 6, 2025

Last time, we were discussing mass gap in 2D lattice gauge theories. We'll prove the following result today:

Theorem 26

Consider lattice gauge theories with free boundary conditions on $\Lambda_N \uparrow \mathbb{Z}^2$, and assume that G is a compact Lie subgroup of $U(n)$ for some n. Then this sequence of theories has a mass gap at any $\beta \geq 0$.

In other words, we get exponential decay of correlations with constant not depending on N.

Proof. It suffices to consider the gauge-fixed theory, in which the matrices attached to all horizontal edges are the identity I – horizontal edges contain no loops, so this is valid. (We can do this same thing in any dimension, but it will not help us as much.) For any plaquette p of vertices $(x, y), (x+1, y), (x+1, y+1), (x, y+1)$, we have

$$\begin{aligned}U_p &= U((x, y), (x+1, y))U((x+1, y), (x+1, y+1))U((x, y+1), U(x+1, y+1))^{-1}U((x, y), (x, y+1))^{-1}\\ &= U((x+1, y), (x+1, y+1))U((x, y), (x, y+1))^{-1}\\ &= U(e_{x+1,y})U(e_{x,y}))^{-1},\end{aligned}$$

where we introduce the notation that $e_{x,y} = ((x, y), (x, y+1))$ is the vertical edge pointing upward from (x, y).

Now fix some N. For each $y \in \mathbb{Z}$, let $S(y)$ be the set of all $x \in \mathbb{Z}$ so that the plaquette we described above is within Λ_N. We can then write the action as

$$S_{\Lambda_N}(U) = \sum_{y\in\mathbb{Z}}\sum_{x\in S(y)} \mathrm{Re}(\mathrm{Tr}(I - U(e_{x+1,y})U(e_x, y)^{-1})).$$

Since Λ_N is finite, all of these are finite sums. Since $S(y) \subseteq \mathbb{Z}$, the set $S(y)$ is always a union of connected components (under the usual graph structure on $\mathbb{Z}$). Let $\mathcal{C}(y)$ be the set of such components, and for each $C = \{x, \cdots, x+k\} \in \mathcal{C}(y)$, define

$$S_C(U) = \sum_{j=0}^{k} \mathrm{Re}\left(\mathrm{Tr}\left(I - U(e_{x+j+1,y})U(e_{x+j,y})^{-1}\right)\right).$$

But now two components (and two components belonging to different ys) don't interact with each other (because

$x+k+1$ won't be in any other component), so the action is just a sum of terms of this form:

$$S_{\Lambda_N} = \sum_{y\in\mathbb{Z}} \sum_{C\in\mathcal{C}(y)} S_C(U).$$

So now let U be a random configuration drawn from the lattice gauge theory. For any $C = \{x, x+1, \cdots, x+k\}$, we then define the collection of random variables

$$U(C) = \{U(e_{x,y}), \cdots, U(e_{x+k+1,y})\}.$$

Assume for simplicity that every edge in Λ_N is in at least one plaquette (otherwise it doesn't contribute to the action so U on that edge is just a Haar matrix). The previous decomposition then says that the $U(C)$s are independent of each other, and the density of $U(C)$ is proportional to

$$\exp\left(-\beta \sum_{j=0}^{k} \mathrm{Re}\left(\mathrm{Tr}\left(I - U(e_{x+j+1,y})U(e_{x+j,y})^{-1}\right)\right)\right).$$

To simplify notation, let $U_1, U_2, \cdots, U_{k+1}$ be random elements of G with joint density (relative to Haar measure) proportional to $\exp\left(-\beta\sum_{j=0}^{k}\mathrm{Re}\left(\mathrm{Tr}\left(I - U_{j+1}U_j^{-1}\right)\right)\right)$. **We claim that** this is a Markov chain with kernel

$$P(V,W) = C^{-1}\exp\left(-\beta\mathrm{Re}\left(\mathrm{Tr}\left(I - WV^{-1}\right)\right)\right),$$

where C is a normalizing constant **that doesn't depend on** V, since by invariance of Haar measure μ we have

$$C = \int_G \exp\left(-\beta\mathrm{Re}\left(\mathrm{Tr}\left(I - WV^{-1}\right)\right)\right) d\mu(W) = \int_G \exp\left(-\beta\mathrm{Re}\left(\mathrm{Tr}\left(I - W\right)\right)\right) d\mu(W).$$

So we have a random walk on a group, and each time we're multiplying by a random matrix chosen by this specified density. The fact that C doesn't depend on V thus means that if U_1 is Haar distributed, and $U_2, U_3, \cdots$ are generated from this Markov kernel, then we indeed have the correct specified joint density for $(U_1, U_2, \cdots, U_{k+1})$.

But now we are almost done: $P(V,W)$ is a strictly positive continuous function on $G\times G$, and G is compact. Thus there is some $\varepsilon > 0$ (depending only on G and β) such that $P(V,W) \geq \varepsilon$ for all V, W. This means that Doeblin's condition is satisfied – we can check that μ is an invariant measure for this Markov chain, and thus we have exponential decay of correlations. (We can complete the proof on our own.) □

Remark 27. *We can get some bounds from this proof, but they'll be exponentially bad in β even if G is of a fixed size. Doeblin usually doesn't give sharp bounds, so we may need something else if we wanted the "correct constants" for exponential decay.*

Next, the area law also holds, but only under an additional assumption on G:

Theorem 28

Suppose, in addition to the assumptions of Theorem 26 (in particular, $d = 2$), that there is no nonzero $x \in \mathbb{C}^n$ fixed by all elements of G. Then the sequence of theories satisfies the area law.

(We should also be able to prove that if this additional condition is not satisfied, then Wilson's area law cannot hold, so this is necessary.) Groups like $U(n)$ or $SO(3)$ satisfy this "irreducibility-type condition," which is basically saying that the standard representation has no *one-dimensional* invariant subspace.

Lemma 29

Let U be Haar distributed on G, and let $E(U)$ be the matrix of expected values of the entries of U. If G fixes no nonzero elements of $\mathbb{C}^n$, then $E(U)$ is identically zero.

(In fact this is an "if and only if" statement: if G fixes a nonzero vector, then the expected value of $Ux = x$ is $E(U)$ times x, so $E(U)$ can't be zero.)

Proof. Let $||M||$ denote the ℓ^2 operator norm of an $n \times n$ matrix; that is,

$$||M|| = \sup\{||Mx|| : x \in \mathbb{C}^n : ||x|| = 1\} = \sup\{||Mx|| : x \in \mathbb{C}^n : ||x|| \leq 1\}.$$

This is indeed a norm and it satisfies the submultiplicativity condition $||MN|| \leq ||M|| \cdot ||N||$. First notice that $||E(U)|| \leq 1$, since if $U_1, U_2, \cdots$ are iid Haar and $x \in \mathbb{C}^n$ of norm at most 1 is fixed, then the law of large numbers says that $\frac{1}{n}\sum_{j=1}^n U_j x$ converges to $E(U)x$ almost surely as $n \to \infty$, and therefore

$$||E(U)x|| \overset{a.s.}{=} \lim_{n\to\infty} \left|\left| \frac{1}{n}\sum_{j=1}^n U_j x \right|\right| \leq \frac{1}{n}\sum_{j=1}^n ||U_j x|| \leq 1,$$

and taking a supremum over all x yields the result. We now further claim that $||E(U)|| < 1$; suppose not so that $||E(U)|| = 1$. Then there exists some $x \in \mathbb{C}^n$ with $||x|| = 1$ and $||E(U)x|| = 1$ (**by compactness** the supremum is achieved). If we now let U_1, U_2 be iid Haar, we have for our chosen x that

$$\begin{aligned}\mathbb{E}[||U_1 x - U_2 x||^2] &= \mathbb{E}\left[||U_1 x||^2 + ||U_2 x||^2 - 2\langle U_1 x, U_2 x\rangle\right] \\ &= 2 - 2\mathbb{E}[\langle U_1 x, U_2 x\rangle] \\ &= 2 - 2\langle E(U_1)x, E(U_2)x\rangle\end{aligned}$$

because all terms in the inner product depend only linearly on the entries of each individual matrix, and this last quantity is then $2 - 2||E(U)x||^2 = 0$. Thus we must have $U_1 x = U_2 x$ with probability 1, so $U_2^{-1}U_1 x = x$ almost surely – since $U_2^{-1}U_1$ is Haar distributed, continuity implies that $Vx = x$ for all $V \in G$. (This just requires showing that every point is in the support of the Haar measure, which is true because the Haar measure is a volume form.)

But now we are almost done: again letting U_1, U_2 be iid from the Haar measure, $U = U_1 U_2$ is also Haar, so $||E(U)|| = ||E(U_1)E(U_2)|| \leq ||E(U)||^2$, which can only occur if $||E(U)|| = 0$, hence $E(U) = 0$ as desired. □

Proof of Theorem 28. Again consider the gauge-fixed theory, and now consider a rectangle with horizontal width T and vertical height R. Take $R \leq T$ without loss of generality, and assume that our rectangle has corners $(0,0), (T,0), (T,R), (0,R)$. We again denote by $U_{j,k}$ the matrix on the edge from (j,k) to $(j,k+1)$; the Wilson loop we care about then satisfies

$$W_\ell = \mathrm{Tr}\left[U_{T,0}U_{T,1}\cdots U_{T,R-1}U_{0,R-1}^* U_{0,R_2}^* \cdots U_{0,0}^*\right]$$

(the inverse is just the adjoint for a unitary matrix, but thinking about it as the adjoint helps us out if we want to write things out in terms of **matrix entries**). If we now let $U_{j,k}^{a,b}$ denote the (a,b)th entry of $U_{j,k}$, we can write out the trace as a sum

$$W_\ell = \sum_{1\leq a_0, a_1, \cdots, a_{2R-1} \leq n} U_{T,0}^{a_0,a_1} U_{T,1}^{a_1,a_2} \cdots U_{T,R-1}^{a_{R-1},a_R} \overline{U_{0,R-1}^{a_{R+1},a_R}}\, \overline{U_{0,R-2}^{a_{R+2},a_{R+1}}} \cdots \overline{U_{0,0}^{a_0,a_{2R-1}}}.$$

Taking expectations of both sides, notice that the different horizontal stripes of vertical edges are independent, and thus this expectation actually factors as a product. Each horizontal stripe only affects two of the terms in our product, and so

$$\mathbb{E}[W_\ell] = \sum_{1\le a_0,\cdots,a_{2R-1}\le n} \prod_{j=0}^{R-1} \mathbb{E}\left[U_{T,j}^{a_j,a_{j+1}} \overline{U_{0,j}^{a_{2R-j},a_{2R-j-1}}}\right].$$

But now we can apply exponential decay of the Markov chains within each strip: we can write

$$\mathbb{E}\left[U_{T,j}^{a_j,a_{j+1}} \overline{U_{0,j}^{a_{2R-j},a_{2R-j-1}}}\right] = \mathbb{E}\left[U_{T,j}^{a_j,a_{j+1}} \overline{U_{0,j}^{a_{2R-j},a_{2R-j-1}}}\right] - \mathbb{E}\left[U_{T,j}^{a_j,a_{j+1}}\right]\mathbb{E}\left[\overline{U_{0,j}^{a_{2R-j},a_{2R-j-1}}}\right]$$

by our previous lemma, since all matrix entries have expectation zero. And this right-hand side is bounded by $C_1 e^{-C_2 T}$ since the two matrices are separated by a distance T, where C_1, C_2 depend only on G, β. We take a product of R of these things and we have a big sum, so we find that

$$\mathbb{E}[W_\ell] \le n^{2R}(C_1 e^{-C_2 T})^R = \exp\left(R(\log C_1 + 2\log n) - C_2 RT\right).$$

So now if $T \ge C_3 = \frac{2(\log C_1 + 2\log n)}{C_2}$, then the first term inside the exponential is at most $\frac{1}{2}C_2 RT$ and thus we have $\mathbb{E}[W_\ell] \le \exp\left(-\frac{1}{2}C_2 RT\right)$ for sufficiently large T. We thus choose $C_0 \ge 1$ so large that $|\langle W_\ell\rangle| \le C_0 \exp\left(-\frac{1}{2}C_2 RT\right)$ even for $T < C_3$, since under the condition $R \le T$ there are only finitely many choices of R and T (which are the only things that our bounds above depend on) and thus we can always find a C_0, as desired. □

The proof in higher dimensions to get from mass gap to area law is somewhat similar, expanding out the trace and using exponential decay. We'll see that later on in the course – indeed, we'll now move into general dimension in the strong coupling (β small) regime. We'll then specialize to $d = 3, 4$ and prove some results there.

6 October 8, 2025

Today, we'll discuss mass gap in the strong coupling regime. We'll first discuss one potential point of confusion: recall that in the continuum, the Yang-Mills action is given by

$$S(A) = -\frac{1}{2g_0^2}\int_{\mathbb{R}^4} \mathrm{Tr}\left(\sum_{0\le i<j\le 3} \eta_i\eta_j(\partial_i A_j(x) - \partial_j A_i(x) + [A_i(X), A_j(x)])^2\right) dx, \quad \eta_0 = 1, \quad \eta_1 = \eta_2 = \eta_3 = -1.$$

This action is often reparameterized so that the fields $A_j \mapsto g_0 A_j$ are rescaled, so that ew now have

$$S(A) = -\frac{1}{2}\int \mathrm{Tr}\left(\sum_{0\le i<j\le 3} \eta_i\eta_j(\partial_i A_j - \partial_j A_i + g_0[A_i, A_j])^2\right) dx.$$

Thus as $g_0 \to 0$, the fields decouple because we approach the electromagnetic theory and then we can do an appropriate gauge-fixing. (That's why g_0 is called the coupling constant.) For g_0 small we thus have "weak coupling," and for g_0 large we say that we have "strong coupling." In lattice gauge theories, β **plays the role of** $\frac{1}{g_0^2}$. Thus, β small indeed corresponds to strong coupling and β large to weak coupling.

Fact 30

We know that small β should make the matrices at different edges more independent, so the notation may be counterintuitive to probabilists. But this is where it all comes from – it's not about the individual sites being coupled, but rather the d fields (in d dimensions) being coupled to each other.

We'll discuss two results:

Theorem 31 (Osterwalder-Seiler 1978)
Consider lattice gauge theory on $\Lambda_N \uparrow \mathbb{Z}^d$ for some $d \geq 2$, and suppose the gauge group G is a subgroup of $U(n)$. Then there exists $\beta_0 > 0$, depending only G on d, such that for all $\beta < \beta_0$, the sequence of theories has mass gap.

There are different approaches for proving this – this result was originally proved using cluster expansions, but we'll show a different, more probabilistic approach.

Definition 32
Let μ_1, μ_2 be two probability measures on $(\Omega, \mathcal{F})$. The **total variation distance** between μ_1 and μ_2 is given by

$$d_{\text{TV}}(\mu_1, \mu_2) = \sup_{A \in \mathcal{F}} |\mu_1(A) - \mu_2(A)|.$$

This quantity has lots of alternative characterizations as well: for example, we have

$$d_{\text{TV}}(\mu_1, \mu_2) = \inf_{X \sim \mu_1, Y \sim \mu_2} \mathbb{P}(X \neq Y),$$

where we're allowed to couple X and Y in any arbitrary way as long as their marginals are specified. And for a third description, suppose that μ is a measure on $(\Omega, \mathcal{F})$ such that μ_1, μ_2 have densities ρ_1, ρ_2 with respect to μ. (We can always find such a μ, for example by setting $\mu = \mu_1 + \mu_2$.) Then

$$d_{\text{TV}}(\mu_1, \mu_2) = \frac{1}{2} \int_\Omega |\rho_1(x) - \rho_2(x)| d\mu(x).$$

The proofs that these are all the same are a bit tricky but not too difficult.

Example 33
Let μ, μ' be two probability measures on a product space Ω^n. For $1 \leq i \leq n$, let $\mu_i(\cdot|(x_j)_{j \neq i})$ denote the conditional law of X_i given $X_j = x_j$ for all other j, where $X = (X_1, \cdots, X_n)$ is distributed according to μ. Define $\mu_i'(\cdot|(x_j)_{j \neq i})$ similarly. Assume that for all $x, y \in \Omega^n$ and all $1 \leq i \leq n$, we have

$$d_{\text{TV}}\left(\mu_i(\cdot|(x_j)_{j \neq i}), \mu_i'(\cdot|(y_j)_{j \neq i})\right) \leq \sum_{j=1}^n \alpha_{ij} 1\{x_j \neq y_j\} + h_i$$

for some nonnegative constants α_{ij}, h_i (assume $\alpha_{ii} = 0$).

The idea is that if we have Markov random fields, the conditional distributions should only depend on a few neighboring coordinates, and this condition says that if the neighbors are the same, then the conditional distributions are the same except for some coordinates (think of h_i as being associated to the boundary values).

If $s = \max_{1 \leq i \leq n} \sum_{j=1}^n \alpha_{ij} < 1$ (this is often called the **Dobrushin condition**), then the matrix $Q = (\alpha_{ij})_{i,j=1}^n$ is substochastic, and in fact $\frac{1}{s}Q$ is substochastic as well. Thus there exists some Markov transition matrix P on $\{1, 2, \cdots, n\}$ such that $Q \leq sP$ elementwise; consider a random walk with those transition probabilities. Let τ_{ij} be the first hitting time at j when the walk starts from i.

Theorem 34

Take all of the notation above. Suppose $Z \sim \mu$ and $Z \sim \mu'$, and let $A \subseteq \{1, \cdots, n\}$. Let ν and ν' be the laws of $(Z_i)_{i \in A}$ and $(Z'_i)_{i \in A}$. Then

$$d_{\text{TV}}(\nu, \nu') \leq \frac{1}{1-s} \sum_{i \in A} \sum_{j=1}^{n} \mathbb{E}\left[s^{\tau_{ij}}\right] h_j.$$

The idea is that if our system is a field on a grid, and the h_js are nonzero on the boundary, then in particular τ_{ij} has to be at least the distance from a point i to the boundary. So all terms here will be exponentially small in that distance, and this lets us compare Gibbs distributions under different boundary conditions.

Proof. This is a coupling argument. Construct two Markov chains $X_0, X_1, \cdots$ and $X'_0, X'_1, \cdots$ as follows (remember that our state space is some product space Ω^n). First generate X_0, X'_0 independently with $X_0 \sim \mu$ and $X'_0 \sim \mu'$. Then at step k, given X_k and X'_k, we generate X_{k+1} and X'_{k+1} by "refreshing a coordinate in the best possible coupling" as follows. Pick a coordinate $i \in \{1, \cdots, n\}$ uniformly at random, and let $\gamma = \mu(\cdot|(X_{k,j})_{j \neq i})$ and $\gamma' = \mu'(\cdot|(X_{k,j})_{j \neq i})$. Then there exists a joint law θ with marginals γ, γ' (meaning $(X, Y) \sim \theta$, $X \sim \gamma$ and $Y \sim \gamma'$), coupled specifically so that $\mathbb{P}(X \neq Y) = d_{\text{TV}}(\gamma, \gamma')$. Generate $(X_{k+1,i}, X'_{k+1,i})$ from θ. Finally, keep all other coordinates the same, meaning that $X_{k+1,j} = X_{k,j}$ and $X'_{k+1,j} = X'_{k,j}$ for all $j \neq i$.

Since we're just regenerating from μ, μ' at each step from the conditional distribution and we start X_0, X'_0 according to those distributions, we have $X_k \sim \mu$ and $X'_k \sim \mu'$ for all k. Define the vector $\ell_k = (\ell_{k,1}, \cdots, \ell_{k,n})$ where $\ell_{k,i} = \mathbb{P}(X_{k,i} \neq X'_{k,i})$; our goal is to bound this vector. We know that

$$\mathbb{P}(X_{k+1,i} \neq X'_{k+1,i} | X_k, X'_k) = \left(1 - \frac{1}{n}\right) 1\{X_{k,i} \neq X'_{k,i}\} + \frac{1}{n} d_{\text{TV}}(\gamma, \gamma'),$$

since with probability $(1 - \frac{1}{n})$ the coordinate doesn't change, and otherwise we generate from the best possible coupling. Recall from the setup of Example 33 that we thus have the bound

$$\mathbb{P}(X_{k+1,i} \neq X'_{k+1,i} | X_k, X'_k) \leq \left(1 - \frac{1}{n}\right) 1\{X_{k,i} \neq X'_{k,i}\} + \frac{1}{n} \left(\sum_{j=1}^{n} \alpha_{ij} 1\{X_{k,j} \neq X'_{k,j}\} + h_i \right),$$

so taking expectations on both sides yields

$$\ell_{k+1,i} \leq \left(1 - \frac{1}{n}\right) \ell_{k,i} + \frac{1}{n} \sum_{j=1}^{n} \alpha_{ij} \ell_{k,j} + \frac{1}{n} h_i.$$

If we now define $\ell_{(i)} = \limsup_{k \to \infty} \ell_{k,i}$, we find that

$$\ell_{(i)} \leq \left(1 - \frac{1}{n}\right) \ell_{(i)} + \frac{1}{n} \sum_{j=1}^{n} \alpha_{ij} \ell_{(j)} + \frac{1}{n} h_i \implies \ell_{(i)} \leq \sum_{j=1}^{n} \alpha_{ij} \ell_{(j)} + h_i.$$

Recalling that Q was our matrix of α_{ij}s and $Q \leq sP$ for some stochastic matrix P and s the maximum row sum, we thus find that the vector ℓ of limsups satisfies $\ell \leq Q\ell + h$ coordinate-wise, so that $\boxed{(I - Q)\ell \leq h}$. Using the matrix norm

$$||M|| = \max_{1 \leq i \leq n} \sum_{j=1}^{n} |M_{ij}|$$

(we can check that indeed $||MN|| \leq ||M|| \cdot ||N||$), we have that $\sum_{k=0}^{\infty} Q^k$ converges because $||Q|| < 1$, it is the inverse of $I - Q$, and it has all nonnegative entries. Thus multiplying both sides of the boxed inequality by this sum maintains

the inequality, and thus we have

$$\ell \le \sum_{k=0}^{\infty} Q^k h \le \sum_{k=0}^{\infty} s^k P^k h$$

again by nonnegativity. This implies that

$$\ell_{(i)} \le \sum_{k=0}^{\infty} s^k \sum_{j=1}^{n} p_{ij}^{(k)} h_j,$$

where $p_{ij}^{(k)}$ is the (i,j)th entry of P^k. Exchanging order of summation yields

$$\ell_{(i)} \le \sum_{j=1}^{n} \left(\sum_{k=0}^{\infty} s^k p_{ij}^{(k)} \right) h_j.$$

But now if we let $Z_0, Z_1, \cdots$ be a Markov chain with transition matrix P, $p_{ij}^{(k)}$ is the probability that $Z_k = j$ given $Z_0 = i$. So

$$\begin{aligned} \sum_{k=0}^{\infty} s^k p_{ij}^{(k)} &= \sum_{k=0}^{\infty} s^k \mathbb{P}(Z_k = j | Z_0 = i) \\ &= \mathbb{E}\left[\sum_{k=0}^{\infty} s^k 1\{Z_k = j\} \middle| Z_0 = i \right] \\ &= \mathbb{E}\left[\sum_{k=\tau_j}^{\infty} s^k \middle| Z_0 = i \right] \end{aligned}$$

where τ_j is the first hitting time of j. Furthermore, this last quantity is just $\frac{1}{1-s}\mathbb{E}[s^{\tau_{ij}}]$. Thus we've shown that

$$\ell_{(i)} \le \frac{1}{1-s} \sum_{j=1}^{n} \mathbb{E}[s^{\tau_{ij}}] h_j.$$

To complete the proof, now we have for our particular coupling and any value of k that (remembering ν, ν' are the laws on A)

$$d_{\mathrm{TV}}(\nu, \nu') \le \mathbb{P}((X_{k,i})_{i \in A} \ne (X'_{k,i})_{i \in A} \le \sum_{i \in A} \mathbb{P}(X_{k,i} \ne X'_{k,i}) = \sum_{i \in A} \ell_{k,i},$$

so then taking limsup over all k yields the desired result. □

We can now apply this to our lattice gauge theories:

Start of proof of Theorem 31. Consider lattice gauge theory on Λ_N, and let f_1, f_2 be two bounded gauge-invariant observables supported on the sets $S_1, S_2 \subseteq \Lambda_N$. (Actually f_1, f_2 don't need to be gauge-invariant, so this is maybe an indication that this cannot be the "physically relevant" result.) Let $n = \mathrm{dist}(S_1, S_2)$ be the graph distance between the two sets. **Assume S_1 and S_2 have bounded sizes and we're taking n very large.** Then we can put a "box" around S_1 of radius $n/2$ (more precisely, the union of points at most $n/2$ away from points in S_1); for sufficiently large n we have S_2 outside this box. We then have

$$\mathbb{E}[f_1 f_2] = \mathbb{E}\left[\mathbb{E}[f_1 | \text{everything outside the box}] f_2\right],$$

since f_2 is dependent only on things outside the box. The point now is that for two different boundary conditions on the box, **if** we have

$$|\mathbb{E}[f_1 | \text{boundary condition 1}] - \mathbb{E}[f_1 | \text{boundary condition 2}]| \le C_1 e^{-C_2 n},$$

then because $\mathbb{E}[f_1]$ itself is some weighted average of $\mathbb{E}[f_1|\text{boundary condition}]$s, **then** we also have

$$|\mathbb{E}[f_1|\text{boundary condition}] - \mathbb{E}[f_1]| \leq C_1 e^{-C_2 n},$$

and thus plugging this back into the covariance expression $\mathbb{E}[f_1 f_2] - \mathbb{E}[f_1]\mathbb{E}[f_2]$ shows that it must be exponentially small. Thus we just want to prove this condition on conditional expectations.

Thus what we have is our arbitrarily-shaped region which we can call A, and we can let μ and μ' be the conditional laws of the lattice gauge theory in A under the two different boundary conditions – let $U \sim \mu$ and $U' \sim \mu'$. (Here, boundary conditions are specifications on the "boundary edges of A" which are part of plaquettes not in A.) For any non-boundary edge $e \in A$ where none of its neighboring edges (that is, edges with which e shares plaquettes) are boundary edges, the conditional density of U_e given all other matrices has density

$$\exp\left(-\beta \sum_{p \ni e} \text{Re}(\text{Tr}(I - U_p))\right);$$

that is, it retains the form of the original density because none of the relevant edges are being fixed by the boundary. The same holds for U' as well. Thus if U, U' have the same neighboring configuration, the distribution is the same, and if they have something different the total variation distance will still be small **if β is small enough**. And the h_is will come in for the edges where a neighboring edge is on the boundary, which will allow us to use our proved bound on total variation. We'll do this next time! □

7 October 13, 2025

Today, we'll first prove a better version of what we showed last time:

Definition 35

Let $(\Omega, \mathcal{F}, \mu)$ be a probability space, and let ν be a probability measure on the product space $(\Omega^n, \mathcal{F}^n, \mu^n)$ such that ν has a density ρ with respect to μ^n. Let $X \sim \nu$. Then the **marginal density** $(X_j)_{j \neq i}$ is defined by integrating out the ith coordinate:

$$\rho_{-i}((x_j)_{j \neq i}) = \int \rho(x_1, \cdots, x_n) d\mu(x_i).$$

The **conditional density** of X_i given $(X_j)_{j \neq i}$ is

$$\rho_i(x_i|(x_j)_{j \neq i} = \frac{\rho(x_1, \cdots, x_n)}{\rho_{-i}((x_j)_{j \neq i}}$$

(where we define this quantity arbitrarily if the denominator is zero).

Proposition 36

With the notation above, suppose the system admits a conditional dependency graph G on $\{1, \cdots, n\}$, meaning that $\rho(x_i|(x_j)_{j\neq i})$ is a function only of $(x_j)_{j\in N(i)\cup\{i\}}$ (that is, the conditional density depends only on i and its neighbors in the graph). Let Δ be the maximum vertex degree of G, and assume the Dobrushin condition

$$s = \sup_{\substack{1\leq i\leq n \\ x\in\Omega^n}} |\rho_i(x_i|(x_j)_{j\neq i}) - 1| < \frac{1}{\Delta}.$$

For disjoint sets $A, B \subseteq \{1, \cdots, n\}$, let $X_A = (X_i)_{i\in A}$ and $X_B = (X_i)_{i\in B}$ and define $r = d(A, B) = \min_{i\in A, j\in B} d(i, j)$, where d is the graph distance on G. Then for any measurable function $f : \Omega^A \to [-1, 1]$ and $g : \Omega^B \to [-1, 1]$,

$$|\mathbb{E}[f(X_A)g(X_B)] - \mathbb{E}[f(X_A)]\mathbb{E}[g(X_B)]| \leq \frac{2s^r\Delta^r|A| \cdot |B|}{1 - s\Delta}.$$

In particular, this quantity will be exponentially decaying in the distance, and indeed we see that the conditional density converges to 1 as $\beta \to 0$ so this will be useful in the high-temperature regime.

Proof. Let $D = \{i \notin B : d(i, B) = 1\}$ be the set of indices that are just outside of B. **We claim that** for all $x, x' \in \Omega^B$, we have

$$\mathbb{E}[f(X_A)|X_B = x] - \mathbb{E}[f(X_A)|X_B = x'] \leq \frac{2s^r\Delta^{r-1}}{1 - s\Delta}|A||D|.$$

Then $|D| \leq \Delta|B|$ because each element of B can have at most Δ neighbors, so in particular this is at most the right-hand side of our theorem. This then proves the desired bound by a similar argument to our proof from last lecture, since first note that we can write the difference to the unconditional expectation as

$$\mathbb{E}[f(X_A)|X_B = x] - \mathbb{E}[f(X_A)] = \int \left(\mathbb{E}\left[f(X_A)|X_B = x\right] - \mathbb{E}[f(X_A)|X_B = x']\right) d\tau_B(x')$$

and so the left-hand side is also bounded by $\frac{2s^{r-1}\Delta^r|A||B|}{1-s\Delta}$. Then

$$\mathbb{E}[f(X_A)g(X_B)] - \mathbb{E}[f(X_A)]\mathbb{E}[g(X_B)] = \mathbb{E}\left[\left(\mathbb{E}[f(X_A)|X_B] - \mathbb{E}[f(X_A)]\right) g(X_B)\right]$$

and now $|g(X_B)|$ is bounded by 1 so the result follows.

So now **to prove the claim**, we'll do a similar coupling argment as last time. Fix some $x, x' \in \Omega^B$. Let τ, τ' be the conditional distributions of X given $X_B = x$ and $X_B = x'$, respectively, and generate $X_0 \sim \tau, Y_0 \sim \tau'$ independently of each other. We can then define a Markov chain starting at (X_0, Y_0) which evolves as follows: from (X_k, Y_k), we choose a coordinate I uniformly from $\{1, \cdots, n\} \setminus B$, and then generate $(U, V) \subset \Omega^2$ so that $U \sim \rho_I(\cdot|(X_{k,j})_{j\neq I}$ and $V \sim \rho_I(\cdot|(Y_{k,j})_{j\neq I}$ are "resamplings of the Ith coordinate given the current value of X_k, Y_k, and where we couple so that $\mathbb{P}(U \neq V)$ is exactly the total variation density between the two conditional laws. We then let $X_{k+1,I} = U$ and $Y_{k+1,I} = V$ and keep all other coordinates the same.

By construction, each X_k is still distributed as τ and each Y_k is distributed as τ' (since we never update the coordinates in B). We wish to calculate the total variation distance between U and V. First of all note that if $X_{k,j} = Y_{k,j}$ for all $j \in N(I)$, then by assumption the conditional distributions U, V are exactly identical, so the total variation distance is 0. And if not, then

$$\mathbb{P}(U \neq V) = \frac{1}{2}\int \left|\rho_I(z|(X_{k,j})_{j\neq I}) - \rho_I(z|(Y_{k,j})_{j\neq I})\right| d\mu(z),$$

but the quantity inside the absolute values is at most $2s$ by assumption and the triangle inequality, and thus the whole

probability is at most s. Thus

$$\mathbb{P}(U \neq V) \leq s \sum_{j \in N(i)} 1\{X_{k,j} \neq Y_{k,j}\}$$

by a crude union bound, and thus we get the inequality (first case if coordinate i is not selected, and second case otherwise)

$$\begin{aligned}\mathbb{P}(X_{k+1,i} \neq Y_{k+1,i} | X_k, Y_k) &= \left(1 - \frac{1}{n-|B|}\right) 1\{X_{k,i} \neq Y_{k,i}\} + \frac{1}{n-|B|} s \sum_{j \in N(i)} 1\left\{(X_{k,j})_{j \in N(i)} \neq (Y_{k,j})_{j \in N(i)}\right\} \\ &\leq \left(1 - \frac{1}{n-|B|}\right) 1\{X_{k,i} \neq Y_{k,i}\} + \frac{s}{n-|B|} \sum_{j \in N(i) \setminus B} 1\left\{X_{k,j} \neq Y_{k,j}\right\} + \frac{s}{m} 1\left\{(X_{k,j})_{j \in N(i) \cap B} \neq Y_{k,j})_{j \in N(i) \cap B}\right\}.\end{aligned}$$

But we can bound the last indicator by $h_i = 1\{i \in D\}$, where recall $D = \{i : d(i, B) = 1\}$. So then define the vector $\ell_k = (\ell_{k,i})_{i \notin B}$ for $\ell_{k,i} = \mathbb{P}(X_{k,i} \neq Y_{k,i})$; taking expectations on both sides,

$$\ell_{k+1,i} \leq \left(1 - \frac{1}{m}\right) \ell_{k,i} + \frac{s}{m} \sum_{j \in N(i) \setminus B} \ell_{k,j} + \frac{s}{m} h_i.$$

As before, letting $\ell = \limsup_{k \to \infty} \ell_k$ (so that $\ell^{(i)}$ denotes the ith coordinate of this limsup), we have

$$\ell \leq \left(1 - \frac{1}{m}\right) \ell + \frac{s}{m} Q\ell + \frac{s}{m} h$$

where Q is the adjacency matrix of G restricted to $\{1, \cdots, n\} \setminus B$. This rearranges to $(I - sQ)\ell \leq sh$, and now we claim that $(I - sQ)^{-1} = \sum_{k=0}^{\infty} (sQ)^k$. For this it's sufficient to show that the right-hand side is convergent (because then we get a telescoping sum for $\sum_{k=0}^{\infty} (sQ)^k (I - sQ)$), remembering that convergence is the same as absolute convergence because our entries are nonnegative. But because $Q^k = (q_{ij}^{(k)})_{1 \leq i,j \leq n}$ is a power of the adjacency matrix, $q_{ij}^{(k)}$ is the number of k-step paths from i to j that avoid B, which is in particular at most Δ^k. Thus $s^k q_{ij}^{(k)} \leq (s\Delta)^k$ will decay exponentially and this proves the claim; this means that (again using nonnegativity of entries so inequality is preserved)

$$\ell \leq (I - sQ)^{-1} sh = \sum_{k=0}^{\infty} s^{k+1} Q^k h,$$

so for any coordinate $i \in A$ we have

$$\begin{aligned}\ell^{(i)} &= \sum_{k=0}^{\infty} \sum_{j \notin B} s^{k+1} q_{ij}^{(k)} h_j \\ &= \sum_{j \notin B} \sum_{k=0}^{\infty} s^{k+1} q_{ij}^{(k)} 1\{j \in D\} \\ &= \sum_{j \in D} \sum_{k=0}^{\infty} s^{k+1} q_{ij}^{(k)}.\end{aligned}$$

But for any $j \in D$ we know that $q_{ij}^{(k)} = 0$ for $k < r - 1$ (since we cannot reach the set in that few steps) and $q_{ij}^{(k)} \leq \Delta^k$

otherwise. Thus

$$\begin{aligned}\ell^{(i)} &\leq \sum_{j\in D}\sum_{k=r-1}^{\infty} s^{k+1}\Delta^k \\ &\leq |D|\sum_{k=r-1}^{\infty} s^{k+1}\Delta^{k+1} \\ &= \frac{s^r\Delta^r}{1-s\Delta}|B|.\end{aligned}$$

So to conclude the proof of our claim, we have

$$\begin{aligned}\left|\mathbb{E}[f(X_A)|X_B = x] - \mathbb{E}[f(X_A)|X_B = x']\right| &= \left|\mathbb{E}\left[f((X_{k,i})_{i\in A}) - f((Y_{k,i})_{i\in A}\right]\right| \\ &\leq 2\mathbb{P}((X_{k,i})_{i\in A} \neq (Y_{k,i})_{i\in A}) \\ &\leq 2\sum_{i\in A}\mathbb{P}(X_{k,i} \neq Y_{k,i}) \\ &= 2\sum_{i\in A}\ell_{k,i},\end{aligned}$$

and since k was arbitrary we take the limsup as $k \to \infty$ and apply the bound to get our result. □

So we now see how we can prove mass gap (Theorem 31) from this:

Proof of Theorem 31. For edges $e, f \in E(\Lambda_N)$, say that e and f are neighbors if they belong to some common plaquette. It is easy to see that this yields a conditional dependency graph for the model (because the contributions to the density only depend on values on plaquettes), and the conditional density has the expression

$$\rho(U_e|(U_f)_{f\neq e}) = \frac{1}{Z(e)}\exp\left(-\beta\sum_{p\in\mathcal{P}(e)}\mathrm{Re}(\mathrm{Tr}(I - U_p))\right)$$

where $\mathcal{P}(e)$ is the set of plaquettes in Λ_N containing e, and where the normalizing constant is

$$Z(e) = \int_G \exp\left(-\beta\sum_{p\in\mathcal{P}(e)}\mathrm{Re}(\mathrm{Tr}(I - U_p))\right).$$

But the number of plaquettes in $\mathcal{P}(e)$ is uniformly bounded by something like $2d$, so for any $\varepsilon > 0$ we can choose some $\beta_0(d, G)$ so that for all $\beta < \beta_0$ the conditional density is close to 1; that is, we can make $\left|\exp\left(-\beta\sum_{p\in\mathcal{P}(e)}\mathrm{Re}(\mathrm{Tr}(I - U_p))\right) - 1\right| < \varepsilon$ (which also implies $|Z(e) - 1| < \varepsilon$; thus for ε small enough this ensures the Dobrushin condition holds, proving exponential decay of correlations as desired. □

Next, we'll prove the area law with a similar argument:

Theorem 37

Suppose that there is no $x \in \mathbb{C}^n \setminus \{0\}$ fixed by all elements of G. Then there is some $\beta \leq \beta_0(d, G)$ so that the area law holds for all $\beta \leq \beta_0(d, G)$.

Recall that in two dimensions, we proved this by taking a large T by R loop and gauge-fixed so that we only have horizontal slices. We then further decomposed each slice into adjacent blocks which were Markov chains, and we used exponential decay of correlations to bound terms from different sides of the rectangle. We don't get a Markov chain in higher dimensions anymore, so we use the following argument instead:

Proof. Let $e_1, \cdots, e_d$ be the standard basis vectors of $\mathbb{R}^d$, and let ℓ be the rectangular loop with vertices $0, Te_1, Re_2$, and $Te_1 + Re_2$. Define the oriented matrices along the lower edge

$$U_{j,0} = U((j, 0, 0, \cdots, 0), (j+1, 0, 0, \cdots, 0))$$

and along the upper edge

$$U_{j,R} = U((j, R, 0, \cdots, 0), (j+1, R, 0, \cdots, 0))$$

for $0 \le j \le T-1$, except the avoid notational overlap let $U((0,0,0,\cdots,0),(1,0,0,\cdots,0))$ be written $\tilde{U}_{0,0}$. Also define the matrices along the left edge

$$U_{0,k} = U((0, k, 0, \cdots, 0), (0, k+1, 0, \cdots, 0))$$

and right edge

$$U_{T,k} = U((T, k, 0, \cdots, 0), (T, k+1, 0, \cdots, 0))$$

for $0 \le k \le R-1$. Then the Wilson loop we are interested in can be written in this notation as

$$W_\ell = \mathrm{Tr}\Big(\tilde{U}_{0,0} U_{1,0} \cdots U_{T-1,0} U_{T,0} U_{T,1} \cdots U_{T,R-1} U^*_{T-1,R} U^*_{T-2,R} \cdots U^*_{0,R} U^*_{0,R-1} \cdots U^*_{0,1} U^*_{0,0}\Big),$$

and this trace can be written out as a huge sum

$$\sum_{1 \le a_0, \cdots, a_{2T+2R-1} \le n} \tilde{U}^{a_0,a_1}_{0,0} U^{a_1 a_2}_{1,0} \cdots U^{a_{T-1},a_T}_{T-1,0} U^{a_T,a_{T+1}}_{T,0} \cdots$$

over products of matrix entries. But now for each $j \in \mathbb{Z}$ we can define E_j to be the set of edges in Λ_N from $(x_1, j, x_3, \cdots, x_d)$ to $(x_1, j+1, x_3, \cdots, x_d)$, and let E' be the union of all such edges. Conditional on $(U(e))_{e \in E(\Lambda_N) \setminus E'}$, the collections $(U(e))_{e \in E_j}$ are independent across different js, since they don't have any common plaquettes. Additionally, the conditional density of one such slab $(U(e))_{e \in E_j}$ given $(U(e))_{e \in E'}$ satisfies the Dobrushin condition with some deterministic s. So we have exponential decay within each slab **as long as β is large enough** (remember that in $d = 2$ we had it for all β because we always had Doeblin's condition). So the rest of the proof is exactly the same as before: taking the conditional expectation of W_ℓ given $(U(e))_{e \in E'}$ yields dependence in pairs, so that expectations factor and then we get the necessary exponential decay for each one. □

8 October 20, 2025

Today, we'll discuss confinement in 3D gauge theories. It's believed that confinement holds quite generally for three-dimensional lattice gauge theories, and the area law has been proved for 3D $U(1)$ theory with a different action called the Villain action by Göpfert and Mack in 1982. But to prove confinement, remember that we don't need something as strong as the area law – we just need something better than the perimeter law. That's what the following (new) result says:

Theorem 38

Let G be a compact Lie subgroup of $U(n)$ for some n, and assume that $zI \in G$ for all $z \in \mathbb{C}$ with $|z| = 1$ (meaning that it "contains $U(1)$") – for example, $U(n)$ satisfies this but not $SU(n)$ (though the following result should also be true for $SU(n)$). Let $\Lambda_N \uparrow \mathbb{Z}^3$ be an increasing sequence of lattices and consider the usual lattice gauge theory with gauge group G on Λ_N with inverse coupling β. Take any rectangular loop ℓ with side lengths $R \leq T$, and pick N large enough that $\ell \subseteq \Lambda_N$. Then

$$|\langle W_\ell \rangle| \leq C_1 e^{-C_2 T \log R}$$

for positive constants c_1, c_2 depending only on G and β.

We'll see in the proof where things go wrong if we don't have $U(1)$ as a subgroup – the actual mild necessary condition conjectured is "there is no nonzero vector in the kernel of the Lie algebra," but that hasn't been worked out yet.

To prove this, we'll begin with the following simple lemma about probability measures on $U(1)$:

Lemma 39

Take any $w \in \mathbb{C}$ and let μ be the probability measure on the unit circle $U(1)$ with density proportional to $\exp(\mathrm{Re}(z\overline{w}))$ with respect to Haar measure. Then the measure is spread out in the sentence that the "variance of the measure" is

$$\iint |z_1 - z_2|^2 d\mu(z_1) d\mu(z_2) \geq C \min\left\{1, \frac{1}{|w|}\right\}$$

for some universal constant C.

There are many possible proofs, but this particular one generalizes to other Lie groups (where maybe we care about the real part of the trace of the product of two matrices – then the result is less obvious because it becomes less clear what the projections look like).

Proof. For all $x \in \mathbb{R}$, we have the Taylor approximation $|e^{ix} - i - ix| \leq \frac{1}{2}x^2$, so that for all $z \in U(1)$ and $x \in \mathbb{R}$,

$$\begin{aligned} |z - ze^{ix}| &= |z(1 - e^{ix})| \\ &= |z(1 - e^{ix} + ix) - zix| \\ &\geq |zix| - |z(1 - e^{ix} + ix)| \\ &\geq |x| - \frac{x^2}{2} \end{aligned}$$

by the reverse triangle inequality whenever the former term is larger. So for $\varepsilon = \min\left\{1, \frac{1}{\sqrt{|w|}}\right\} \leq 1$, we have

$$|z - ze^{\pm i\varepsilon}| \geq \varepsilon - \frac{1}{2}\varepsilon^2 \geq \frac{1}{2}\varepsilon$$

for all $|\varepsilon| \leq 1$; that is, shifting by a small amount on the unit circle changes the distance by some noticeable amount. If we then define $\rho(z) = \exp(\mathrm{Re}(z\overline{w}))$ and define

$$Z = \int_{U(1)} \rho(z) d\sigma_0(z)$$

for σ_0 the normalized Haar measure (that is, the uniform distribution on the unit circle), then we have $\frac{dZ}{d\sigma_0(z)} = \frac{1}{Z}\rho$.

If we now define

$$\rho^+(z) = \rho(ze^{i\varepsilon}), \quad \rho^-(z) = \rho(ze^{-i\varepsilon}),$$

then by invariance of σ_0 we have that $\int_{U(1)} \rho^+(z)d\sigma_0(z) = \int_{U(1)} \rho^-(z)d\sigma_0(z) = Z$ is the same normalizing constant. Letting μ^+, μ^- be the shifted probability measures on $U(1)$ given by $\frac{d\mu^+}{d\sigma_0} = \frac{1}{Z}\rho^+$ and $\frac{d\mu^-}{d\sigma_0} = \frac{1}{Z}\rho^-$, we have

$$\mu^+(A) = \frac{1}{Z}\int_A \rho(ze^{i\varepsilon})d\sigma_0(z),$$

then the change of variables $z' = ze^{i\varepsilon}$ yields

$$\mu^+(A) = \frac{1}{Z}\int_{Ae^{i\varepsilon}} \phi(z')d\sigma_0(z') = \mu(Ae^{i\varepsilon})$$

(which is not a surprise – we're just rotating the density) and similarly $\mu^-(A) = \mu(Ae^{-i\varepsilon})$. Thus for all $z \in U(1)$ we have

$$\begin{aligned}
\sqrt{\rho^+(z)\rho^-(z)} &= \exp\left(\frac{1}{2}\mathrm{Re}(\overline{w}z(e^{i\varepsilon} + e^{-i\varepsilon})\right) \\
&= \rho(z)\exp\left(\frac{1}{2}\mathrm{Re}(\overline{w}z(e^{i\varepsilon} + e^{-i\varepsilon} - 2)\right) \\
&\geq \rho(z)\exp\left(\frac{1}{2}|\overline{w}z(e^{i\varepsilon} + e^{-i\varepsilon} - 2)|\right) \\
&\geq \rho(z)\exp\left(\frac{1}{2}|\overline{w}| \cdot |e^{i\varepsilon} + e^{-i\varepsilon} - 2|\right),
\end{aligned}$$

and now $|e^{i\varepsilon} + e^{-i\varepsilon} - 2| \leq |e^{i\varepsilon} - 1 - i\varepsilon| + |e^{-i\varepsilon} - 1 + i\varepsilon| \leq \varepsilon$ by the inequality we've proved. So actually we have

$$\sqrt{\rho^+(z)\rho^-(z)} \geq \rho(z)e^{-\frac{1}{2}\varepsilon^2|w|} \geq \rho(z)e^{-1/2}$$

by our choice of ε. So for all measurable sets $A \subseteq U(1)$, we have

$$\begin{aligned}
\sqrt{\mu^+(A)\mu^-(A)} &= \frac{1}{Z}\sqrt{\int_A \rho^+(z)d\sigma_0(z)\int_A \rho^-(z)d\sigma_0(z)} \\
&\geq \frac{1}{Z}\sqrt{\rho^+(z)\rho^-(z)}d\sigma_0(z) \\
&\geq \frac{e^{-1/2}}{Z}\int_A \rho(z)d\sigma_0(z) \\
&= e^{-1/2}\mu(A),
\end{aligned}$$

second line by Cauchy-Schwarz in the reverse direction. So then for any $z \in U(1)$, if we define the set $\{A = z' \in U(1) : |z - z'| < \frac{1}{4}\varepsilon\}$, then the above inequality shows that either $\mu^+(A)$ or $\mu^-(A)$ is at least $e^{-1/2}\mu(A)$ (since the geometric mean is at least this value); without loss of generality suppose it is $\mu^+(A)$. Therefore

$$\begin{aligned}
\mu(A \cap Ae^{i\varepsilon}) &\geq \mu(A) + \mu(Ae^{i\varepsilon}) - 1 \\
&= \mu(A) + \mu^+(A) - 1 \\
&\geq \mu(A)(1 + e^{-1/2}) - 1.
\end{aligned}$$

We claim now that actually $A \cap Ae^{i\varepsilon})$ must be empty; indeed if $z' \in A \cap Ae^{i\varepsilon}$, then $|z' - z| < \frac{1}{4}\varepsilon$ and $|z' - ze^{i\varepsilon}| < \frac{1}{4}\varepsilon$,

meaning that $|z - ze^{i\varepsilon}| < \frac{1}{2}\varepsilon$, which is impossible by our earlier bounds. Therefore

$$0 \geq \mu(A)(1 + e^{-1/2}) - 1 \implies \mu(A) \leq \frac{1}{1 + e^{-1/2}},$$

meaning we've gotten a bound on any $\frac{1}{4}\varepsilon$-neighborhood of any given point. And now we're almost done, since

$$\int |z - z'|^2 d\mu(z') \geq \frac{1}{16}\varepsilon^2 \mu(A^c)$$

by only integrating over the part where $|z - z'|$ is large, and this is at least $\frac{1}{16}\varepsilon^2 \left(1 - \frac{1}{1+e^{-1/2}}\right) = C\varepsilon^2$ for some universal constant C, as desired. □

This proof didn't have to use projections anywhere, and so it goes through basically verbatim for other Lie groups as well.

Corollary 40

Suppose ξ is a $U(1)$-valued random variable with density proportional to ρ (as above). Then

$$|\mathbb{E}[\xi]| \leq 1 - C \min\left\{1, \frac{1}{|w|}\right\}.$$

Proof. Let ξ_1, ξ_2 be iid copies of ξ. Then by a simple computation we have

$$|\xi_1 - \xi_2|^2 = 2 - 2\mathrm{Re}(\xi_1 \overline{\xi_2})$$

so taking expectations on both sides yields

$$\begin{aligned}\mathbb{E}\left[|\xi_1 - \xi_2|^2\right] &= 2 - 2\mathbb{E}\left[\mathrm{Re}(\xi_1\overline{\xi_2})\right] \\ &= 2 - 2\mathrm{Re}\left(\mathbb{E}[\xi]\overline{\mathbb{E}[\xi]}\right) \\ &= 2 - 2|\mathbb{E}[\xi]|^2.\end{aligned}$$

Therefore $|\mathbb{E}[\xi]|^2 \leq 1 - \frac{1}{2}\mathbb{E}[(|\xi_1 - \xi_2|)^2] \leq 1 - C\min\left(1, \frac{1}{|w|}\right)$, and take square root on both sides and use $\sqrt{1-a} \leq 1 - \frac{a}{2}$ to conclude. □

We can now use this to prove the following result, which is a generalization of the famous Mermin–Wagner theorem (using a different proof technique).

Theorem 41

Let Λ be a finite subset of $\mathbb{Z}^2$ (not $\mathbb{Z}^3$) and let E be the set of positively oriented nearest-neighbor edges with both endpoints in Λ. Let $(w_e)_{e\in E}$ be some collection of complex numbers, and consider the probability measure γ on $U(1)^\Lambda$ (that is, a complex number at each **vertex**) with density proportional to $\exp\left(\sum_{e\in(x,y)\in E} \mathrm{Re}(w_e \xi_x \overline{\xi_y})\right)$ with respect to product Haar measure on $\xi \in U(1)^\Lambda$. (This is called the XY model.) Then for all $x \in \Lambda$ and all $R > 0$ so that $y = x + (R, 0) \in \Lambda$, we have (for ϕ distributed as γ)

$$\mathbb{E}[\phi_x \overline{\phi_y}] \leq C_1 e^{-C_2 \log R},$$

where C_1, C_2 depend only on the maximum weight $\max_{e\in E} |w_e|$. That is, correlation decays polynomially in distance.

The original result had the same w at all sites, and the symmetry was important in that argument (for confinement we do really need arbitrary weights for the argument to go through). And some other proofs use tools like contour integrals or other probabilistic arguments, but we'll do our particular proof because it seems like it might generalize to nonabelian groups. Indeed, something like this should be true not just for $U(1)$; it should hold as long as we have **continuous** symmetry. (Note that this result is not true for something like the Ising model, where we do have a phase transition past where it doesn't hold.)

Also, note that these correlations do not subtract off the $\mathbb{E}[\phi_x]\mathbb{E}[\overline{\phi_y}]$ term like we've had in some past exponential decay results.

Start of proof. Fix x, R as above, and without loss of generality take $\Lambda = x + \{-N, \cdots, N\}^2$ for some large N. (Indeed, we can put $w_e = 0$ for all extra edges when we expand the lattice, so we just get our original system coupled with independent spins.) For all $0 \le k \le N$, define $y_k = x + (k, 0)$, and define (for ease of notation)

$$L_k = \{y \in \Lambda : \ell^\infty(y, x) = k\}$$

(so basically the boundary of the square centered at x which goes up to y_k). Define the map $\tau : U(1)^\Lambda \to U(1)^\Lambda$ between configurations as follows. At the center, we define $\tau(\xi)_x = \xi_x$ (keeping the value the same), and then at each y_k we define

$$\tau(\xi)_{y_k} = \xi_{y_k}\overline{\xi_{y_{k-1}}} \text{ for } 1 \le k \le N.$$

Then for all $1 \le k \le N$ and all $y \in L_k \setminus \{y_k\}$, we define

$$\tau(\xi)_y = \xi_y\overline{\xi_{y_k}}.$$

For any edge $e \in E$, either e connects two neighboring vertices in some L_k, or it connects a vertex in L_k to a vertex in L_{k-1} for some $1 \le k \le N$; call the edges in the former set E_k and the edges in the latter set $F_k \cup \{(y_{k-1}, y_k)\}$ (so we exclude the distinguished edge that we are most interested in). For any configuration we can then define (this is like what appears in the density but without the real part)

$$\begin{aligned} f(\xi) &= \sum_{e \in (y,y') \in E} w_e \xi_y \overline{\xi_{y'}} \\ &= \sum_{k=1}^N w_{(y_{k-1}, y_k)} \xi_{y_{k-1}} \overline{\xi_{y_k}} + \sum_{k=1}^N \sum_{e=(y,y') \in E_k} w_e \xi_y \overline{\xi_{y'}} + \sum_{k=1}^N \sum_{e=(y,y') \in F_k} w_e \xi_y \overline{\xi}_{y'}. \end{aligned}$$

We now consider a different configuration $\chi = \tau(\xi)$, and we will rewrite f in terms of χ instead of ξ. In the case $(y, y') \in F_k$, we have

$$\xi_y \overline{\xi_{y'}} = \xi_y \overline{\xi_{y_{k-1}}}\, \overline{\xi_{y'}} \xi_{y_k} \overline{\xi_{y_k}} \xi_{y_{k-1}} = \chi_y \overline{\chi_{y'}}\, \overline{\chi_{y_k}}$$

(everything is commutative here because it's complex numbers – what helps for more complicated groups in the all-identity case is that we can swap things around with traces), so that last term is easy to replace. Next if $(y, y') \in E_k$, we have to consider cases: if neither y, y' are y_k, then

$$\xi_y \overline{\xi_{y'}} = \xi_y \overline{\xi_{y_k}}\, \overline{\xi_{y'}} \xi_{y_k} = \chi_y \overline{\chi}_{y'}.$$

Besides that, we also have to consider the edges attached to y_k going up and down: if $y = y_k, y' = y_k + e_2$ then $\xi_y \overline{\xi}_{y'} = \xi_{y_k} \overline{\xi}_{y'} = \overline{\chi_{y'}}$. Similarly if $y' = y_k, y = y_k - e_2$, then $\xi_y \overline{\xi_{y'}} = \xi_y \overline{\xi}_{y_k} = \chi_y$.

Finally, the first term is easy because we just have $\xi_y \overline{\xi_{y'}} = \overline{\chi_{y_k}}$ by definition. So putting this all together, we find

that

$$f(\xi) = \sum_{k=1}^{N} w_{(y_{k-1},y_k)}\overline{\chi_{y_k}} + \sum_{k=1}^{N} \sum_{\substack{e=(y,y')\in E_k \\ y,y'\neq y_k}} w_e \chi_y \overline{\chi_{y'}} + \sum_{k=1}^{N} \sum_{e=(y,y')\in F_k} w_e \chi_y \overline{\chi_{y'}}\,\overline{\chi_{y_k}}$$
$$+ \sum_{k=1}^{N} w_{(y_k,y_k+e_2)}\overline{\chi_{y_k+e_2}} + \sum_{k=1}^{N} w_{(y_k-e_2,y_k)}\chi_{y_k+e_2},$$

meaning that if we define the right-hand function as $g(\chi)$, then $f = g \circ \tau$. So what we'll do is pick from the original density, then apply τ; the new configuration has the density on the right-hand side. But now if we consider χ at $y_0, y_1, \cdots, y_n$ and look at the conditional density given all other χs, then actually χ_{y_k} and χ_{y_j} never interact so all of the values at different states are conditionally independent. So then χ_{y_k} is a successive product of terms and thus we end up multiplying conditionally independent variables to calculate correlations; densities then come in the form from our corollary above, so the conditional expectation when multiplying k things together is bounded by something like $1 - \frac{c}{k}$. We'll discuss some details of this next time! □

9 October 22, 2025

Last time, we proved a result about an inhomogeneous $U(1)$-valued spin system on a lattice $x + \{-N, \cdots, N\}^2$, and we're trying to prove specifically that for $y = x + (R, 0)$, we have

$$\left|\mathbb{E}[\phi_x \overline{\phi_y}\right| \leq e^{-C/L \log(R+1)}$$

for some universal constant C and where $L = 1 + \max_{e\in E} |w_e|$. What we did so far is define $f(\xi) = \sum_{e=(x,y)\in E} w_E \xi_x \xi_y$ and then define a certain reparametrization τ, which lets us write $f(\xi)$ fully in terms of $\chi = \tau(\xi)$ (as we did at the end of last lecture). We cal this new expression g, so that

$$g(\chi) = \sum_{k=1}^{N} w_{(y_{k-1},y_k)}\overline{\chi_{y_k}} + \sum_{k=1}^{N} \sum_{\substack{e=(y,y')\in E_k \\ y,y'\neq y_k}} w_e \chi_y \overline{\chi_{y'}} + \sum_{k=1}^{N} \sum_{e=(y,y')\in F_k} w_e \chi_y \overline{\chi_{y'}}\,\overline{\chi_{y_k}}$$
$$+ \sum_{k=1}^{N} w_{(y_k,y_k+e_2)}\overline{\chi_{y_k+e_2}} + \sum_{k=1}^{N} w_{(y_k-e_2,y_k)}\chi_{y_k+e_2}.$$

We now claim that if ϕ is a random configuration with density proportional to $e^{\mathrm{Re}(f)}$, then the new spin configuration $\psi = \tau(\phi)$ has density proportional to $e^{\mathrm{Re}(g)}$. (In effect, the point is that Haar measure means the change of variable determinant constant will be 1.) To do that, we prove the following fact:

Lemma 42

The transformation τ is measure-preserving, meaning that

$$\int F(\tau(\xi)) \prod_{y\in\Lambda} d\sigma_0(\xi_y) = \int F(\xi) \prod_{y\in\Lambda} d\sigma_0(\xi_y)$$

for any F and for σ_0 normalized Haar measure on $U(1)$.

Proof. Recall that we defined $\tau(\xi)_{y_0} = \xi_{y_0}$, then $\tau(\xi)_{y_k} = \xi_{y_k}\overline{\xi_{y_{k-1}}}$, and finally (for all $y \in L_k \setminus \{y_k\}$) $\tau(\xi)_y = \xi_y \overline{\xi_{y_k}}$. So in the left-hand integral, we can replace $\tau(\xi)_y$ with ξ_y for all $y \notin \{y_0, \cdots, y_N\}$ by doing a single change of variable

and using invariance of σ_0. Once we do that, our integrand now looks like

$$F\left(\tau(\xi)_{y_0}, \cdots, \tau(\xi)_{y_N}, \tau(y)_{y \notin \{y_0, \cdots, y_N\}}\right) = F\left(\xi_{y_0}, \xi_{y_1}\overline{\xi_{y_0}}, \xi_{y_2}\overline{\xi_{y_1}}, \cdots, \xi_{y_N}\overline{\xi_{y_{N-1}}}, (\xi_y)_{y \notin \{y_0, \cdots, y_N\}}\right).$$

If we now fix everything else besides the y_ks, we can observe that ξ_{y_N} only appears in one place, so we can replace $\xi_{y_N}\overline{\xi_{y_{N-1}}}$ by just ξ_{y_N} by invariance. Then we do the same for $\xi_{y_{N-1}}$ (since that now only shows up in one place), then $\xi_{y_{N-2}}$, and so on, and that removes all of the extra factors and gets us the thing on the right-hand side, as desired. □

So now we can apply this lemma by noting that

$$\begin{aligned}\int F(\tau(\xi)) e^{\mathrm{Re}(f(\xi))} \prod_{y \in \Lambda} d\sigma_0(\xi_y) &= \int F(\tau(\xi)) e^{\mathrm{Re}(g(\tau(\xi))} \prod d\sigma_0(\xi_y) \\ &= \int F(\xi) e^{\mathrm{Re}(g(\xi))} \prod d\sigma_0(\xi_y)\end{aligned}$$

by applying our lemma, and now if we set $F = 1$ we see that the normalizing constants are the same: $\int e^{\mathrm{Re}(f)} = \int e^{\mathrm{Re}(g)}$. Thus the expectation $\mathbb{E}[f(\psi)] = \frac{\int F(\tau(\xi)) e^{\mathrm{Re}(f(\xi))}}{\int e^{\mathrm{Re}(f(\xi))}}$ is exactly $\frac{\int F(\xi) e^{\mathrm{Re}(g(\xi))}}{\int e^{\mathrm{Re}(g(\xi)}}$, proving our claim.

Returning to our main proof now, we now let $\phi \sim e^{\mathrm{Re}(f)}$ so that $\psi = \tau(\phi) \sim e^{\mathrm{Re}(g)}$. We then have

$$\psi_{y_0} = \phi_{y_0}, \quad \psi_{y_k} = \phi_{y_k}\overline{\phi_{y_{k-1}}} \text{ for } k \geq 1.$$

Thus for $x = y_0$ and $y = x + (R, 0) = y_R$, we have that (remember everything is still commutative here)

$$\psi_{y_1} \cdots \psi_{y_R} = phi_{y_1}\overline{\phi_{y_0}}\phi_{y_2}\overline{\phi_{y_1}} \cdots \phi_{y_R}\overline{\phi_{y_{R-1}}} = \overline{\phi_{y_0}}\phi_{y_R},$$

so that the quantity we're interested in involves a product of sequential ψs. Looking back at our expression for g, we can now write it as

$$g(\chi) = \sum_{k=1}^{N} w_{(y_{k-1}, y_k)}\overline{\chi_{y_k}} + \sum_{k=1}^{N} M_k \overline{\chi_{y_k}} + \mathcal{R}$$

for $\mathcal{R}$ a remainder term with no dependence on χ_{y_0} through χ_{y_N} and where

$$M_k = \sum_{e=(y,y') \in F_k} w_e \chi_y \overline{\chi_{y'}}.$$

also has no dependence on χ_{y_0} through χ_{y_N}. Thus we have **conditional** independence of the random variables $\psi_{y_0}, \cdots, \psi_{y_N}$ given all other ψ_ys, and under this conditioning $\psi_{y_0} \sim \sigma_0$ and each other ψ_{y_k} has density proportional to

$$\exp\left(\mathrm{Re}\left((w_{(y_{k-1}, y_k)} + M_k)\overline{\psi_{y_k}}\right)\right) = \exp\left(\mathrm{Re}\left(\overline{(w_{(y_{k-1}, y_k)} + M_k)}\psi_{y_k}\right)\right)$$

But because all ws are bounded by our constant L, we can crudely bound the absolute value of the conjugated term by Ck for some universal constant. Thus if $\mathbb{E}'$ denotes the conditional expectation, we have by Corollary 40 that

$$|\mathbb{E}'(\psi_{y_k})| \leq 1 - C_0 \min\left\{1, \frac{1}{C_1 k L}\right\}$$

for some constants C_0, C_1. We can thus prove the result we want:

$$\mathbb{E}'[\overline{\phi_{y_0}}\phi_{y_R}] = \mathbb{E}'[\psi_{y_1} \cdots \psi_{y_R}] = \mathbb{E}'[\psi_{y_1}] \cdots \mathbb{E}'[\psi_{y_R}],$$

and plugging in our bound above yields the desired generalization of Mermin-Wagner.

Remark 43. *This doesn't work in higher dimensions – if we use the $O(n)$ model we actually don't have decay of correlations at large β at all by the "infrared bound," so in fact no argument of this type can work. So the question is*

how we can prove confinement for four-dimensional lattice gauge theories

From here, we can finally prove the confinement result we want:

Proof of Theorem 38. First we do the proof for $G = U(1)$. Consider a rectangle of side lengths $R \leq T$ in three dimensions, and now condition on all horizontal edges (so now we are conditioning on all edges sitting on parallel planes cutting through the rectangle). Then the vertical edges within one slab are independent of the vertical edges in another slab, and within each slab the vertical edges actually follow a density like in the lemma we just proved (thinking of each edge as a spin). So the correlation between a pair of edges on the two vertical sides will decay as $e^{-C \log R}$ for some C, and we multiply this over all T edges to get the desired result. (Remember that now the R edges are fixed because we've conditioned on horizontal edges.)

Now for the general case, recall that our density takes the form

$$\exp\left(-\beta \sum_p \mathrm{Re}(\mathrm{Tr}(I - U_p))\right).$$

Consider a new system now with space of configurations $(G \times U(1))^E$ (so on each directed edge we have a matrix along with a point on the unit circle, and we define $U(y,x) = U(x,y)^{-1}$ and $\xi(y,x) = \xi(x,y)^{-1}$). With this joint system (U, ξ), now consider the density

$$\exp\left(-\beta \sum_p \mathrm{Re}(\mathrm{Tr}(I - \xi_p U_p))\right),$$

where similarly ξ_p is the product of the $U(1)$ values along the plaquette. Specifically, suppose U is distributed according to the original lattice gauge theory, and (V, ξ) is distributed according to this new lattice gauge theory. Then we can define a new configuration ξV in the original space where $(\xi V)_e = \xi_e V_e$ (since G contains $U(1)$); **we claim** ξV has the same distribution as U. If we prove this claim, then the result follows because the ξs just come out of the expectation and so we can condition on all of the Vs and get the intended decay.

But this claim is very similar to what we did before with the change of variables: if $\tilde{Z}$ is the normalizing constant for the new model and Z is the one for the original model (with respect to the normalized Haar measure), we can show that $\tilde{Z} = Z$. Indeed, for any ξ, we can integrate over V and write

$$\int \exp\left(-\beta \sum \mathrm{Re}(\mathrm{Tr}(I - \xi_p V_p))\right) \int \prod_e d\sigma_0(V_e) = \int \exp\left(-\beta \sum \mathrm{Re}(\mathrm{Tr}(I - V_p))\right) \int \prod_e d\sigma_0(V_e) = Z.$$

by invariance of Haar measure (replacing $\xi_{e_1} V_{e_1} \xi_{e_2} V_{e_2} \xi_{e_3} V_{e_3} \xi_{e_4} V_{e_4}$ by the version without ξs by a change of variable). So if we then integrate over the base measure we get $\tilde{Z} = Z$. The same logic then yields for any fixed ξ

$$\int f(\xi U) \exp\left(-\beta \sum \mathrm{Re}(\mathrm{Tr}(I - \xi_p U_p))\right) \prod d\sigma_0(U_e) = \int f(U) \exp\left(-\beta \sum \mathrm{Re}(\mathrm{Tr}(I - U_p))\right) \prod d\sigma_0(U_e),$$

and so integrating over ξ yields the same thing and then dividing by the corresponding normalizing constants shows expectations are equal, as desired. □

The challenge is next to do this for nonabelian groups like $SU(n)$ – it seems like we can generalize this argument, but it's not so clear yet how we actually do so.

Our next discussion is the **deconfinement transition**, showing that for large enough β we actually get the perimeter law in four dimensions in $U(1)$ theory. This requires some preparation – we have to study some discrete differential geometry first, since it makes use of duality. (It would be nice to have a useful duality for nonabelian lattice gauge theories – we can do some things with the character expansion, but it becomes a huge mess.)

Definition 44

Take any $n \geq 1$ and $x \in \mathbb{Z}^n$. There are n positively oriented edges coming out of x; we denote them $dx_1, \cdots, dx_n$. For all $1 \leq k \leq n$ and any collection of indices $1 \leq i_1 < i_2 < \cdots < i_k \leq n$, the edges dx_{i_1} through dx_{i_k} define a positively oriented ***k*-cell** of $\mathbb{Z}^n$, which we denote $dx_{i_1} \wedge \cdots \wedge dx_{i_k}$. (For example, vertices are 0-cells, edges are 1-cells, and plaquettes are 2-cells.)

The point of discrete differential geometry is to do bookkeeping in higher dimensions, since otherwise it's very hard to visualize. We use the convention that if $j_1, \cdots, j_k$ are obtained by applying a permutation π to $i_1, \cdots, i_k$, then

$$dx_{j_1} \wedge \cdots \wedge dx_{j_k} = \operatorname{sgn}(\pi) dx_{i_1} \wedge \cdots \wedge dx_{i_k},$$

where $-dx_{i_1} \wedge \cdots \wedge dx_{i_k}$ is the negatively oriented version of that k-cell $dx_{i_1} \wedge \cdots \wedge dx_{i_k}$.

Definition 45

Let R be a commutative ring (often we'll take $\mathbb{R}$ or $\mathbb{C}$ or $\mathbb{Z}$). An ***R*-valued *k*-form** on $\mathbb{Z}^n$ is a map from the set of k-cells into R, such that $f(c) = 0$ for all but finitely many k-cells c. We write this map as a sum over all possible cells and points x

$$f(x) = \sum_{1 \leq i_1 < \cdots < i_k \leq n} f_{i_1, \cdots, i_k}(x) dx_{i_1} \wedge \cdots \wedge dx_{i_k},$$

where $f_{i_1, \cdots, i_k}(x)$ is f evaluated at the k-cell $dx_{i_1} \wedge \cdots \wedge dx_{i_k}$. If $k < 0$ or $k > N$, the only k-form is denoted 0 (to be compatible with various other definitions we'll soon have).

So a 0-form is just a map from $\mathbb{Z}^n$ into R, and a 1-form is a map from the set of positively oriented edges into R, and so on.

Definition 46

Let $h : \mathbb{Z}^n \to R$ be a function. For each $x \in \mathbb{Z}^n$, define the difference operator $\partial_i h(x) = h(x + e_i) - h(x)$. For a k-form f (for some $0 \leq k \leq n-1$), the **discrete exterior derivative** df is a $(k+1)$-form defined as follows:

$$df(x) = \sum_{1 \leq i_1 < \cdots < i_k \leq n} \sum_{i=1}^{n} \partial_i f_{i_1, \cdots, i_k}(x) dx_i \wedge dx_{i_1} \wedge \cdots \wedge dx_{i_k}.$$

Notice that we do sum over $(k+1)$ indices, but many of the terms can be combined or simplified. Indeed, in the inner sum if i is any of $i_1, \cdots, i_k$ then the term is just zero, and otherwise we can do some increasing rearrangement while picking up some -1 factors. We'll see a more concrete example next time and go from there!

10 October 27, 2025

Last time, we started discussing discrete differential geometry and considered the cell complex of $\mathbb{Z}^n$ consisting of k-cells $dx_{i_1} \wedge \cdots \wedge dx_{i_k}$ for $0 \leq k \leq n$. The idea is that dx_i represents an edge pointing in the $+e_i$ direction starting from x, and $-dx_i$ is its reversal. Then $dx_1 \wedge dx_2$ would be the plaquette traversing along $(x, x+e_1, x+e_1+e_2, x+e_2)$, and $-dx_1 \wedge dx_2$ would be that same plaquette in the reverse (clockwise) orientation. Then $dx_1 \wedge dx_2 \wedge dx_3$ would be a positively oriented cube, but it's a little harder to think about what orientation means and it's now just more of an abstract concept.

More generally, for arbitrary $i_1, \cdots, i_k \leq n$, we said that $dx_{i_1} \wedge dx_{i_k} = 0$ if $i_1, \cdots, i_k$ are not all distinct, and otherwise it is s times the increasing rearrangement where s is the sign of the permutation (so $dx_1 \wedge dx_1 = 0$ and $dx_2 \wedge dx_1 \wedge dx_3 = -dx_1 \wedge dx_2 \wedge dx_3$). We'll be using these conventions soon – we often do have situations where the indices aren't distinct or are out of order.

We defined R-valued k-forms in Definition 45 – here is a quick related definition:

Definition 47

The **support** of a k-form f, denoted $\operatorname{supp} f$, is the set of all $x \in \mathbb{Z}^n$ with $f_{i_1,\cdots,i_k}(x) \neq 0$ for some $i_1, \cdots, i_k$.

We also defined discrete exterior derivatives in Definition 46, and in particular the latter definition

$$df(x) = \sum_{1 \leq j_1 < \cdots < j_k \leq n} \sum_{i=1}^{n} \partial_i f_{j_1,\cdots j_k}(x) dx_i \wedge dx_{j_1} \wedge \cdots \wedge dx_{j_k}$$

requires repeated and out-of-order indices. For example if f is a 0-form, then we get something like the ordinary derivative

$$df(x) = \sum_{i=1}^{n} \partial_i f(x) dx_i,$$

but then with larger k we often have i equal to one of $i_1, \cdots, i_k$. Indeed, writing $g = df$, we have

$$g(x) = \sum_{1 \leq i_1 < \cdots < i_{k+1} \leq n} g_{i_1,\cdots,i_{k+1}}(x) dx_{i_1} \wedge \cdots \wedge dx_{i_{k+1}},$$

and we're curious what each component $g_{i_1,\cdots,i_{k+1}}(x)$ looks like in terms of f. For $k = n$ we know all components must be zero because we never get distinct indices, and otherwise we have casework over all of the possible indices that could have been the i in the definition:

$$g_{i_1,\cdots,i_{k+1}}(x) = \sum_{1 \leq j \leq k+1} (-1)^{j-1} \partial_{i_j} f_{i_1,\cdots,\widehat{i_j},\cdots,i_{k+1}}(x),$$

where the notation $f_{i_1,\cdots,\widehat{i_j},\cdots,i_{k+1}}(x)$ means that we omit the coordinate i_j so that we only have k indices. (The factor of $(-1)^{j-1}$ comes from how many times we must swap i over to be in the right relative order, since $(i, j_1, \cdots, j_{k+1})$ is of the form $(i_j, i_1, \cdots, \widehat{i_j}, \cdots, i_{k+1})$.)

Example 48

Suppose $k = 2$ and $n = 4$, and we consider the 2-form with just two terms

$$f(x) = f_{12}(x) dx_1 \wedge dx_2 + f_{13}(x) dx_1 \wedge dx_3.$$

Then $df(x)$ should be a 3-form, and specifically we have

$$\begin{aligned} df(x) &= \sum_{i=1}^{4} (\partial_i f_{12}(x))\, dx_i \wedge dx_1 \wedge dx_2 + \sum_{i=1}^{4} (\partial_i f_{13}(x))\, dx_i \wedge dx_1 \wedge dx_3 \\ &= (\partial_3 f_{12}(x))\, dx_3 \wedge dx_1 \wedge dx_2 + (\partial_4 f_{12}(x))\, dx_4 \wedge dx_1 \wedge dx_2 \\ &\quad + (\partial_2 f_{13}(x))\, dx_2 \wedge dx_1 \wedge dx_3 + (\partial_4 f_{13}(x))\, dx_4 \wedge dx_1 \wedge dx_3 \end{aligned}$$

because the other terms have repeated indices, and then we can combine the first and third term and reorder up to

signs to get

$$df(x) = (\partial_3 f_{12}(x) - \partial_2 f_{13}(x))\, dx_1 \wedge dx_2 \wedge dx_3 + (\partial_4 f_{12}(x))\, dx_1 \wedge dx_2 \wedge dx_4 + (\partial_4 f_{13}(x))\, dx_1 \wedge dx_3 \wedge dx_4.$$

Definition 49

A k-form is **closed** if $df = 0$, and it is **exact** if $f = dg$ for some g.

Lemma 50

For all f, we have $ddf = 0$.

This corresponds to the result from ordinary differential geometry saying that $d^2 = 0$.

Proof. If f is a k-form for some $k \geq n-1$, then we already have $ddf = 0$ because the only k-form for $k > n$ is 0. Thus we can assume $k < n-1$. Since the exterior derivative is a linear operator, we can check that

$$ddf(x) = \sum_{1 \leq i_1 < \cdots < i_k \leq n} \sum_{1 \leq i,j \leq n} \partial_j \partial_i f_{i_1,\cdots,i_k}(x) dx_j \wedge dx_i \wedge dx_{i_1} \wedge \cdots \wedge dx_{i_k},$$

and we wish to show that all components are zero. But if $i = j$ then this term is zero because of a repeated index, and otherwise we claim the contributions from (i,j) and (j,i) cancel each other out. Indeed, we have a negative sign coming from $dx_i \wedge dx_j$ versus $dx_j \wedge dx_i$, so it suffices to check that $\partial_j \partial_i f_{i_1,\cdots,i_k}(x) = \partial_i \partial_j f_{i_1,\cdots,i_k}(x)$. And this can just be verified directly, since for any function h we have

$$\begin{aligned} \partial_i \partial_j h &= \partial_i (h(x + e_j) - h(x)) \\ &= h(x + e_j + e_i) - h(x + e_j) - h(x + e_i) + h(x) \end{aligned}$$

which is symmetric in i and j. □

Definition 51

A set $B \subseteq \mathbb{Z}^n$ is a **hypercube** if it is of the form $([a_1, b_1] \times [a_n, b_n]) \cap \mathbb{Z}^n$, where all a_i, b_i are integers and $b_i - a_i$ is the same for all i. The **interior** of B is B without its boundary.

Lemma 52 (Discrete Poincaré lemma)

Take any $1 \leq k \leq n-1$, and let R be any (commutative, though this is not necessary) ring. If f is an R-valued closed k-form on $\mathbb{Z}^n$ (meaning $df = 0$), then there exists an R-valued $(k-1)$-form g with $dg = f$, such that $\text{supp}(g)$ is contained in the smallest hypercube containing $\text{supp}(f)$ in its interior.

In words, the previous result says "any exact form is closed," and this says "any closed form is exact." This would not be true if our manifold had genus greater than 0, but it's true in our current setting.

Proof. Let B be the smallest hypercube containing $\text{supp}(f)$ in its interior. Without loss of generality, suppose $B = [a,b]^n \cap \mathbb{Z}^n$ for integers a, b. For each $a \leq r \leq b$, define the slabs

$$B_r = B \cap (\mathbb{Z}^{n-1} \times \{r\}).$$

We can now define the desired $(k-1)$-form g as follows. Let $g_{i_1,\cdots,i_{k-1}}(x) = 0$ for all $i_1, \cdots, i_{k-1}$ and all $x \notin B$. Also, if $k \geq 2$ and the last coordinate is $i_{k-1} = n$, then set $g_{i_1,\cdots,i_{k-1}}(x) = 0$. Now if $k = 1$ or if $i_{k-1} < n$, suppose $x \in B_r$. We define the value of g inductively on the slab: for $r = a$ again $g_{i_1,\cdots,i_{k-1}}(x) = 0$, and for the inductive step define

$$g_{i_1,\cdots,i_{k-1}}(x) = g_{i_1,\cdots,i_{k-1}}(x - e_n) + (-1)^{k-1} f_{i_1,\cdots,i_{k-1},n}(x).$$

(So if $k = 1$ we would have no indices on the g, and we have $g(x) = g(x - e_n) + f_n(x)$. Of course, we don't really need to distinguish these two cases because we'd automatically repeatedly get zero in the inductive step.)

This is supported on B (**note: actually, we have to modify it a bit, and this is elaborated on more next lecture**), and now we need to check that it works – that is, we wish to show that $h = dg$ is actually f. Fix indices $i_1, \cdots, i_k$ and take $x \in B_r$; we will prove by induction on r that $h_{i_1,\cdots,i_k}(x) = f_{i_1,\cdots,i_k}(x)$, and that's enough because $h = f = 0$ outside B. (Here, we use that $dg(x)$ involves going only one step in the positive direction from x, so we only end up getting nonzero things on the boundary because g is zero on all of the boundaries which matter and otherwise supported on the interior.)

For **case 1**, suppose $r = a$ and $i_k < n$. Then $dx_{i_1} \wedge dx_{i_k}$ is a k-cell contained in the slab B_a, but we've defined $h = 0$ on that slab and f is zero on the slab by definition.

Next, for **case 2**, suppose $r = a$ (though this isn't actually necessary) and $i_k = n$. Then

$$h_{i_1,\cdots,i_k}(x) = \sum_{1 \leq j \leq k} (-1)^{j-1} \partial_j g_{i_1,\cdots,\widehat{i_j},\cdots,i_k}(x),$$

But we know $g_{j_1,\cdots j_k}(x) = 0$ if $j_k = n$, and $i_k = n$, so the only way to get a nonzero contribution is if we take $j = k$ in that sum to cancel out the last index:

$$h_{i_1,\cdots,i_k}(x) = (-1)^{k-1} \partial_n g_{i_1,\cdots,i_{k-1}}(x),$$

and by the recursive way we defined g we have $\partial_n g_{i_1,\cdots,i_k}(x) = (-1)^{k-1} f_{i_1,\cdots,i_{k-1},n}(x)$, so substituting this in yields $f_{i_1,\cdots,i_{k-1},n}(x) = f_{i_1,\cdots,i_k}(x)$ as desired.

In **case 3**, now consider $a < r \leq b$ and $i_k = n$. Then the same proof goes through as in case 2 – we don't need the induction. But finally we have **case 4** where $a < r \leq b$ and $i_k < n$. Then letting $u = dh$ and $v = df$, we find that $u = ddg = 0$ by construction and we are also given that $v = df = 0$ by assumption. Then for any $x \in B_r$, setting $y = x - e_n \in B_{r-1}$, we get the equations

$$0 = u_{i_1,\cdots,i_k,n}(y) = (-1)^k \partial_n h_{i_1,\cdots,i_k}(y) + \sum_{1 \leq j \leq k} (-1)^{j-1} \partial_{i_j} h_{i_1,\cdots,\widehat{i_j},\cdots,i_k,n}(y)$$

and

$$0 = v_{i_1,\cdots,i_k,n}(y) = (-1)^k \partial_n f_{i_1,\cdots,i_k}(y) + \sum_{1 \leq j \leq k} (-1)^{j-1} \partial_{i_j} f_{i_1,\cdots,\widehat{i_j},\cdots,i_k,n}(y).$$

We claim the two sums on the right-hand side are actually equal, because

$$\partial_{i_j} h_{i_1,\cdots,\widehat{i_j},\cdots,i_k,n}(y) - \partial_{i_j} f_{i_1,\cdots,\widehat{i_j},\cdots,i_k,n}(y) = \left(h_{i_1,\cdots,\widehat{i_j},\cdots,i_k,n}(y + e_{i_j}) - h_{i_1,\cdots,\widehat{i_j},\cdots,i_k,n}(y)\right) - \left(f_{i_1,\cdots,\widehat{i_j},\cdots,i_k,n}(y + e_{i_j}) - f_{i_1,\cdots,\widehat{i_j},\cdots,i_k,n}(y)\right),$$

and because $y \in B_{r-1}$ and $y + e_{i_j} \in B_{r-1}$ as well (we're not moving out of the slab), **by the inductive hypothesis** we get that $f = h$ on B_{r-1} so the whole term goes away. Thus we can take those two equations for h and f and subtract them to get

$$0 = \partial_n h_{i_1,\cdots,i_k}(y) - \partial_n f_{i_1,\cdots,i_k}(y),$$

which rearranges to

$$h_{i_1,\cdots,i_k}(y+e_n) - h_{i_1,\cdots,i_k}(y) = f_{i_1,\cdots,i_k}(y+e_n) - f_{i_1,\cdots,i_k}(y).$$

Now the terms with y are equal by inductive hypothesis, so the terms with $y+e_n = x$ are also equal, completing the induction and proof. □

This is a pretty powerful tool – we'll come to it soon, but this basically says that no matter what closed curve we have, we can always find a surface bounding it in the lattice. And that's why the inductive step is needed here. Next time, we'll discuss the discrete coderivative and Hodge dual and do the corresponding Poincaré lemma there as well.

11 October 29, 2025

We proved the discrete Poincaré lemma last time – we didn't actually need to have a hypercube with all equal side lengths, so we can instead consider $\prod_{i=1}^n [a_i, b_i] \cap \mathbb{Z}^n$ (and let the interior be the same thing but with open intervals instead of closed intervals). In particular, note that we must actually have $1 \le k \le n-1$ (so it's not true that $f = dg$ for any n-form f, even though f is automatically closed). But we have to elaborate a little bit more on one of the details to complete the proof.

As a reminder, the key components of the proof of Lemma 52 are as follows (now allowing for a more general hypercube): we define slabs $S_r = \mathbb{Z}^{n-1} \times \{r\}$ for all $r \in \mathbb{Z}$. For all $x \in S_r$ with $x \le a_n$, we can let $g(x) = 0$. Then for $r > a_n$, we define $g(x)$ inductively via

$$g_{i_1,\cdots,i_{k-1}}(x) = g_{i_1,\cdots,i_{k-1}}(x-e_n) + (-1)^{k-1} f_{i_1,\cdots,i_{k-1},n}(x-e_n)$$

if $i_{k-1} < n$ and 0 otherwise.

Fact 53

To avoid needing to write out these different cases, we'll adopt the convention that for $f = \sum_{1\le i_1<\cdots<i_k\le n} f_{i_1,\cdots,i_k} dx_{i_1} \wedge \cdots \wedge dx_{i_k}$, we write $f_{j_1,\cdots j_k} = s f_{i_1,\cdots,i_k}$ if $(i_1,\cdots,i_k)$ is the increasing rearrangement of the distinct indices $j_1,\cdots,j_k$ and s is the sign of the associated permutation. (We also write $f_{j_1,\cdots j_k} = 0$ if the js are not distinct.)

With this definition, we indeed have $dg = f$, and this doesn't even require the hypercube. Furthermore, for any $x = (x_1,\cdots,x_n)$ with $x_i \notin [a_i, b_i]$ for some $i < n$ we have inductively that $g(x) = 0$ (since we start off with zero and we never get any nonzero contributions from the f term in the recursive formula). So now we just have to worry about the nth coordinate. If $x_n \le a_n$, then we do still have $G(x) = 0$. But if $x_n \ge b_n$, nothing here ensures that g will vanish, and in fact g doesn't have to vanish as stated.

If $x_n \ge b_n + 1$, then $f_{i_1,\cdots,i_{k-1},n}(x-e_n) = 0$ and thus g becomes a constant as we keep going up in the nth coordinate; all we need to do is ensure that it is zero. We first claim that **if** $k = 1$, then we do actually always get $g(x) = 0$ when $x_n \ge b_n$ (so nothing goes wrong). By our logic we just have to check that $g(x) = 0$ on the slab S_{b_n}; indeed, $dg(x) = f(x) = 0$ on the slab, and since g is a 0-form we have

$$dg(x) = \sum_{i=1}^n \partial_i g(x) dx_i = \sum_{i=1}^n (g(x+e_i) - g(x)) dx_i,$$

so g must actually be constant on the slab S_{b_n} because all directional derivatives vanish. And since we know g is zero outside a finite box, that means g is zero everywhere, as desired. (So in particular this means the proof is completely correct for $n = 2$.)

But for $2 \le k \le n-1$ (and in particular $n \ge 3$), we will need an induction. Supposing that the lemma holds up to dimension $(n-1)$, the strategy is to produce a $(k-1)$-form w possibly with infinite support with $dw = 0$, such that **(1)** $w(x) = g(x)$ except on the slab S_{r-1}, and **(2)** w vanishes outside the hypercube on S_{b_n-1}. Then $g' = g - w$ will do the job because $dg' = dg - dw = dg = f$ and g' is indeed vanishing outside the hypercube.

Example 54

To illustrate what might be going wrong, suppose we have a big square in the xy-plane in three dimensions, and then we attach perpendicular plaquettes to each of its edges (in either the xz or yz-directions). Suppose we let our ring R just be $\mathbb{Z}/2\mathbb{Z}$ to not worry about signs. Then if we have a 2-form f which is 1 on all of the attached plaquettes and 0 on all other plaquettes, then this is closed (because if we add up the numbers on the sides of any cube, we always get 0).

But if we try to find a g by our algorithm (which is a 1-form assigning values to edges), then g will be all zero up until the slab containing our big square, and then we get 1s on the top edges of all of our perpendicular plaquettes, since we are adding the values of f from one slab to the next. Then as we continue going up, g will continue being 1 on those vertical columns, which is not good. So what we should do instead is modify things a little bit to get finite support, and we do this by using the "vertical edges" popping out from the big square.

What we do is define a $(k-1)$-form q on $\mathbb{Z}^{n-1}$ by just looking on the particular slab

$$q(y) = (-1)^k g(y, b_n).$$

Then $dq = dg|_{S_{b_n}} = f|_{S_{B_n}} = 0$ (since we only look along the horizontal directions when computing dq), and q is supported on $\prod_{i=1}^{n-1}(a_i, b_i)$. But remember this time it does not mean q is constant, so we're not immediately done. Instead, by the inductive hypothesis, we have some $(k-2)$-form p on $\mathbb{Z}^{n-1}$, supported only on $\prod_{i=1}^{n-1}[a_i, b_i]$, so that $dp = q$. Therefore we can define w via the following: if $x \in S_r$ for $r \ge b_n$, then $w(x) = g(x)$. Otherwise if $x = (x_1, \cdots, x_{n-1}, b_n - 1)$, then $w_{i_1,\cdots,i_{k-2},n} = p_{i_1,\cdots,i_{k-2}}(x_1, \cdots, x_{n-1})$, and otherwise $w = 0$ everywhere. (So we're just defining w using p along the vertical edges pointing in the e_n direction.) We thus get $g' = g - w$ supported on the hypercube, and so it only remains to show that $dw = 0$. Indeed, we break into cases.

1. For $x \in S_r$ where $r \ge b_n$, we have $dw(x) = dg(x) = f(x) = 0$ so everything is okay (dw only looks in positive directions).

2. For $x \in S_r$ for some $r \le b_n - 2$, we have

$$(dw)_{i_1,\cdots,i_k}(x) = \sum_{j=1}^{n} (-1)^{j-1} \partial_{i_j} w_{i_1,\cdots,\widehat{i_j},\cdots,i_k}(x).$$

   For $i_k < n$ we have all terms restricted within a slab (since $x + e_{i_j} \in S_r$) and below $b_n - 1$ we defined w to be identically zero. Thus the whole sum here will vanish. On the other hand if $i_k = n$, the only partial derivative that can contribute is the last one, so that

$$(dw)_{i_1,\cdots,i_{k-1},n}(x) = (-1)^{n-1}(w_{i_1,\cdots,i_{k-1}}(x + e_n) - w_{i_1,\cdots,i_{k-1}}(x)),$$

   but now the latter term is zero, and $x + e_n \in S_{r+1}$ is at level $b_n - 1$ or below, and the last index satisfies $i_{k-1} < n$ so $w_{i_1,\cdots,i_{k-1}}(x + e_n)$ is also zero.

3. Finally for $x \in S_{b_n-1}$, we again consider $i_k < n$ and $i_k = n$ separately. For $i_k < n$ again all partial derivatives are zero because we stay within the slab S_{b_n-1}, and w is only nonzero if it has an index n. And finally for the

remaining case $i_k = n$, we have for all $1 \le j \le k-1$ that

$$\begin{aligned}\partial_{i_j} w_{i_1,\cdots,\widehat{i_j},\cdots,i_{k-1},n}(x) &= w_{i_1,\cdots,\widehat{i_j},\cdots,i_{k-1},n}(x+e_{i_j}) - w_{i_1,\cdots,\widehat{i_j},\cdots,i_{k-1},n}(x) \\ &= p_{i_1,\cdots,\widehat{i_j},\cdots,i_{k-1}}(x'+e_{i_j}) - p_{i_1,\cdots,\widehat{i_j},\cdots,i_{k-1}}(x') \\ &= \partial_{i_j} p_{i_1,\cdots,\widehat{i_j},\cdots,i_{k-1}}(x')\end{aligned}$$

where x' is the first $(n-1)$ coordinates of x. Furthermore we also have a term $\partial_n w_{i_1,\cdots,i_{k-1}}(x) = w_{i_1,\cdots,i_{k-1}}(x+e_n) - w_{i_1,\cdots,i_{k-1}}(x) = g_{i_1,\cdots,i_{k-1}}(x+e_n) - 0$. So we can now plug this all back into our definition of dw to find

$$\begin{aligned}(dw)_{i_1,\cdots,i_{k-1},n}(x) &= \sum_{j=1}^{k-1}(-1)^{j-1}\partial_{i_j} p_{i_1,\cdots,\widehat{i_j},\cdots,i_{k-1}}(x') + (-1)^{k-1} g_{i_1,\cdots,i_{k-1}}(x+e_n) \\ &= q_{i_1,\cdots,i_{k-1}}(x') + (-1)^{k-1} g_{i_1,\cdots,i_{k-1}}(x+e_n) \\ &= (-1)^k g_{i_1,\cdots,i_{k-1}}(x+e_n) + (-1)^{k-1} g_{i_1,\cdots,i_{k-1}}(x+e_n) \\ &= 0.\end{aligned}$$

Thus our modification satisfies all of our requirements and we've finished the proof.

Remark 55. *Even this proof, we can modify it a bit so we don't need the "interior" condition: support of g should be contained in the smallest cube containing the support of f.*

This result is useful for certain "duality arguments" that we'll soon see. For that, we'll need to introduce another notion called the **codifferential** or **coderivative**.

Definition 56

For any function h, define $\overline{\partial_i} h(x) = h(x) - h(x-e_i)$. For a k-form f with $1 \le k \le n$, the **codifferential** or **coderivative** δ is defined via δf being the $(k-1)$-form

$$\delta f(x) = \sum_{1 \le i_1 < \cdots < i_k \le n} \sum_{\ell=1}^{k} (-1)^{\ell} \overline{\partial_{i_\ell}} f_{i_1,\cdots,i_k}(x) dx_{i_1} \wedge \cdots \wedge \widehat{dx_{i_\ell}} \wedge \cdots \wedge dx_{i_k}.$$

In coordinates, we can write down $g = \delta f$ and we want to compute $g_{i_1,\cdots,i_{k-1}}(x)$ for any choice of i indices. We have to group together terms where we have some additional index, and the formula works out to

$$g_{i_1,\cdots,i_{k-1}}(x) = -\sum_{i=1}^{n} \overline{\partial_i} f_{i,i_1,\cdots,i_{k-1}}(x).$$

Pictorially, what this says is the following. For a 1-form (numbers on edges), the **differential** operator gives us a 2-form which is numbers on plaquettes given by adding and subtracting numbers along the edges. Meanwhile, the **codifferential** must go from a 2-form (numbers on plaquettes) to a 1-form, and what it does is add and subtract the values of all plauqettes that contain our particular edge. So in general the codifferential yields some linear combination of all of the $(k+1)$-cells containing our k-cell.

Definition 57

Let $*\mathbb{Z}^n$ be the dual lattice of $\mathbb{Z}^n$, meaning that it is the set of midpoints $(x_1 + \frac{1}{2}, \cdots, x_n + \frac{1}{2})$ of n-cells in $\mathbb{Z}^n$. For any $y \in *\mathbb{Z}^n$, let $dy_i = (y, y - e_i)$ point in the negative direction instead. A **dual k-cell** is the $(n-k)$-cell defined as follows. For n-cells we have

$$*(dx_1 \wedge * \wedge dx_n) = y = (x_1 + \frac{1}{2}, \cdots, x_N + \frac{1}{2}),$$

and more generally

$$*(dx_{i_1} \wedge \cdots \wedge dx_{i_k}) = s dy_{j_1} \wedge \cdots \wedge dy_{j_{n-k}},$$

where $j_1, \cdots, j_{n-k}$ are the increasing rearrangement of the remaining indices and s is the sign of the permutation $(i_1, \cdots, i_k, j_1, \cdots, j_{n-k})$ (so here the is are increasing and separately the js are increasing).

Example 58

In two dimensions, the dual of dx_1 would be $*dx_1 = dy_2$ (so we take our right-pointing edge and rotate it 90 degrees clockwise around the midpoint); similarly the dual of dx_2 is actually $-dy_1$. And in three dimensions, the dual of dx_1 would be a plaquette, and it's $dy_2 \wedge dy_3$, which is the plaquette perpendicularly bisecting that edge dx_1.

Definition 59

We similarly define the dual of a dual cell via

$$*(dy_{j_1} \wedge \cdots \wedge dy_{j_{n-k}}) = (-1)^{k(n-k)} s dx_{i_1} \wedge \cdots \wedge dx_{i_k}$$

for s the sign of the permutation $(i_1, \cdots, i_k, j_1, \cdots, j_{n-k})$ since the sign of the permutation $(j_1, \cdots, j_{n-k}, i_1, \cdots, i_{n-k})$ is $(-1)^{k(n-k)}$ times the sign of $(i_1, \cdots, i_k, j_1, \cdots, j_{n-k})$.

Definition 60

The dual of a k-form f is the $(n-k)$-form $*f$ on the dual lattice, given by

$$*f(y) = \sum_{1 \le i_1 < \cdots < i_k \le n} f_{i_1, \cdots, i_k}(x) s dy_{j_1} \wedge \cdots \wedge dy_{j_{n-k}}.$$

That is, we take each k-cell and put the value we had on the k-cell on the same dual cell with carefully adjusted signs (and this is okay because $\binom{n}{k} = \binom{n}{n-k}$).

Our definitions here are carefully set up so that $** f = (-1)^{k(n-k)} f$ for all f. And next time, we'll use this to write the coderivative in terms of the derivative so that we won't have to reprove the Poincaré lemma.

12 November 3, 2025

Last time, we discussed the Hodge dual and coderivative operator.

Lemma 61

Let f be a k-form for $1 \leq k \leq n$. Then the coderivative can be written

$$\delta f = (-1)^{n(k+1)+1} * d * f(y),$$

where $y = *(dx_1 \wedge \cdots \wedge dx_n)$ is the dual of x.

(We can check that the dimensions work out correctly when we apply $*$, then d, then $*-$ we do get a $(k-1)$-form.)

Proof. Recall the definition

$$*f(y) = \sum_{1 \leq i_1 < \cdots < i_k \leq n} f_{i_1, \cdots, i_k}(x) s dy_{j_1} \wedge \cdots \wedge dy_{j_{n-k}}$$

for s the corresponding sign of the permutation. Thus taking the derivative on both sides, we have

$$d*f(y) = \sum_{1 \leq i_1 < \cdots < i_k \leq n} \sum_{1 \leq i \leq n} \overline{\partial_i} f_{i_1, \cdots, i_k}(x) s dy_i \wedge dy_{j_1} \wedge \cdots \wedge dy_{j_{n-k}}.$$

Let's think now about what happens if we try to take $*$ on both sides. If $i \notin \{i_1, \cdots, i_k\}$, then we always have $i \in j_1, \cdots, j_{n-k}\}$, meaning that $dy_i \wedge dy_{j_1} \wedge \cdots \wedge dy_{j_{n-k}} = 0$ and we don't have to consider those terms. And if $i = i_\ell$ for some ℓ, then

$$*(dy_i \wedge dy_{j_1} \wedge \cdots \wedge dy_{j_{n-k}}) = (-1)^{k(n-k)+n-k+\ell-1} s dx_{i_1} \wedge \cdots \wedge \widehat{dx_{i_\ell}} \wedge \cdots \wedge dx_{i_k}.$$

The reason for the power of -1 is that it is the sign of the permutation $(i_\ell, j_1, \cdots, j_{n-k}, i_1, \cdots, \widehat{i_\ell}, i_k)$ (since it takes $(n-k+\ell-1)$ swaps to put i_ℓ back where it's omitted, and then $(n-k)$ swaps to move each i past all of the js). Therefore

$$\begin{aligned} *d*fy = (-1)^{k(n-k)+n-k-1} \sum_{1 \leq i_1 < \cdots < i_k \leq n} \sum_{\ell=1}^{k} (-1)^\ell \overline{\partial_{i_\ell}} f_{i_1, \cdots, i_k}(x) dx_{i_1} \wedge \cdots \wedge \widehat{dx_{i_\ell}} \wedge \cdots \wedge dx_{i_k} \\ = (-1)^{n(k+1)-1} \delta f(x), \end{aligned}$$

where in the last step we use that $k(n-k)+n-k-1 = n(k+1)-1-k(k+1)$ and the latter term is always even. □

The utility of this is that we automatically get the Poincaré lemma for the coderivative using the one for the exterior derivative:

Proposition 62 (Poincaré lemma for the coderivative)

Take any $1 \leq k \leq n-1$, and let f be an R-valued k-form on $\mathbb{Z}^n$ (for some ring R) with $\delta f = 0$. Then there exists some $(k+1)$-form h such that $f = \delta h$ and where $\text{supp}(h)$ is contained in the smallest hypercube containing $\text{supp}(f)$.

Proof. Since $\delta f = 0$, the lemma above implies that $d*f = 0$. Letting B be the smallest hypercube containing $\text{supp}(f)$, we can define $*B$ to be the dual hypercube $\{*(dx_1 \wedge \cdots \wedge dx_n) : x \in B\}$, and by construction we have $\text{supp}(*f) \subseteq *B$. Therefore the (usual) Poincaré lemma implies the existence of an $(n-k-1)$-form g on $*\mathbb{Z}^n$ with $dg = *f$ and $\text{supp}(g) \subseteq *B$; we can then define $h = (-1)^{-(k-1)(n-k+1)-k(n-k)-nk-1} * g$. Then we can again directly check that

$\operatorname{supp}(h) \subseteq B$, and using that $**f = (-1)^{k(n-k)}f$ yields

$$*h = (-1)^{-k(n-k)-nk-1}g.$$

Therefore

$$\delta h = (-1)^{nk+1} * d * h = (-1)^{-k(n-k)} * dg = (-1)^{-k(n-k)} * (*f) = f,$$

as desired. □

We can now give a geometric interpretation of some of our results. Taking the ring to be $\{0,1\}$, any 1-form is a set of edges and a 2-form is a set of plaquettes. If f is a 1-form, then δf is a function on vertices, and $\delta f = 0$ means that every vertex has an even number of incident edges. That must mean that f is a union of closed loops (start with any vertex and follow some edge we haven't used, and repeat until we've used everything up; we can't get stuck). What we're saying now is that there is some collection of plaquettes h such that $\delta h = f$; what this means is that the number of plaquettes adjacent to each edge of f is odd. And that's exactly what a surface bounded by each loop encodes. Thus this proves that "for any closed loop, we can find a surface whose boundary is that loop, and we can contain it in the smallest hypercube containing that loop."

Definition 63

For k-forms f, g, we define the inner product

$$(f,g) = \sum_{x\in\mathbb{Z}^n} \sum_{1\le i_1<\cdots<i_k\le n} f_{i_1,\cdots,i_k}(x) g_{i_1,\cdots,i_k}(x)$$

Note that this is the first time we've really actually used the ring structure of the k-forms.

Lemma 64

For any $(k+1)$-form f and any k-form g (for $0 \le k \le n-1$, we have

$$(f, dg) = (\delta f, g).$$

Proof. Let $h = \delta f$ and substitute in the formula for the coderivative (the version in coordinates stated after Definition 56). We find that

$$\begin{aligned}
(h,g) &= -\sum_{x\in\mathbb{Z}^n} \sum_{1\le i_1<\cdots<i_k\le n} \sum_{i=1}^{n} \overline{\partial}_i f_{i,i_1,\cdots,i_k}(x) g_{i_1,\cdots,i_k}(x) \\
&= -\sum_{x\in\mathbb{Z}^n} \sum_{1\le i_1<\cdots<i_k\le n} \sum_{i=1}^{n} (f_{i,i_1,\cdots,i_k}(x) - f_{i,i_1,\cdots,i_k}(x-e_i)) g_{i_1,\cdots,i_k}(x) \\
&= -\sum_{x\in\mathbb{Z}^n} \sum_{1\le i_1<\cdots<i_k\le n} \sum_{i=1}^{n} f_{i,i_1,\cdots,i_k}(x) g_{i_1,\cdots,i_k}(x) + \sum_{x\in\mathbb{Z}^n} \sum_{1\le i_1<\cdots<i_k\le n} \sum_{i=1}^{n} f_{i,i_1,\cdots,i_k}(x-e_i) g_{i_1,\cdots,i_k}(x)
\end{aligned}$$

(all forms are finitely-supported, so there's no issue with moving sums around). But now reindexing the second sum by a change of variable $x \mapsto x + e_i$ doesn't change the overall sum, and then we can combine them back together again to get

$$(h,g) = \sum_{x\in\mathbb{Z}^n} \sum_{1\le i_1<\cdots<i_k\le n} \sum_{i=1}^{n} f_{i,i_1,\cdots,i_k}(x) \partial_i g_{i_1,\cdots,i_k}(x).$$

We only need to consider the case where all indices $i, i_1, \cdots, i_k$ are distinct (otherwise the contribution is zero); let $j_1 < j_2 < \cdots < j_{k+1}$ be the increasing rearrangement of them. Then we have $i = j_\ell$ for some ℓ, and then each term in the sum can be written

$$f_{i,i_1,\cdots,i_k}(x)\partial_l g_{i_1,\cdots,i_k}(x) = (-1)^{\ell-1} f_{j_1,\cdots j_{k+1}}(x)\partial_{j_\ell} g_{j_1,\cdots \widehat{j_\ell},\cdots j_{k+1}}(x)$$

and then summing this over all indices and all ℓ yields exactly the definition of (f, dg) as desired. □

So this proof is basically "summation by parts" similarly to how we normally do "integration by parts."

This is all of the "bookkeeping" we'll do for now – this discrete differential geometry has been very useful for a lot of Professor Chatterjee's work. We'll now think about the **Villain action**, which is easier to think about for abelian gauge theories. Specifically, we'll think about $U(1)$ lattice gauge theory on $\mathbb{Z}^d$ for $d \geq 2$.

Definition 65

Let $\Lambda \subseteq \mathbb{Z}^d$ be finite, and let E be the set of positively oriented edges with both vertices in Λ and P the set of all plaquettes. We previously considered the configuration space $U(1)^E$, and now we'll replace that with the space $[-\pi, \pi)^E$ (via the correspondence of $e^{i\theta(e)}$ on the unit circle to the angle θ). For any $(x, y) \in E$, we then also define $\theta(y, x) = -\theta(x, y)$ (so that $e^{i\theta(y,x)}$ is the inverse of $e^{i\theta(x,y)}$. For a plaquette with vertices x_1, x_2, x_3, x_4 in counterclockwise order, define

$$\theta(p) = \theta(x_1, x_2) + \theta(x_2, x_3) + \theta(x_3, x_4) + \theta(x_4, x_1)$$

(note that this quantity may not be in $[-\pi, \pi)$). Then for a given $\beta > 0$, define the **Villain action**

$$S_\beta(\theta) = -\sum_{p \in P} \log \left[\sum_{n \in \mathbb{Z}} \exp\left(-\frac{\beta}{2}(\theta_p - 2\pi n)^2 \right) \right].$$

This series is always rapidly convergent so there are no problems, and furthermore this takes care of the issue of making a choice of interval $[-\pi, \pi)$ – because of the 2π-periodicity inside the sum it doesn't actually matter. (This is actually the heat kernel on the circle for Brownian motion, and that's how we generalize for other groups.)

Previously in $U(1)$ theory the Wilson action tries to make each plaquette variable close to 1, and this does the same – θ_p will try to concentrate around some multiple of 2π under the Villain action to make this sum close to its maximum value. And the point is that the Villain action works quite well with discrete differential geometry and duality properties.

Lemma 66

The Villain action S_β is gauge-invariant.

Proof. Take any configuration $\theta \in [-\pi, \pi)^E$. (This corresponds to the configuration $U \in U(1)^E$ with $U(e) = e^{i\theta(e)}$ for all e.) For any gauge transformation $g \in U(1)^\Lambda$, let $V = U^g$, meaning that $V(x, y) = g(x)U(x, y)g(y)^{-1}$. So if $\xi \in [-\pi, \pi)^E$ represents V and $\eta \in [-\pi, \pi)^\Lambda$ represents g, the gauge transformation equation means that

$$e^{i\xi(x,y)} = e^{i(\eta(x)+\theta(x,y)-\eta(y))} \implies \xi(x, y) \equiv \eta(x) + \theta(x, y) - \eta(y) \bmod 2\pi.$$

Therefore for any plaquette as in the definition of the Villain action, we have

$$\begin{aligned}\xi_p &= \xi(x_1,x_2)+\xi(x_2,x_3)+\xi(x_3,x_4)+\xi(x_4,x_1)\\ &\equiv \theta(x_1,x_2)+\theta(x_2,x_3)+\theta(x_3,x_4)+\theta(x_4,x_1) \bmod 2\pi\end{aligned}$$

because all of the ηs cancel out (each one is added once and subtracted once). This means $\xi_p \equiv \theta_p \bmod 2\pi$, and therefore

$$\sum_{n\in\mathbb{Z}} \exp\left(-\frac{\beta}{2}(\theta_p-2\pi n)^2\right) = \sum_{n\in\mathbb{Z}} \exp\left(-\frac{\beta}{2}(\xi_p-2\pi n)^2\right).$$

Applying this to every plaquette, the Villain action stays the same under this transformation, as desired. □

This action also makes it possible to "scale to the continuum limit:"

Theorem 67

(**Note**: this was corrected during the subsequent lecture.) Take $d \in \{2,3,4\}$ and $g_0 > 0$. Let $A = \sum_{j=1}^d A_j dx_j$ be a compactly supported smooth $\mathfrak{u}(1)$-valued 1-form on $\mathbb{R}^d$. For any $\varepsilon \in (0,1)$, define the configuration

$$\theta(x, x+\varepsilon e_j) = -i\log(e^{\varepsilon A_j(x)}),$$

where we use the convention that $\log e^{ix} = ix$ for $-\pi \le x < \pi$. (We're writing it this way because εA_j may not necessarily be in $[-\pi,\pi)$.) This means we have $e^{i\theta(x,x+\varepsilon e_j)} = e^{\varepsilon A_j(x)}$. Now for the Villain action, we need to make a small modification

$$\tilde{S}_\beta(\theta) = -\sum_{p\in P}\left(\log\left(\sum_{n\in\mathbb{Z}}\exp\left(-\frac{\beta}{2}(\theta_p-2\pi n)^2\right)\right) - \log\left(\sum_{n\in\mathbb{Z}}\exp\left(-\frac{\beta}{2}(-2\pi n)^2\right)\right)\right)$$

(this just makes it so that we get a finite sum over plaquettes because $\theta_p = 0$ outside of a finite region). Let $\beta = \frac{\varepsilon^{d-4}}{g_0^2}$. If $d<4$, then we have $\lim_{\varepsilon\to 0}\tilde{S}_{\beta(\varepsilon)}(\theta) = S_E(A)$, and if $d=4$ then $\lim_{\varepsilon\to 0}\tilde{S}_{\beta(\varepsilon)}(\theta) = \kappa S_E(A)$ for the constant

$$\kappa = 1 - \frac{4\pi^2}{g_0^2}\frac{\sum_{n\in\mathbb{Z}} n^2 e^{-2\pi^2 n^2/g_0^2}}{\sum_{n\in\mathbb{Z}} e^{-2\pi^2 n^2/g_0^2}}$$

(and we can choose a slightly modified g_0' to get the desired constant instead). Meanwhile if $d \ge 5$, we actually just get $\lim_{\varepsilon\to 0}\tilde{S}_{\beta(\varepsilon)}(\theta) = 0$. Furthermore we indeed have $\kappa > 0$.

Start of proof. In this proof, $O(\varepsilon^\alpha)$ will denote any quantity whose absolute value is bounded by $C\varepsilon^\alpha$ for some C depending only on d, g_0, and A. We have

$$e^{i\theta(x,x+\varepsilon e_j)} = e^{\varepsilon A_j(x)} = e^{i(-i\varepsilon A_j(x))},$$

which means that $\theta(x,x+\varepsilon e_j) \equiv -i\varepsilon A_j(x) \bmod 2\pi$. Around a plaquette, we therefore get by a similar calculation as before that

$$\theta_p \equiv -i\varepsilon(A_j(x_1)+A_k(x_2)-A_j(x_4)-A_k(x_1)) \bmod 2\pi.$$

Thus we can write down the Villain action in terms of A_js:

$$\begin{aligned}\sum_{n\in\mathbb{Z}}\exp\left(-\frac{\beta}{2}(\theta_p-2\pi n)^2\right) &= \sum_{n\in\mathbb{Z}}\exp\left(-\frac{\beta}{2}\left(-i\varepsilon(A_j(x_1)+A_k(x_2)-A_j(x_4)-A_k(x_1))-2\pi n\right)^2\right)\\ &= \sum_{n\in\mathbb{Z}}\exp\left(\frac{\varepsilon^{d-2}}{2g_0^2}\left(A_j(x_1)+A_k(x_2)-A_j(x_4)-A_k(x_1)-\frac{2\pi i n}{\varepsilon}\right)^2\right).\end{aligned}$$

Now because A is smooth we have $A_j(x_1) + A_k(x_2) - A_j(x_4) - A_k(x_1) = \varepsilon(\partial_j A_k(x) - \partial_k A_j(x)) + O(\varepsilon^2)$, and we can plug this in to get

$$\sum_{n\in\mathbb{Z}} \exp\left(\frac{\varepsilon^d}{2g_0^2}\left(\partial_j A_k(x) - \partial_k A_j(x) + O(\varepsilon) - \frac{2\pi i n}{\varepsilon^2}\right)^2\right).$$

Expanding out the square here yields that for any fixed n this summand is

$$\exp\left(-\frac{2\varepsilon^{d-4}\pi^2 n^2}{g_0^2}\right)\exp\left(\frac{\varepsilon^d}{2g_0^2}(\partial_j A_k(x) - \partial_k A_j(x) + O(\varepsilon))^2 - \frac{2\pi i \varepsilon^{d-2} n}{g_0^2}(\partial_j A_k(x) - \partial_k A_j(x) + O(\varepsilon))\right).$$

Now note that the $O(\varepsilon)$ term doesn't depend on n, so if we add the summands for some specific n and $-n$ together, everything except the last term is the same and we can use $e^{ix} + e^{-ix} = 2\cos x$ (because those two angles are exactly negatives). Therefore we get a total contribution of

$$2\exp\left(-\frac{2\varepsilon^{d-4}\pi^2 n^2}{g_0^2}\right)\exp\left(\frac{\varepsilon^d}{2g_0^2}(\partial_j A_k(x) - \partial_k A_j(x) + O(\varepsilon))^2\right)\cos\left(\frac{2\pi\varepsilon^{d-2} n}{g_0^2}(\partial_j A_k(x) - \partial_k A_j(x) + O(\varepsilon))\right).$$

We'll see how to get estimates for this next time! □

13 November 5, 2025

We were proving convergence of the Villain action to the Yang-Mills action last time, and in particular (see above) we had a sum over nonnegative integers n with a cosine term in the summand (specifically the summand we care about is $2\exp\left(-\frac{2\varepsilon^{d-4}\pi^2 n^2}{g_0^2}\right)\exp\left(\frac{\varepsilon^d}{2g_0^2}(\partial_j A_k(x) - \partial_k A_j(x) + O(\varepsilon))^2\right)\cos\left(\frac{2\pi\varepsilon^{d-2} n}{g_0^2}(\partial_j A_k(x) - \partial_k A_j(x) + O(\varepsilon))\right)$.). Continuing the proof now, we now have to consider three separate cases:

1. $d = 4$. By a Taylor approximation, we can write the second exponential as $1 + \frac{\varepsilon^4}{2g_0^2}(\partial_j A_k - \partial_k A_j)^2 + O(\varepsilon^5)$, and we can write the cosine term as $1 - \frac{2\pi\varepsilon^4 n^2}{g_0^4}(\partial_j A_k - \partial_k A_j)^2 + O(\varepsilon^5)n^2 + O(\varepsilon^8)n^4$ (here we use that cosine has all bounded derivatives, so we don't need to go on after this). So multiplying these together and also substituting $d = 4$ into the first exponential, we get

$$2\exp\left(-\frac{2\pi^2 n^2}{g_0^2}\right)\left(1 + \varepsilon^4\left(\frac{1}{2g_0^2} - \frac{2\pi^2 n^2}{g_0^4}\right)(\partial_j A_k - \partial_k A_j)^2 + O(\varepsilon^5)n^2 + O(\varepsilon^8)n^4\right),$$

and we can combine the two last error terms into just an $O(\varepsilon^5)n^4$ term. So now we will sum over all n. Defining the constants $K_1 = \sum_{n\in\mathbb{Z}}\exp\left(-\frac{2\pi^2 n^2}{g_0^2}\right)$ and $K_2 = \sum_{n\in\mathbb{Z}}\exp\left(-\frac{2\pi^2 n^2}{g_0^2}\right)\left(\frac{1}{2g_0^2} - \frac{2\pi^2 n^2}{g_0^4}\right)$, we get

$$\begin{aligned}\sum_{n\in\mathbb{Z}}\exp\left(-\frac{\beta}{2}(\theta_p - 2\pi n)^2\right) &= \sum_{n=0}^{\infty} 2\exp\left(-\frac{2\pi^2 n^2}{g_0^2}\right)\left(1 + \varepsilon^4\left(\frac{1}{2g_0^2} - \frac{2\pi^2 n^2}{g_0^4}\right)(\partial_j A_k - \partial_k A_j)^2 + O(\varepsilon^5)n^4\right)\\ &= K_1 + K_2\varepsilon^4(\partial_j A_k - \partial_k A_j)^2 + O(\varepsilon^5),\end{aligned}$$

where we use that the exponential decays very fast so the sum $\sum n^4\exp\left(-\frac{2\pi^2 n^2}{g_0^2}\right)$ is finite. Remember that all of this has been for a specific plaquette (with edges in the j, k direction). Let P' be the set of all plaqeuttes with θ_p nonzero or $\partial_j A_k(x) - \partial_k A_j(x)$ nonzero. Then our modified action sums over all plaquettes and looks like

(remember $\beta = \frac{1}{g_0^2}$)

$$\begin{aligned}\tilde{S}_\beta(\theta) &= -\sum_{p\in P'} \log\left(K_1 + K_2\varepsilon^4 \left(\partial_j A_k - \partial_k A_j\right)^2 + O(\varepsilon^5)\right) - \log(K_1) \\ &= -\sum_{p\in P'} \log\left(1 + \frac{K_2}{K_1}\varepsilon^4 \left(\partial_j A_k - \partial_k A_j\right)^2 + O(\varepsilon^5)\right)\end{aligned}$$

because remember that the subtracted off term just has $\theta_p = 0$ so the sum is exactly the definition of K_1. But now by the expansion of the log this is exactly (now writing out the sum over plaquettes more explicitly)

$$\sum_{x\in\mathbb{Z}^4}\sum_{j,k} \left(\partial_j A_k(x) - \partial_k A_j(X)\right)^2 + O(\varepsilon^5)|P'|.$$

But now $|P'| = O(\varepsilon^{-4})$ because A was assumed to be compactly supported, so the result follows and $\frac{K_2}{K_1}$ exactly results in the desired constant $\frac{1}{2g_0^2}\kappa$. (And we'll show why κ is nonzero later.)

2. $d = 2, 3$. Returning to the summand above, notice that for $n \geq 1$ this whole summand is bounded by $C_1 e^{-C_2\varepsilon^{d-4}n^2}$ (since the second exponential and cosine are basically bounded) and so actually we only need to consider the $n = 0$ term. That is,

$$\sum_{n\in\mathbb{Z}} \exp\left(-\frac{\beta}{2}(\theta_p - 2\pi n)^2\right) = \exp\left(-\frac{\beta}{2}\theta_p^2\right) + O\left(C_1 e^{-C_2\varepsilon^{d-4}}\right),$$

and this main term can be Taylor expanded as

$$\exp\left(-\frac{\beta}{2}\theta_p^2\right) = \exp\left(\frac{\varepsilon^d}{2g_0^2}(\partial_j A_k - \partial_k A_j + O(\varepsilon))^2\right) = 1 + \frac{\varepsilon^d}{2g_0^2}(\partial_j A_k - \partial_k A_j)^2 + O(\varepsilon^{d+1}).$$

On the other hand, we also have to subtract off the K_1 term, which looks like

$$\sum_{n\in\mathbb{Z}} \exp\left(-\frac{\beta}{2}(-2\pi n)^2\right) = \sum_{n\in\mathbb{Z}} \exp\left(-\frac{\pi^2\varepsilon^{d-4}}{g_0^2}n^2\right) = 1 + O\left(C_1 e^{-C_2\varepsilon^{d-4}}\right).$$

So all errors are pretty small and the rest of the proof goes through like in the $d = 4$ case.

3. $d \geq 5$. In this case all terms in the summand will contribute because we have a positive ε power in both exponentials. We have

$$\exp\left(\frac{\varepsilon^d}{2g_0^2}(\partial_j A_k(x) - \partial_k A_j(x) + O(\varepsilon))^2\right) = 1 + \frac{\varepsilon^d}{2g_0^2}(\partial_j A_k - \partial_k A_j)^2 + O(\varepsilon^{d+1})$$

as before. For the other terms we have

$$\cos\left(\frac{2\pi\varepsilon^{d-2}n}{g_0^2}(\partial_j A_k(x) - \partial_k A_j(x) + O(\varepsilon))\right) = 1 - \frac{2\pi^2\varepsilon^{2d-4}n^2}{g_0^4}(\partial_j A_k - \partial_k A_j)^2 + O(\varepsilon^{4d-8})n^4 + O(\varepsilon^{2d-3})n^2,$$

and now we have to multiply these two expansions together. But what saves us now is that $2d - 4 > d$, so we no longer have two different leading-order terms and so we won't get this nonzero κ factor. Unfortunately K_1, K_2 are now dependent on ε, and so the calculation still takes some work. Substituting in our expansions, the summand now looks like

$$\exp\left(-\frac{2\varepsilon^{d-4}\pi^2 n^2}{g_0^2}\right)\left(1 + \left(\frac{\varepsilon^d}{2g_0^2} - \frac{2\pi\varepsilon^{2d-4}n^2}{g_0^4}\right)(\partial_j A_k - \partial_k A_j)^2 + O(\varepsilon^{d+1}) + O(\varepsilon^{2d-3})n^2 + O(\varepsilon^{4d-8})n^4\right),$$

where the other terms are negligible because of our value of d. This time, we need to separate out K_2 into two

different terms because of the different orders of ε: define

$$K_1(\varepsilon) = \sum_{n\in\mathbb{Z}} \exp\left(-\frac{2\varepsilon^{d-4}\pi^2 n^2}{g_0^2}\right),$$

and similarly

$$K_2(\varepsilon) = \sum_{n\in\mathbb{Z}} n^2 \exp\left(-\frac{2\varepsilon^{d-4}\pi^2 n^2}{g_0^2}\right), \quad K_3(\varepsilon) = \sum_{n\in\mathbb{Z}} n^4 \exp\left(-\frac{2\varepsilon^{d-4}\pi^2 n^2}{g_0^2}\right).$$

We thus have

$$\sum_{n\in\mathbb{Z}} \exp\left(-\frac{\beta}{2}(\theta_p - 2\pi n)^2\right)$$
$$= K_1(\varepsilon) + \left(\frac{\varepsilon^d}{2g_0^2}K_1(\varepsilon) - \frac{2\pi^2\varepsilon^{2d-4}}{g_0^4}K_2(\varepsilon)\right)(\partial_j A_k - \partial_k A_j)^2 + O(\varepsilon^{d+1})K_1(\varepsilon) + O(\varepsilon^{2d-3})K_2(\varepsilon) + O(\varepsilon^{4d-8})K_3(\varepsilon).$$

So like before, taking log and subtracting off $\log(K_1(\varepsilon))$, and using that all of the exponents here are positive because $d \geq 5$, the result of that calculation ends up being

$$\log\left(1 + \left(\frac{\varepsilon^d}{2g_0^2} - \frac{2\pi^2\varepsilon^{2d-4}}{g_0^4}\frac{K_2(\varepsilon)}{K_1(\varepsilon)}\right)(\partial_j A_k - \partial_k A_j)^2 + O(\varepsilon^{d+1}) + O(\varepsilon^{2d-3})\frac{K_2(\varepsilon)}{K_1(\varepsilon)} + O(\varepsilon^{4d-8})\frac{K_3(\varepsilon)}{K_1(\varepsilon)}\right).$$

But now viewing our sums K_1, K_2, K_3 as Riemann sum approximations, we have (using $\delta = \varepsilon^{(d-4)/2}$ as the spacing)

$$\lim_{\varepsilon\to 0} \varepsilon^{(d-4)/2} K_1(\varepsilon) = \int_{-\infty}^{\infty} \exp\left(-\frac{2\pi x^2}{g_0^2}\right) dx,$$

and similarly (note the slightly different factor in the exponents of ε because we need to compensate for the powers of n)

$$\lim_{\varepsilon\to 0} \varepsilon^{3(d-4)/2} K_2(\varepsilon) = \int_{-\infty}^{\infty} x^2 \exp\left(-\frac{2\pi x^2}{g_0^2}\right) dx,$$
$$\lim_{\varepsilon\to 0} \varepsilon^{5(d-4)/2} K_3(\varepsilon) = \int_{-\infty}^{\infty} x^4 \exp\left(-\frac{2\pi x^2}{g_0^2}\right) dx.$$

Thus $\frac{K_2(\varepsilon)}{K_1(\varepsilon)} = O(\varepsilon^{-(d-4)})$ and $\frac{K_3(\varepsilon)}{K_1(\varepsilon)} = O(\varepsilon^{-(2d-8)})$, so as $\varepsilon \to 0$ the error terms are negligible and we can repeat the argument from $d = 4$ to take the appropriate limit, but the $\frac{K_2}{K_1}$ term actually still appears in the leading correction, and if we calculate the Gaussian integral it actually cancels out exactly with the other term $\frac{\varepsilon^d}{2g_0^2}$. So instead we get that the **action converges to a constant** for $d \geq 5$.

Furthermore, the quantity

$$\frac{\kappa(\tilde{g}_0)}{2\tilde{g}_0^2} = \frac{1}{2\tilde{g}_0^2} - \frac{2\pi^2}{\tilde{g}_0^4}\frac{\sum n^2 e^{-2\pi^2 n^2/\tilde{g}_0^2}}{\sum e^{-2\pi^2 n^2/\tilde{g}_0^2}}$$

goes to ∞ (resp. 0) as $\tilde{g}_0 \to 0$ (resp. $\tilde{g}_0 \to \infty$) by Riemann sum approximation, so we can always tune it so that we get the correct constant $\frac{1}{2g_0^2}$.

That concludes the proof of the scaling limit result, and the point of all of this is that we want to get the dual of 4D lattice gauge theory with the Villain action. For this, we define a function $\phi_\beta : [-\pi, \pi) \to \mathbb{R}$

$$\phi_\beta(x) = \sum_{n\in\mathbb{Z}} \exp\left(-\frac{\beta}{2}(x - 2\pi n)^2\right).$$

We can prove that $\phi_\beta(x)$ is C^∞ on the open interval $(-\pi, \pi)$ and in fact lifts to a smooth function on the unit circle

because

$$\phi_\beta^{(k)}(x) = \sum_{n\in\mathbb{Z}} \frac{d^k}{dx^k} \exp\left(-\frac{\beta}{2}(x-2\pi n)^2\right) = \sum_{n\in\mathbb{Z}} p_{k,\beta}(x-2\pi n)\exp\left(-\frac{\beta}{2}(x-2\pi n)^2\right)$$

for polynomials $p_{k,\beta}$. Taking the limit of this expression as $x \to \pi$, we get $\sum_{n\in\mathbb{Z}} p_{k,\beta}(\pi - 2\pi n)\exp\left(-\frac{\beta}{2}(n-2\pi n)^2\right)$, and then replacing $n \mapsto n+1$ shows that it is also the limit of the kth derivative as $x \to -\pi$. So all derivatives "wrap around" on both sides and so we lift to a smooth function as desired.

For such functions, we can derive facts about their Fourier transforms: for a function $f : [-\pi, \pi) \to \mathbb{R}$ with those properties above (C^∞ and agreement of limiting kth derivatives at both endpoints for all $k \geq 0$), we can define the **Fourier coefficient**

$$\hat{f}(n) = \int_{-\pi}^{\pi} f(x)e^{inx}dx$$

for all $n \in \mathbb{Z}$.

Lemma 68

The Fourier coefficient $\hat{f}(n)$ satisfies $|\hat{f}(n)| = o(|n|^{-\alpha})$ for any $\alpha > 0$, so the coefficients are in fact rapidly decaying.

Proof. By integration by parts,

$$\hat{f}(n) = \frac{f(x)e^{inx}dx}{in}\Big|_{-\pi}^{\pi} - \frac{1}{in}\int_{-\pi}^{\pi} f'(x)e^{inx}dx$$

and the boundary term is zero because $e^{in\pi} = e^{-in\pi}$ and $f(\pi) = f(-\pi)$. So $\hat{f}(n) = O(|n|^{-1})$ as $n \to \infty$, and then repeat the same integration by parts on f' and so on (an arbitrary number of times). □

What this means is that the sum

$$\frac{1}{2\pi}\sum_{n\in\mathbb{Z}} \hat{f}(n)e^{-inx}$$

is a well-defined smooth function on $[-\pi, \pi)$ (by the dominated convergence theorem).

Lemma 69

The function above is actually exactly $f(x)$; that is, we have the Fourier inversion formula

$$f(x) = \frac{1}{2\pi}\sum_{n\in\mathbb{Z}} \hat{f}(n)e^{-inx}.$$

Applying this to our function ϕ_β, we'll see why the Villain action is actually nice:

Lemma 70

We have

$$\hat{\phi}_\beta(n) = \sqrt{\frac{2\pi}{\beta}}e^{-n^2/(2\beta)}.$$

Proof. We wish to compute

$$\begin{aligned}\hat{\phi}_\beta(n) &= \int_{-\pi}^{\pi} \phi_\beta(x)e^{inx}dx \\ &= \sum_{m\in\mathbb{Z}} \int_{-\pi}^{\pi} \exp\left(-\frac{\beta}{2}(x-2\pi m)^2 + inx\right)dx\end{aligned}$$

and now we can replace inx with $in(x-2\pi m)$ because $e^{2\pi i}=1$. So now by a change of variables on the inner integral, we have

$$\sum_{m\in\mathbb{Z}}\int_{-2\pi m-\pi}^{-2\pi m+\pi}\exp\left(-\frac{\beta}{2}y^2+iny\right)dy$$

which is just an integral over the whole real line, and we can evaluate the resulting Gaussian integral. □

So we have a nice Fourier transform, and we'll use this next time to work towards the dual of 4D $U(1)$ theory.

14 November 10, 2025

We'll be continuing to think about $U(1)$ lattice gauge theory on some hypercube Λ of $\mathbb{Z}^4$ – specifically, we'll consider the Villain action S_β with free boundary conditions. Configurations are then of the form $\theta\in[-\pi,\pi)^E$ for E the set of positively oriented edges (with both endpoints in Λ); letting $\mathcal{P}$ be the set of plaquettes with all vertices in Λ, we're considering the probability density

$$\exp(-S_\beta(\theta))=\exp\left(-\left(-\sum_{p\in\mathcal{P}}\log\left(\sum_{n\in\mathbb{Z}}\exp\left(-\frac{\beta}{2}(\theta_p-2\pi n)^2\right)\right)\right)\right).$$

We can simplify this a bit because we have a sum of exponentials, which cancels out the log and thus yields

$$\exp(-S_\beta(\theta))=\prod_{p\in\mathcal{P}}\left(\sum_{n\in\mathbb{Z}}\exp\left(-\frac{\beta}{2}(\theta_p-2\pi n)^2\right)\right);$$

we will call the term inside the product $\phi_\beta(\theta_p)$. To calculate the normalizing constant, we must compute

$$\begin{aligned}Z=\int_{[-\pi,\pi)^E}e^{-S_\beta(\theta)}d\theta&=\int_{[-\pi,\pi)^E}\prod_{p\in\mathcal{P}}\phi_\beta(\theta_p)d\theta\\&=\frac{1}{(2\pi)^{|\mathcal{P}|}}\int_{[-\pi,\pi)^E}\prod_{p\in\mathcal{P}}\left(\sum_{n\in\mathbb{Z}}\hat{\phi}_\beta(n)e^{-in\theta_p}\right)d\theta\end{aligned}$$

where we've substituted in the expression for ϕ_β using the Fourier coefficients from last time. Now we can use the distributivity (noting that we only have a finite product of plaquettes but an infinite sum over n)

$$Z=\frac{1}{(2\pi)^{|\mathcal{P}|}}\int_{[-\pi,\pi)^E}\sum_{n\in\mathbb{Z}^{\mathcal{P}}}\left(\prod_{p\in\mathcal{P}}\hat{\phi}_\beta(n_p)e^{-in_p\theta_p}\right)d\theta,$$

where we've really used that the Fourier coefficients are rapidly decaying in n to allow this operation. Now again using the rapid decay of these coefficients and that $\mathcal{P}$ is a finite set, we have a countable sum so we can move it past the integral by the dominated convergence theorem:

$$\begin{aligned}Z&=\frac{1}{(2\pi)^{|\mathcal{P}|}}\sum_{n\in\mathbb{Z}^{\mathcal{P}}}\int_{[-\pi,\pi)^E}\left(\prod_{p\in\mathcal{P}}\hat{\phi}_\beta(n_p)e^{-in_p\theta_p}\right)d\theta\\&=\frac{1}{(2\pi)^{|\mathcal{P}|}}\sum_{n\in\mathbb{Z}^{\mathcal{P}}}\left(\prod_{p\in\mathcal{P}}\hat{\phi}_\beta(n_p)\right)\int_{[-\pi,\pi)^E}\exp\left(-\sum_{p\in\mathcal{P}}in_p\theta_p\right)d\theta.\end{aligned}$$

Now for any configuration $\theta \in [-\pi, \pi)^E$, we can define an $\mathbb{R}$-valued 1-form on $\mathbb{Z}^4$, which we will also call $\theta = \sum_{j=1}^4 \theta_j dx_j$, by

$$\theta_j(x) = \begin{cases} \theta(x, x+e_j) & \text{if } (x, x+e_j) \in E, \\ 0 & \text{otherwise.} \end{cases}$$

For a plaquette $p \in \mathcal{P}$ with vertices $x, x+e_j, x+e_j+e_k, x+e_k$ (for $1 \le j < k \le 4$), we have (remember $d\theta$ is a 2-form)

$$\begin{aligned} d\theta_{j,k}(x) &= \partial_j \theta_k(x) - \partial_k \theta_j(x) \\ &= \theta_k(x+e_j) - \theta_k(x) - \theta_j(x+e_k) + \theta_j(x) \\ &= \theta(x+e_j, x+e_j+e_k) - \theta(x, x+e_k) - \theta(x+e_k, x+e_k+e_j) + \theta(x, x+e_j) \\ &= \theta_p; \end{aligned}$$

that is, the 2-form $d\theta$ keeps track of the values θ_p (for plaquettes in $\mathcal{P}$; the others don't matter). Now for any $n \in \mathbb{Z}^p$, define the $\mathbb{R}$-valued 2-form on $\mathbb{Z}^4$ (also called n) as

$$n_{j,k}(x) = \begin{cases} n_p & \text{if } p \in \mathcal{P}, \\ 0 & \text{Otherwise.} \end{cases}$$

(remember that p is exactly determined by x, j, k). With these definitions, the sum we care about is really an inner product on 2-forms

$$\sum_p n_p \theta_p = (d\theta, n) = (\theta, \delta n)$$

(and that's why we don't have to worry about $d\theta$ outside of $\mathcal{P}$, since we've extended n to be zero outside). But now

$$(\theta, \delta_n) = \sum_{e \in E} \theta(e) \delta n(e)$$

because θ is zero for all other edges outside of E, and here remembering that $\delta n(e) = \delta_{n_j}(x)$ for $e = (x, x+e_j)$. Plugging this into our previous formula, we find that the integral now factors:

$$\begin{aligned} Z &= \frac{1}{(2\pi)^{|\mathcal{P}|}} \sum_{n \in \mathbb{Z}^{\mathcal{P}}} \left(\prod_{p \in \mathcal{P}} \hat{\phi}_\beta(n_p) \right) \int_{[-\pi,\pi)^E} \exp\left(-i \sum_{e \in E} \theta(e) \delta n(e) \right) d\theta \\ &= \frac{1}{(2\pi)^{|\mathcal{P}|}} \sum_{n \in \mathbb{Z}^{\mathcal{P}}} \left(\prod_{p \in \mathcal{P}} \hat{\phi}_\beta(n_p) \right) \prod_{e \in E} \int_{[-\pi,\pi)} \exp\left(-i\theta(e)\delta n(e) \right) d\theta. \end{aligned}$$

Now each $\delta n(e)$ is an integer, and we know that $\int_{-\pi}^{\pi} e^{-ikx} dx = 2\pi$ if $k = 0$ and 0 otherwise. Thus we really have

$$Z = (2\pi)^{|E| - |\mathcal{P}|} \sum_{\substack{n \in \mathbb{Z}^{\mathcal{P}} \\ \delta n(e) = 0 \ \forall e \in E}} \prod_{p \in \mathcal{P}} \hat{\phi}_\beta(n_p).$$

We now claim that if $\delta n(e) = 0$ for all $e \in E$, then actually $\delta n = 0$ everywhere as a 2-form. For any edge $e = (x, x+e_j)$, we have

$$\delta n_j(x) = -\sum_{i=1}^4 \overline{\partial}_i n_{ij}(x) = \sum_{i=1}^4 (n_{ij}(x - e_i) - n_{ij}(x))$$

by the formula for the coderivative; this is exactly the difference between the values on the plaquettes on the two sides of e (and summing over all directions except j). But if e is not in our hypercube, all plaquettes containing e will be outside the hypercube as well, so this whole sum will indeed be zero. Thus $\delta n_j(x)$ is always zero in all such cases, and

thus we can really simplify to

$$Z = (2\pi)^{|E|-|\mathcal{P}|} \sum_{\substack{n\in\mathbb{Z}^{\mathcal{P}} \\ \delta n=0}} \prod_{p\in\mathcal{P}} \hat{\phi}_\beta(n_p).$$

(Remember that $\delta n = 0$ means n can be thought of as a signed combination of plaquettes in a closed surface.) This whole argument is fairly general, so we can do this whole thing with other actions (like the Wilson action). But what's important now is that we have an explicit formula for those coefficients from Lemma 70 $\hat{\phi}_\beta(n) = \sqrt{\frac{2\pi}{\beta}} \exp\left(-\frac{n^2}{2\beta}\right)$, so

$$Z = (2\pi)^{|E|-|\mathcal{P}|} \left(\frac{2\pi}{\beta}\right)^{\frac{1}{2}|\mathcal{P}|} \sum_{\substack{n\in\mathbb{Z}^{\mathcal{P}} \\ \delta n=0}} \exp\left(-\frac{1}{2\beta}\sum_{p\in\mathcal{P}} n_p^2\right).$$

Call $\mathcal{N}$ the set of all $n \in \mathbb{Z}^{\mathcal{P}}$ with $\delta n = 0$, and let $\mathcal{A}$ be the set of all $\mathbb{Z}$-valued 1-forms α supported on the dual hypercube $*\Lambda$, such that $*d\alpha(p) = 0$ for all plaquettes outside $\mathcal{P}$. (This is almost true for anything in $*\Lambda$, except it might go slightly outside at the very boundary.) We can then define an equivalence relation

$$\alpha \sim \alpha' \quad \text{if} \quad d(\alpha - \alpha') = 0;$$

by discrete Poincaré this is equivalent to saying that there is some 0-form γ supported on $*\Lambda$ with $\alpha' = \alpha + d\gamma$. Letting $\mathcal{A}_0 = \mathcal{A}/\sim$ (whose elements we typically write as $[\alpha]$).

Lemma 71

The map $\xi : [\alpha] \mapsto *d\alpha$ is a bijection between $\mathcal{A}_0$ and $\mathcal{N}$. (This is well-defined because the result only depends on $d\alpha$.)

This lemma will allow us to transfer a sum over $\mathcal{N}$ to a sum over $\mathcal{A}_0$. And $\mathcal{A}$ is perhaps easier to think about than $\mathcal{N}$, where it's important to note that the dual of a plaquette is also a plaquette in four dimensions.

Proof. We've already observed that the map is well-defined. Note that for any $\alpha \in \mathcal{A}$, remembering that δ is some sign times $*d*$,

$$\delta(*d\alpha) = (-1)^{(\cdots)} *d**d\alpha,$$

and now because $**f = (-1)^{(\cdots)} f$ this simplifies to $(-1)^{(\cdots)} *dd\alpha = 0$ because $d^2 = 0$. Furthermore, by assumption, $*d\alpha(p) = 0$ for all $p \in \mathcal{P}$. Thus $*(d\alpha)$ is indeed always an element of $\mathcal{N}$, so ξ indeed maps $\mathcal{A}_0$ into $\mathcal{N}$.

Next, we prove this map is surjective. For any $n \in \mathcal{N}$, we have $\delta n = 0$ and so there is some $\mathbb{Z}$-valued 3-form m on $\mathbb{Z}^4$ such that $n = \delta m$ and m is supported on Λ. We can then define $\alpha = *(-m)$ (this may not be unique, but that's okay); this is a $\mathbb{Z}$-valued 1-form on the dual lattice $*\mathbb{Z}^4$ supported on $*\Lambda$. Then

$$n = \delta m = (-1)^{4(3+1)+1} *d**m = -*d**m = *d*(-m) = *d\alpha,$$

and $\alpha \in \mathcal{A}$ because all conditions are satisfied. Thus we can indeed reach any n with this map. (Note that we use crucially that we are in four dimensions here.)

Finally, we prove injectivity, which is easier. If $*d\alpha = *d\alpha'$, then $*d(\alpha - \alpha') = 0$, and therefore $d(\alpha - \alpha') = 0$ and $[\alpha] = [\alpha']$. □

So now for any $n \in \mathbb{Z}^{\mathcal{P}}$ with $\delta n = 0$, we can find some $\alpha \in \mathcal{A}$ such that $n = *d\alpha$. Then

$$\begin{aligned}\sum_{p\in\mathcal{P}} n_p^2 &= \sum_{x\in\mathbb{Z}^4} \sum_{1\le j<k\le 4} n_{jk}(x)^2 \\ &= \sum_{y\in *\mathbb{Z}^4} \sum_{1\le j<k\le 4} d\alpha_{jk}(y)^2 \\ &= (d\alpha, d\alpha).\end{aligned}$$

(Remember that we might get a sign coming from the permutation when we take $*$, but we square things here so there's nothing to worry about.) That yields the following normalizing constant:

Lemma 72

For the Villain action, we have

$$Z = (2\pi)^{|E|-|\mathcal{P}|} \left(\frac{2\pi}{\beta}\right)^{\frac{1}{2}|\mathcal{P}|} \sum_{[\alpha]\in\mathcal{A}_0} \exp\left(-\frac{1}{2\beta}(d\alpha, d\alpha)\right).$$

Thus we can consider the **dual model**, which is a probability measure on $\mathcal{A}_0$ (which notably is a **discrete** set) with probability mass at $[\alpha] \in \mathcal{A}_0$ proportional to $\exp\left(-\frac{1}{2\beta}(d\alpha, d\alpha)\right)$. That is, the model selects, for each edge in Λ^*, an integer, and this is kind of like the integer Gaussian free field but it's valued on edges rather than vertices. And the key now is that for β large this becomes close to the "comfortable region" where we are in the high-temperature regime and thus things are actually tractable!

Example 73

What we can do now is write Wilson loop expectations in terms of this new model. Let ℓ be a rectangular Wilson loop in Λ of the form which first goes in the $+e_j$, then $+e_k$, then $-e_j$, then $-e_k$ directions (for $j < k$), and suppose it has n vertices $x_0, x_1, \cdots, x_n$ and then returns to $x_{n+1} = x_0$. The Wilson loop expectation is then

$$W_\ell(\theta) = \prod_{j=0}^{n} e^{i\theta(x_j, x_{j+1})}$$

Lemma 74

We can alternatively express

$$W_\ell(\theta) = \exp\left(-i\sum_{p\in\Sigma}\theta_p\right),$$

where Σ is the set of plaquettes enclosed by the rectangle ℓ.

This crucially uses the abelian nature of the loop – this is basically a discrete version of Green's theorem.

Proof. Define a 2-form η on $\mathbb{Z}^4$ by $\eta(p) = 1$ if $p \in \Sigma$ and 0 otherwise. Then because $d\theta$ is exactly encoding the values of θ on plaquettes,

$$\sum_{p\in\Sigma}\theta_p = (d\theta, \eta) = (\theta, \delta\eta),$$

and now $\delta\eta$ (a signed linear combination of plaquettes next to each edge) is exactly -1 for all edges on the bottom and right and $+1$ on the top and left. So then $\exp\left(-i\sum_{p\in\Sigma}\theta_p\right) = W_\ell(\theta)$ as desired. □

So now we can go back to what we were doing before: we want to evaluate the expectation of W_ℓ, and

$$\langle W_\ell \rangle = \frac{Z_\ell}{Z}, \quad Z_\ell = \int_{[-\pi,\pi)^E} W_\ell(\theta) e^{-S_\beta}(\theta) d\theta.$$

Proceeding in exactly the same way as before, we can compute

$$Z_\ell = \frac{1}{(2\pi)^{|\mathcal{P}|}} \sum_{n \in \mathbb{Z}^{\mathcal{P}}} \left(\prod_{p \in \mathcal{P}} \hat{\phi}_\beta(n_p) \right) \int_{[-\pi,\pi)^E} \exp\left(-i \sum_{p \in \mathcal{P}} n_p \theta_p - i \sum_{p \in \Sigma} \theta_p \right) d\theta.$$

So in other words, we now have either n_p or $n_p + 1$ for each plaquette, depending on whether $p \in \Sigma$ or not. But then by a change of variables for only the plaquette variables in Σ, this can be rewritten as

$$\begin{aligned} Z_\ell &= \frac{1}{(2\pi)^{|\mathcal{P}|}} \sum_{n \in \mathbb{Z}^{\mathcal{P}}} \left(\prod_{p \in \mathcal{P} \setminus \Sigma} \hat{\phi}_\beta(n_p) \right) \left(\prod_{p \in \Sigma} \hat{\phi}_\beta(n_p - 1) \right) \int_{[-\pi,\pi)^E} \exp\left(-i \sum_{p \in \mathcal{P}} n_p \theta_p \right) d\theta \\ &= (2\pi)^{|E|-|\mathcal{P}|} \sum_{n \in \mathbb{Z}^{\mathcal{P}}, \delta n = 0} \left(\prod_{p \in \mathcal{P} \setminus \Sigma} \hat{\phi}_\beta(n_p) \right) \left(\prod_{p \in \Sigma} \hat{\phi}_\beta(n_p - 1) \right). \end{aligned}$$

Plugging in the value of $\hat{\phi}$, this simplifies to

$$(2\pi)^{|E|-|\mathcal{P}|} \left(\frac{2\pi}{\beta} \right)^{\frac{1}{2}|\mathcal{P}|} \sum_{n \in \mathbb{Z}^{\mathcal{P}}, \delta n = 0} \exp\left(-\frac{1}{2\beta} \sum_{p \in \mathcal{P} \setminus \Sigma} n_p^2 - \frac{1}{2\beta} \sum_{p \in \Sigma} (n_p - 1)^2 \right).$$

Defining the **disorder operator**

$$D_\ell(m) = \exp\left(\frac{1}{\beta} \sum_{p \in \Sigma} m_p - \frac{|\Sigma|}{2\beta} \right),$$

we can thus write this entire expression inside the n-sum as $D_\ell(m) \exp\left(-\frac{1}{2\beta} \sum_{p \in \mathcal{P}} m_p^2 \right)$. Transferring in the exact same way as we did for Z, we now get the following:

Lemma 75

We have

$$Z_\ell = (2\pi)^{|E|-|\mathcal{P}|} \left(\frac{2\pi}{\beta} \right)^{|\frac{1}{2}|\mathcal{P}} \sum_{[\alpha] \in \mathcal{A}_0} D_\ell(*d\alpha) \exp\left(-\frac{1}{2\beta} (d\alpha, d\alpha) \right).$$

So dividing Z_ℓ by Z makes all factors cancel out, and we get the following result:

Theorem 76

We have the equality of expectations

$$\langle W_\ell(\theta) \rangle = \langle D_\ell(*d\alpha) \rangle^*,$$

where θ is chosen from the lattice gauge theory and α is chosen from the dual theory.

15 November 12, 2025

We proved last time that for a rectangular loop ℓ in Λ under $U(1)$ lattice gauge theory with the Villain action and free boundary condition, we can write $\langle W_\ell \rangle$ in terms of a certain dual model defined on the state space $\mathcal{A}_0$ (an equivalence

quotient of $\mathbb{Z}$-valued 1-forms on $*\mathbb{Z}^4$ with $*d\alpha(p) = 0$ for all p not in the dual lattice, where $\alpha \sim \alpha'$ if $d\alpha = d\alpha'$). Specifically, with the probability measure μ_0 proportional to $\exp\left(-\frac{1}{2\beta}(d\alpha, d\alpha\right)$ on $\mathcal{A}_0$, we have $\langle W_\ell \rangle = \langle D_\ell \rangle^*$, where $D_\ell(\alpha) = \exp\left(\frac{1}{\beta}\sum_{p\in\Sigma} *d\alpha(p) - \frac{|\Sigma|}{2\beta}\right)$ was the disorder operator.

Remark 77. *Note that we didn't really use the fact that it's a rectangular loop – we could run this argument for any general loop as well if we choose a surface bounded by it, and we can double-check that the result D_ℓ doesn't depend on the choice of bounding surface and so we still get this relationship between the primal and dual model.*

Also, we can think of $*d\alpha(p) = 0$ as being "zero boundary conditions" on the dual lattice when we expand out by $\frac{1}{2}$ unit in all directions.

So now we want to compute $\langle D_\ell \rangle^*$, and it's very much like a Gaussian measure but it's supported only on integers. Thus, we'll define a new model to help us:

Example 78

In the **Gaussian $U(1)$ model**, let $\mathcal{U}$ be the space of all $\mathbb{R}$-valued (rather than $\mathbb{Z}$-valued) 1-forms on $*\Lambda$ such that $*d\alpha(p) = 0$ for all $p \notin P$. This is a finite-dimensional vector space over $\mathbb{R}$, and we similarly define a quotient relation $\alpha \sim \alpha'$ if $d(\alpha - \alpha') = 0$. To define a Gaussian space, first fix $\varepsilon > 0$ and define γ_ε to be the Gaussian measure on $\mathcal{U}$ with density proportional to $\exp\left(-\frac{1}{2\beta}(d\alpha, d\alpha) - \frac{\varepsilon}{2\beta}(\alpha, \alpha)\right)$ (so that we're actually integrable after quotienting).

Write $||\alpha|| = (\alpha, \alpha)^{1/2}$ for convenience, and let $\mathcal{H}$ be the space of all $\mathbb{R}$-valued 1-forms on $*\mathbb{Z}^4$, not necessarily with finite support, where $||\alpha|| < \infty$. Then $\mathcal{H}$ is a Hilbert space under this inner product, and $\mathcal{U}$ is a finite-dimensional closed subspace of $\mathcal{H}$. Let Π be the orthogonal projection onto $\mathcal{U}$.

Lemma 79

For any $\varepsilon > 0$, the restriction of the map $\Pi(\delta d + \varepsilon) : \mathcal{H} \to \mathcal{H}$ to $\mathcal{U}$ maps into $\mathcal{U}$, and it is self-adjoint, positive definite, and invertible. Furthermore, the covariance matrix of γ_ε is exactly β times the inverse of this map.

Proof. Clearly $\Pi(\delta d + \varepsilon)$ maps into $\mathcal{U}$ because Π projects onto $\mathcal{U}$. Let S denote this restriction $\Pi(\delta d + \varepsilon)|_{\mathcal{U}}$. Now for any $\alpha, \alpha' \in \mathcal{U}$ we have

$$(\alpha, S\alpha') = (\alpha, \Pi(\delta d + \varepsilon)\alpha') = (\alpha, (\delta d + \varepsilon)\alpha')$$

because α is already in $\mathcal{U}$ anyway so the projection doesn't change the inner product, and now by summation by parts this becomes $(d\alpha, d\alpha') + \varepsilon(\alpha, \alpha')$. So from this we see clearly that the map is self-adjoint and positive definite; for invertibility note that $S\alpha = 0$ implies $(\alpha, S\alpha) = 0$ and therefore we must have $(\alpha, \alpha) = 0$, hence $\alpha = 0$. So S is injective, and since we have a finite-dimensional space S is invertible. Finally from this formula, the probability density of γ_ε is proportional to $\exp\left(-\frac{1}{2\beta}(d\alpha, d\alpha) - \frac{\varepsilon}{2\beta}(\alpha, \alpha)\right) = \exp\left(-\frac{1}{2\beta}(\alpha, S\alpha)\right)$, as desired. □

Now let V_ε be this covariance $(\Pi(\delta d + \varepsilon)|_{\mathcal{U}})^{-1}$; since we have a Gaussian random vector we have the characteristic function for γ_ε

$$\int e^{i(\alpha,\eta)} d\gamma_\varepsilon(\alpha) = \exp\left(-\frac{\beta}{2}(\eta, V_\varepsilon \eta)\right)$$

for all η supported on $*\Lambda$ and we want to understand what happens to this as $\varepsilon \to 0$. (Eventually, we'll essentially restrict this probability measure to integer points, but we haven't done that yet.) For this, first define the subspaces

$$\mathcal{V}_0 = \{\alpha \in U : d\alpha = 0\}, \quad \mathcal{U}_0 = \mathcal{V}_0^\perp \text{ in } \mathcal{U}.$$

(so we can think of $\mathcal{U}_0$ as choosing one element of each equivalence class, now that we're in the continuous setting).

Lemma 80

The map $\Pi\delta d$ maps $\mathcal{U}$ into $\mathcal{U}_0$, and furthermore $\Pi\delta d|_{\mathcal{U}_0}$ is self-adjoint, positive definite, and invertible.

The point is that on $\mathcal{U}_0$ we can actually take $\varepsilon \to 0$ without problems.

Proof. For $\alpha \in \mathcal{U}$ and $\eta \in \mathcal{V}_0$, we have

$$(\eta, \Pi\delta d\alpha) = (\eta, \delta d\alpha) = (d\eta, d\alpha) = 0$$

by the same logic as before (removing the projection and then using summation by parts). Thus $\Pi\delta d\alpha$ is orthogonal to $\mathcal{V}_0$, meaning it maps into $\mathcal{U}_0$.

Next, let T be the restriction of $\Pi\delta d$ to $\mathcal{U}_0$. Proceeding like before, for $\alpha, \alpha' \in \mathcal{U}_0$ we have

$$(\alpha, T\alpha') = (\alpha', \Pi\delta d\alpha') = (\alpha, \delta d\alpha') = (d\alpha, d\alpha'),$$

which proves self-adjointness and positive semidefiniteness. Now if $(\alpha, T\alpha) = 0$ we must have $(d\alpha, d\alpha) = 0$, meaning that $\alpha \in \mathcal{V}_0$ but also $\alpha \in \mathcal{U}_0$ so we must have $\alpha = 0$ by orthogonality; thus we actually have positive definiteness and this proves invertibility as well. □

If we add on an ε to this, we also get the following:

Corollary 81

For all $\varepsilon > 0$, the operator $\Pi(\delta d + \varepsilon)$ maps $\mathcal{U}_0$ into $\mathcal{U}_0$, and the restriction to $\mathcal{U}_0$ is invertible. Thus the same claims also hold for V_ε.

Proof. Take $\alpha \in \mathcal{U}_0$. Then $\Pi(\delta d + \varepsilon)\alpha = \Pi\delta d\alpha + \varepsilon\alpha$ is the sum of two terms in $\mathcal{U}_0$, hence in $\mathcal{U}_0$. And $\Pi\delta(d\alpha + \varepsilon)$ is injective on the bigger space $\mathcal{U}$, so it is also injective on $\mathcal{U}_0$, hence invertible. Thus the inverse of this on $\mathcal{U}_0$, which is V_ε, also satisfies those same properties. □

Lemma 82

There exists some $c > 0$ such that $||d\alpha|| \geq c||\alpha||$ for all $\alpha \in \mathcal{U}_0$.

Proof. It suffices to prove the result when $||\alpha|| = 1$ by scaling. But by compactness this means it suffices to show that $d\alpha \neq 0$ for all $\alpha \in \mathcal{U}_0$ for all $||\alpha|| = 1$, which is true by orthogonality to $\mathcal{V}_0$. □

Corollary 83

There is some C such that for all $\varepsilon > 0$ and all $\alpha \in \mathcal{U}_0$, $||V_\varepsilon\alpha|| \leq C||\alpha||$.

Proof. Take any $\alpha \in \mathcal{U}_0$ and take c as above. For any $\varepsilon > 0$ we can define $\eta = V_\varepsilon\alpha$; note that $\eta \in \mathcal{U}_0$ and therefore

$$(\alpha, \eta) = (\alpha, V_\varepsilon\alpha) = (V_\varepsilon^{-1}\eta, \eta) = (\eta, V_\varepsilon^{-1}\eta) = (\eta, \Pi(\delta d + \varepsilon)\eta),$$

and this last expression is $||d\eta||^2 + \varepsilon||\eta||^2 \geq ||d\eta||^2 \geq c^2||\eta||^2$. Therefore by Cauchy-Schwarz,

$$||\alpha|| \cdot ||\eta|| \geq (\alpha, \eta) \geq c^2||\eta||^2$$

and so $||\eta|| \leq \frac{1}{c}||\alpha||$, as desired (we can take $C = \frac{1}{c^2}$). □

Lemma 84

Let V denote the inverse of the restriction $(\Pi\delta d|_{\mathcal{U}_0})$. Then for all $\eta \in \mathcal{U}$, we have the limit of the characteristic functions

$$\lim_{\varepsilon\to 0}\int e^{i(\alpha,\eta)}d\gamma_\varepsilon(\alpha) = \begin{cases}\exp\left(-\frac{\beta}{2}(\eta, V\eta)\right) & \text{if } \eta\in\mathcal{U}_0,\\ 0 & \text{if } \eta\in\mathcal{U}\setminus\mathcal{U}_0.\end{cases}$$

So we have a sequence of characteristic functions which converges to something which is not a characteristic function (since it's vanishing everywhere except a subspace) – instead it's the characteristic function on a quotient space. What we're basically doing as $\varepsilon\to 0$ is getting something uniform in the orthogonal direction (think about having density e^{-x^2} in $\mathbb{R}^2$).

Proof. Recall that $\int e^{i(\alpha,\eta)}d\gamma_\varepsilon(\alpha) = \exp\left(-\frac{\beta}{2}(\eta, V_\varepsilon\eta)\right)$. First note that for any $\eta\in\mathcal{V}_0$, we have $\Pi(\delta d+\varepsilon)\eta = \varepsilon\eta$ because $d\eta = 0$. Therefore applied to $\mathcal{V}_0$ we just have the simple expression

$$V_\varepsilon\eta = (\Pi(\delta d+\varepsilon))^{-1}\eta = \varepsilon^{-1}\eta.$$

So for any $\eta\in\mathcal{U}$ we can do the orthogonal decomposition $\eta = \eta_0+\eta_1$ with $\eta_0\in\mathcal{U}_0$ and $\eta_1\in\mathcal{V}_0$, and we get

$$\begin{aligned}(\eta, V_\varepsilon\eta) &= (\eta_0, V_\varepsilon\eta_0)+(\eta_0, V_\varepsilon\eta_1)+(\eta_1, V_\varepsilon\eta_0)+(\eta_1, V_\varepsilon\eta_1)\\ &= (\eta_0, V_\varepsilon\eta_0)+2(\eta_0, V_\varepsilon\eta_1)+(\eta_1, V_\varepsilon\eta_1)\\ &= (\eta_0, V_\varepsilon\eta_0)+2\varepsilon^{-1}(\eta_0,\eta_1)+\varepsilon^{-1}(\eta_1,\eta_1)\\ &\geq 0+0+\varepsilon^{-1}||\eta_1||^2.\end{aligned}$$

Thus this whole quantity tends to infinity if $\eta_1\neq 0$, so as long as we move out of $\mathcal{U}_0$ the thing appearing in the exponent of the characteristic function goes to infinity and thus we do indeed get zero.

On the other hand, if we're inside $\mathcal{U}_0$, write $\xi_\varepsilon = V_\varepsilon\eta$ (we know this is an element of $\mathcal{U}_0$) and so $\eta = V_\varepsilon^{-1}\xi_\varepsilon = \Pi(\delta d+\varepsilon)\xi_\varepsilon$. Now remembering that $||\xi_\varepsilon||$ is uniformly bounded by $C||\eta||$ by the previous lemma, if ξ is any subsequential limit of ξ_ε as $\varepsilon\downarrow 0$, then this expression shows $\eta = \Pi(\delta d+\varepsilon)\xi$, so $\xi = V\eta$. Therefore ξ_ε converges to V_η, meaning that $(\eta, V_\varepsilon\eta)$ converges to $(\eta, V\eta)$ and this is what we wanted. □

Remark 85. *To explain why we're taking this approach, we have some lattice of integer points, and we have a restriction of the Gaussian measure not to the lattice but to certain fibers. The issue is that if we just define a Gaussian measure here, we may not contain the projections of the integer points if we directly take the Gaussian measure on $\mathcal{U}_0$, since it's difficult to connect it to the integer model (where the αs are all integers, and we don't want to get non-integer valued forms if we just project to $\mathcal{U}_0$).*

We'll now use the Poisson summation formula to get the desired restriction to integers (in some sense).

Theorem 86 (Poisson summation formula)

Let $f:\mathbb{R}\to\mathbb{C}$ be a Schwartz function (meaning that the function and its derivatives are rapidly decaying faster than any polynomial). Then

$$\sum_{n\in\mathbb{Z}} f(n) = \sum_{n\in\mathbb{Z}}\hat{f}(2\pi n)$$

for $\hat{f}(t) = \int_{-\infty}^{\infty} e^{itx}f(x)dx$ the Fourier transform (which is also Schwartz).

Proof. Define the "periodization" function $g : [-\pi, \pi) \to \mathbb{C}$ via $g(x) = \sum_{n\in\mathbb{Z}} f\left(\frac{x}{2\pi} + n\right)$. We can check that

$$g^{(k)}(x) = \frac{1}{(2\pi)^k} \sum_{n\in\mathbb{Z}} f^{(k)}\left(\frac{x}{2\pi} + n\right),$$

meaning that

$$\begin{aligned}\lim_{x\to\pi} g^{(k)}(x) &= \frac{1}{(2\pi)^k} \sum_{n\in\mathbb{Z}} f^{(k)}\left(\frac{1}{2} + n\right) \\ &= \frac{1}{(2\pi)^k} \sum_{n\in\mathbb{Z}} f^{(k)}\left(-\frac{1}{2} + n\right) \\ &= \lim_{x\to-\pi} g^{(k)}(x),\end{aligned}$$

so the periodization has left and right boundary derivatives agreeing and thus we can write $g(x) = \frac{1}{2\pi}\sum_{n\in\mathbb{Z}} e^{-inx}\hat{g}(n)$. Plugging in $x = 0$, we see that

$$\begin{aligned}\sum_{n\in\mathbb{Z}} f(n) = g(0) &= \frac{1}{2\pi}\sum_{n\in\mathbb{Z}} \hat{g}(n) \\ &= \frac{1}{2\pi}\sum_{n\in\mathbb{Z}} \int_{-\infty}^{\infty} \sum_{m\in\mathbb{Z}} f\left(\frac{x}{2\pi} + m\right) e^{inx} dx \\ &= \frac{1}{2\pi}\sum_{n\in\mathbb{Z}}\sum_{m\in\mathbb{Z}} \int_{-\infty}^{\infty} f\left(\frac{x}{2\pi} + m\right) e^{inx} dx \\ &= \sum_{n\in\mathbb{Z}}\sum_{m\in\mathbb{Z}} \int_{-1/2}^{1/2} f(y+m) e^{2\pi i n y} dy.\end{aligned}$$

But now we can replace y with $y + m$ in the exponential, which doesn't change the value, and then do another change of variable to write the integral over the whole real line:

$$\sum_{n\in\mathbb{Z}}\sum_{m\in\mathbb{Z}} \int_{m-1/2}^{m+1/2} f(y) e^{2\pi i n y} dy = \sum_{n\in\mathbb{Z}} \hat{f}(2\pi n),$$

completing the proof. □

This is often used in comparing continuous to discrete models, and we'll see that soon!

16 November 17, 2025

Let's start with a quick review. We've been working with the $U(1)$ gauge theory on a four-dimensional hypercube Λ, and considering a rectangular loop ℓ we've previously shown that the Wilson loop expectation $\langle W_\ell \rangle$ agrees with the expectation $\langle D_\ell \rangle^*$ in the dual model, where the disorder operator $D_\ell = D_\ell(\alpha)$ takes in an **integer**-valued 1-form supported on the dual lattice $*\Lambda$ (meaning that it's nonzero only on edges that are in a slightly-enlarged hypercube around Λ) and where $*d\alpha(p) = 0$ for all plaquettes p inside Λ. (As a reminder, $d\alpha$ is a 2-form on the dual lattice, so $*d\alpha$ is a $(4-2)$-form on the original lattice.) And in fact we can write this set of allowed 1-forms $\mathcal{A}$ in another way:

Lemma 87

Say that $\Lambda = [a_1, b_1] \times \cdots \times [a_4, b_4]$. The dual lattice $*\Lambda$ is the shifted lattice $[a_1 + \frac{1}{2}, b_1 + \frac{1}{2}] \times \cdots \times [a_4 + \frac{1}{2}, b_4 + \frac{1}{2}]$, while Λ^* is instead the slightly enlarged lattice $[a_1 - \frac{1}{2}, b_1 + \frac{1}{2}] \times \cdots \times [a_4 - \frac{1}{2}, b_4 + \frac{1}{2}]$. Let E^* is the set of all edges with both endpoints in Λ^*, and its boundary ∂E^* is then the subset of edges with both endpoints having a neighbor outside Λ^*. Then $\mathcal{A}$ is the set of all $\mathbb{Z}$-valued 1-forms on $*\mathbb{Z}^4$ such that $\alpha(e) \neq 0$ only if $e \in E^* \setminus \partial E^*$.

(We won't prove this in class because it's rather technical, but it's in the notes.) This is a much nicer description than what we had before – it makes it clear that we're basically just working with the internal edges of a slightly enlarged cube. Recall also that we had the equivalence relation $\alpha \sim \alpha' \iff d(\alpha - \alpha') = 0$, and we defined the quotient space $\mathcal{A}_0 = \mathcal{A}/\sim$; we then defined the dual model as a probability measure μ_0 on $\mathcal{A}_0$ with probability mass function proportional to $\exp\left(-\frac{1}{2\beta}(d\alpha, d\alpha)\right)$. (The quotient space is countable, and we can confirm that we get a finite sum because of the derivations we've done earlier for the duality formula and the partition function.)

So this is almost like a Gaussian measure except it's forced to be integer-valued. Thus, let $\mathcal{U}$ be the set of $\mathbb{R}$-valued 1-forms on $*\mathbb{Z}^4$ supported on $*\Lambda$, such that $*d\alpha(p) = 0$ for all $p \notin P$; we can equivalently write this as $\mathbb{R}^{E^* \setminus \partial E^*}$. We can then define the Gaussian measures γ_ε on $\mathcal{U}$ with density proportional to $\exp\left(-\frac{1}{2\beta}(d\alpha, d\alpha) - \frac{\varepsilon}{2\beta}(\alpha, \alpha)\right)$ (this ε factor is there to make things integrable), and we also have the analogous object to $\mathcal{A}_0$, which is

$$\mathcal{U}_0 = \{\alpha \in \mathcal{U} : (\alpha, \eta) = 0 \text{ for all } \eta \in \mathcal{U} \text{ with } d\eta = 0\}$$

(This gets us one object from each equivalence class, because $\alpha' = \alpha + \eta$ for some η with $d\eta = 0$.) The analog to γ_ε is then μ_ε, which is the probability mass function on $\mathcal{A}$ (not $\mathcal{A}_0$) proportional to $\exp\left(-\frac{1}{2\beta}(d\alpha, d\alpha)\frac{\varepsilon}{2\beta}(\alpha, \alpha)\right)$, again the ε term makes things decay fast so that we have a valid probability measure. Let $\{z_q\}_{q=1}^\infty$ be positive real numbers such that

$$\sum_{q=1}^{\infty} \frac{2}{z_q} = 1;$$

define the partial sums $a_N = \sum_{q=1}^N \frac{2}{z_Q}$. Let $E_0^* = E^* \setminus \partial E^*$ be the internal edges, and let $\mathcal{Q} = \{1, 2, \cdots, \}^{E_0^*}$ and $\mathcal{Q}_N = \{1, 2, \cdots, N\}^{E_0^*}$ be assignments of **positive** integers to those edges Then for any $q \in \mathcal{Q}$, we can define

$$c(q) = \prod_{e \in E_0^*} \frac{2}{z_{q(e)}}.$$

Now for any $\alpha \in \mathcal{U} = \mathbb{R}^{E_0^*}$, we can define

$$\rho_N(\alpha) = \sum_{q \in \mathcal{Q}_N} c(q) \prod_{e \in E_0^*} \left(1 + a_N z_{q(e)} \cos\left(2\pi q(e)\alpha(e)\right)\right),$$

which is just some nonnegative function.

Lemma 88

With the notation above, and for any $\tau \in \mathcal{U}$, let $f(\alpha) = e^{2\pi i(\alpha, \tau)}$. Then for all $\varepsilon > 0$, we have that the discrete dual model with μ_ε can be written as

$$\langle f \rangle_\varepsilon^* = \lim_{N \to \infty} \frac{\int_U f(\alpha)\rho_N(\alpha) d\gamma_\varepsilon(\alpha)}{\int_U \rho_N(\alpha) d\gamma_\varepsilon(\alpha)}.$$

The point is that μ_ε cannot have a density with respect to γ_ε (which is a real Gaussian measure), but we can approximate it with something of the sort. So the quantity on the left side is an expectation with respect to the

integer measure, but the integrals on the right are with respect to real 1-forms. (This is kind of like approximating by an integer-valued Gaussian by something times the continuous Gaussian density, where that something gets spikier and spikier.)

Proof. We have that $f\mu_\varepsilon$ is a Schwarz function, so (letting Z_ε be the normalizing constant for μ_ε)

$$\begin{aligned}\langle f\rangle_\varepsilon^* &= \sum_{\alpha\in\mathbb{Z}^{E_0^*}} f(\alpha)\mu_\varepsilon(\alpha)\\ &= \frac{1}{Z_\varepsilon}\sum_{\alpha\in\mathbb{Z}^{E_0^*}} \widehat{f\mu_\varepsilon}(2\pi\alpha)\\ &= \frac{1}{Z_\varepsilon}\sum_{\alpha\in\mathbb{Z}^{E_0^*}}\int_{\mathcal{U}} e^{2\pi i(\alpha,\eta)} f(\eta)\exp\left(-\frac{1}{2\beta}(d\eta,d\eta)-\frac{\varepsilon}{2\beta}(\eta,\eta)\right)d\eta\\ &= \frac{1}{Z_\varepsilon}\sum_{\alpha\in\mathbb{Z}^{E_0^*}}\int_{\mathcal{U}} e^{2\pi i(\alpha+\tau,\eta)}\exp\left(-\frac{1}{2\beta}(d\eta,d\eta)-\frac{\varepsilon}{2\beta}(\eta,\eta)\right)d\eta\end{aligned}$$

by a higher-dimensional generalization of the Poisson summation formula we discussed last time. And now if we take f to be identically 1 (that is, take $\tau = 0$), the left-hand side is 1 and so we can solve for Z_ε: we thus find that

$$\langle f\rangle_\varepsilon^* = \frac{\sum_{\alpha\in\mathbb{Z}^{E_0^*}}\int_{\mathcal{U}} e^{2\pi i(\alpha+\tau,\eta)}d\gamma_\varepsilon(\eta)}{\sum_{\alpha\in\mathbb{Z}^{E_0^*}}\int_{\mathcal{U}} e^{2\pi i(\alpha,\eta)}d\gamma_\varepsilon(\eta)}.$$

But we have already made the evaluation that

$$\int_{\mathcal{U}} e^{2\pi i(\alpha,\eta)}d\gamma_\varepsilon(\eta) = \exp\left(-2\pi^2\beta(\alpha,V_\varepsilon\alpha)\right)$$

for some positive definite operator V_ε, so if we sum this left-hand side over all α we know that it converges and in fact

$$\sum_{\alpha\in\mathcal{A}}\int_{\mathcal{U}} e^{2\pi i(\alpha,\eta)}d\gamma_\varepsilon(\eta) = \lim_{N\to\infty}\sum_{\alpha\in\mathcal{A}_N}\int_{\mathcal{U}} e^{2\pi i(\alpha,\eta)}d\gamma_\varepsilon(\eta);$$

the same is true in the numerator. This means that in fact we have

$$\begin{aligned}\langle f\rangle_\varepsilon^* &= \lim_{N\to\infty}\frac{\sum_{\alpha\in\mathcal{A}_N}\int_{\mathcal{U}} e^{2\pi i(\alpha+\tau,\eta)}d\gamma_\varepsilon(\eta)}{\sum_{\alpha\in\mathcal{A}_N}\int_{\mathcal{U}} e^{2\pi i(\alpha,\eta)}d\gamma_\varepsilon(\eta)}\\ &= \lim_{N\to\infty}\frac{\int_{\mathcal{U}} e^{2\pi i(\tau,\eta)}\sum_{\alpha\in\mathcal{A}_N} e^{2\pi i(\alpha,\eta)}d\gamma_\varepsilon(\eta)}{\int_{\mathcal{U}}\sum_{\alpha\in\mathcal{A}_N} e^{2\pi i(\alpha,\eta)}d\gamma_\varepsilon(\eta)},\end{aligned}$$

, so now we just need to show that the sum we have in the integrand is proportional to the ρ_N we had before. But indeed

$$\begin{aligned}\sum_{\alpha\in\mathcal{A}_N} e^{2\pi i(\alpha,\eta)} &= \sum_{\alpha\in\mathcal{A}_N}\prod_{e\in E_0^*} e^{2\pi i\alpha(e)\eta(e)}\\ &= \prod_{e\in E_0^*}\left(\sum_{q(e)=-N}^{N} e^{2\pi i q(e)\eta(e)}\right)\\ &= \prod_{e\in E_0^*}\left(1+\sum_{q(e)=1}^{N}\cos(2\pi q(e)\eta(e))\right),\end{aligned}$$

where the second line is by the distributive property because $\mathcal{A}_N = \{-N, \cdots, N\}^{E_0^*}$. And finally we can write

$$1 + \sum_{q(e)=1}^{N} \cos(2\pi q(e)\eta(e)) = \sum_{q(e)=1}^{N} \frac{2}{z_{q(e)}} \left(\frac{1}{a_N} + z_{q(e)} \cos\left(2\pi q(e)\eta(e)\right) \right),$$

substituting this in and bringing out the a_Ns yields the result. □

We'll now start working some more with this ρ_N to understand what we can say about it – it's not supposed to converge to anything as $N \to \infty$, again because μ_ε shouldn't have a density with respect to γ_ε.

Definition 89

For any $\alpha \in \mathcal{A} = \mathbb{Z}^{E_0^*}$, define the **edge-support** of α, denoted $E(\alpha)$, to be the set of all edges $e \in E_0^*$ where $\alpha(e) \neq 0$. For two edges $e, e' \in \mathbb{Z}^{E_0^*}$, let $\mathrm{dist}(e, e')$ be the minimum of the Euclidean distances between the vertices of the edges (in particular this can be zero for different edges if there is a common vertex). Similarly for two subsets of vertices $E, E' \subseteq E_0^*$, we can define $\mathrm{dist}(E, E')$ to be the minimum of $\mathrm{dist}(e, e')$ over all $e \in E$ and $e' \in E'$.

We say that a set of 1-forms $\mathcal{E} \subset \mathcal{A}$ is an **ensemble** if for all distinct $\alpha, \alpha' \in \mathcal{E}$, the edge-supports $E(\alpha)$ and $E(\alpha')$ are disjoint, and we say that $\mathcal{E}$ is a ***k*-ensemble** if for all distinct α, α', we have $\mathrm{dist}(E(\alpha), E(\alpha')) \geq 2^{k/2}$.

We will only need to take $k = 1$ in our case (meaning we need vertices from $E(\alpha)$ and $E(\alpha')$ to not be overlapping and also not be side-by-side), but in more generality we may need higher k.

Lemma 90

Take any $N \geq 1$. Given any $q \in \mathcal{Q}_N = \{1, \cdots, N\}^{E_0^*}$, there exists a finite set $\Gamma(q)$ which indexes strictly positive coefficients $\{c_\gamma\}_{\gamma \in \Gamma(q)}$, 1-ensembles $\{\mathcal{E}_\gamma\}_{\gamma \in \Gamma(q)}$ (possibly the same for different γ), and strictly positive coefficients $\{K_\gamma(\eta)\}_{\gamma \in \Gamma(q), \eta \in \mathcal{E}_\gamma}$, so that

$$\prod_{e \in E_0^*} \left(1 + a_N z_{q(e)} \cos(2\pi q(e)\alpha(e))\right) = \sum_{\gamma \in \Gamma(q)} c_\gamma \prod_{\eta \in \mathcal{E}_\gamma} \left(1 + K_\gamma(\eta) \cos(2\pi(\alpha, \eta))\right),$$

and furthermore we have that $K_\gamma(\eta) \leq 3^{N_1(\eta)} \prod_{e \in E(\eta)} z_{q(e)}$ for $N_1(\eta)$ the number of edges $e \in E_0^*$ of distance at most 1 from $E(\eta)$ (so this is the size of the support, expanded by a little bit).

In particular, taking a combination over all q will get us the ρ_N we had above. Notice that we have a product over all e in the left-hand side, and different edges may interact (since they have common vertices or plaquettes). But on the right-hand side, the product is a product over 1-forms with disjoint support, and because the supports are even separated by a bit more there is no interaction between different ηs.

Start of proof. The key fact is the following identity: for real numbers $K_1, K_2, \theta_1, \theta_2 \in \mathbb{R}$, we have (by the sum-to-product formula)

$$\begin{aligned}(1 + K_1 \cos\theta_1)(1 + K_2 \cos\theta_2) \\ = \frac{1}{3}(1 + 3K_1 \cos(\theta_1)) + \frac{1}{3}(1 + 3K_2 \cos(\theta_2)) + \frac{1}{6}(1 + 3K_1 K_2 \cos(\theta_1 - \theta_2)) + \frac{1}{6}(1 + 3K_1 K_2 \cos(\theta_1 + \theta_2)).\end{aligned}$$

Since our object of interest is

$$I = \prod_{e \in E_0^*} \left(1 + a_N z_{q(e)} \cos\left(2\pi q(e)\alpha(e)\right)\right),$$

we can write it in the form

$$\sum_{\gamma\in\Gamma} c_\gamma \prod_{\eta\in\mathcal{E}_\gamma}\Big(1+K_\gamma(\eta)\cos(2\pi(\alpha,\eta)\Big),$$

where Γ is just a single element $\{0\}$, we have the constant $c_0=1$, the ensemble is the set $\mathcal{E}_0=\{\eta_e\}_{e\in E_0^*}$ where $\eta_e(e')=q(e)$ if $e=e'$ and 0 otherwise, and $K_0(\gamma_e)=a_N z_{q(e)}$. Now $\mathcal{E}_0$ is an ensemble, but it is not a 1-ensemble; thus, we will gradually make it into one by repeatedly using the identity we wrote down (and we keep track of things to make sure the constants don't blow up too much).

This is an inductive construction: in the first step, take any distinct $a,b\in E_0^*$ with $d(a,b)<\sqrt{2}$ (that is, $d(a,b)\le 1$). We can then observe that

$$\begin{aligned}
&\Big(1+a_N z_{q(a)}\cos(2\pi q(a)\alpha(a))\Big)\Big(1+a_N z_{q(b)}\cos(2\pi q(b)\alpha(b))\Big)\\
&\quad=\frac{1}{3}\Big(1+3a_N z_{q(a)}\cos(2\pi q(a)\alpha(a))\Big)+\frac{1}{3}\Big(1+3a_N z_{q(b)}\cos(2\pi q(b)\alpha(b))\Big)\\
&\qquad+\frac{1}{6}\Big(1+3a_N^2 z_{q(a)}z_{q(b)}\cos(2\pi(q(a)\alpha(a)-q(b)\alpha(b)))\Big)\\
&\qquad+\frac{1}{6}\Big(1+3a_N^2 z_{q(a)}z_{q(b)}\cos(2\pi(q(a)\alpha(a)+q(b)\alpha(b)))\Big).
\end{aligned}$$

So we can now take this four-term expression and substitute it back into our product; we'll now have four different ensembles. Indeed, we will have $\Gamma'=\{1,2,3,4\}$, the first ensemble $\mathcal{E}_1$ will consist of $\{\eta_e\}_{e\neq b}$ (so we threw out one of the one-forms of $\mathcal{E}_0$), $\mathcal{E}_2$ will consist of $\{\eta_e\}_{e\neq a}$, $\mathcal{E}_3=\{\eta_e\}_{e\neq a,b}\cup\{\xi_1\}$ for the form $\xi_1=\eta_a-\eta_b$, and $\mathcal{E}_4=\{\eta_e\}_{e\neq a,b}\cup\{\xi_2\}$ for the form $\xi_1=\eta_a+\eta_b$. None of these ensembles witness the problem between a and b, and everything is still an ensemble.

But now we just keep repeating this: whenever there is a pair of edges a,b of this form, we do this to remove the problem, and so our ensembles will grow like a tree downward. Furthermore, notice that $|\mathcal{E}_i|<|\mathcal{E}|$ for all $i=1,2,3,4$ in the above argument, since the size of the ensemble decreases by 1 in all cases. The inductive hypothesis is then as follows: if we've arrived at some Γ and collection $\{\mathcal{E}_\gamma\}_{\gamma\in\Gamma}$, and at least one $\mathcal{E}_\gamma$ is not a 1-ensemble, then there exist some distinct 1-forms $\eta,\eta'\in\mathcal{E}_\gamma$ with $\operatorname{dist}(E(\eta),E(\eta'))<\sqrt{2}$. This means there is some $a\in E(\eta)$ and $b\in E(\eta')$ with $\operatorname{dist}(a,b)<\sqrt{2}$; we can then break up $(1+K_\gamma(\eta)\cos(2\pi(\alpha,\eta)))(1+K_\gamma(\eta')\cos(2\pi(\alpha,\eta')))$ in the same way and substitute in; this replaces $\mathcal{E}_\gamma$ with $\mathcal{E}_{\gamma_1},\mathcal{E}_{\gamma_2},\mathcal{E}_{\gamma_3},\mathcal{E}_{\gamma_4}$ where $\mathcal{E}_{\gamma_1}=\mathcal{E}\setminus\{\eta'\}$, $\mathcal{E}_{\gamma_2}=\mathcal{E}\setminus\{\eta\}$, $\mathcal{E}_{\gamma_3}=(\mathcal{E}\setminus\{\eta,\eta'\})\cup\{\eta-\eta'\}$, and $\mathcal{E}_{\gamma_4}=(\mathcal{E}\setminus\{\eta,\eta'\})\cup\{\eta+\eta'\}$. And at each such step, the ensemble sizes decrease, so this process must stop at some point if we have reached a collection of 1-ensembles. We'll show next time how to get the bound on the constants $K_\gamma(\eta)$! □

17 November 19, 2025

We'll start today by explaining "how to send ε to zero." Recall that we had defined $\mathcal{U}=\mathbb{R}^{E_0^*}$ and decomposed it into the orthogonal spaces $\mathcal{V}_0=\{\alpha\in\mathcal{U}\}$ and $\mathcal{U}_0=\mathcal{V}_0^\perp$; we also had a discrete analog $\mathcal{A}=\mathbb{Z}^{E_0^*}$ and $\mathcal{A}_0=\mathcal{A}/\sim$ where $\alpha\sim\alpha'$ if $d\alpha=d\alpha'$. **We'll now choose a specific representative from each equivalence class**, letting $\mathcal{A}_0$ be one minimum norm element from each class (all norms are integers here).

Remember that we had a probability measure μ_0 on $\mathcal{A}_0$ with probability mass proportional to $\exp\left(-\frac{1}{2\beta}||d\alpha||^2\right)$, and we also defined the probability measure on $\mathcal{A}$ with probability mass function $\exp\left(-\frac{1}{2\beta}||d\alpha||^2-\frac{\varepsilon}{2\beta}||\alpha||^2\right)$. We then also have the analogs of these in the continuous space: γ_ε is the measure on $\mathcal{U}$ with probability density proportional to $\exp\left(-\frac{1}{2\beta}||d\alpha||^2\frac{\varepsilon}{2\beta}||\alpha||^2\right)$, and so analogously we now also define γ_0 to be the measure on the subspace $\mathcal{U}_0$ with probability density proportional to $\exp\left(-\frac{1}{2\beta}||d\alpha||^2\right)$.

The goal will be to relate the μs to the γs fully rigorously through a rather complicated lemma. Recall that we have expectations $\langle f\rangle^* = \int_{\mathcal{A}_0} f\, d\mu_0$.

Proposition 91

Let η_0 be a random 1-form with law γ_0. Let $\mathcal{B}_0 = \mathcal{A}_0 \cap \mathcal{U}_0$ (note that $\mathcal{A}_0$ is the minimum norm element among integer-valued 1-forms; $\mathcal{A}_0$ would actually be $\mathcal{U}_0$ if we were minimizing over real-valued 1-forms instead). Take any $\tau \in \mathcal{U}_0$ and let $f(\alpha) = e^{(\tau,\alpha)}$ for $\alpha \in \mathcal{U}$ (this is the kind of function we want to take the expectation of, since we care about this disorder operator $D_\ell(\alpha)$ which depends only on $d\alpha$, and using the adjoint rule $\delta\tau$ is always in $\mathcal{U}_0$). Then

$$\langle f\rangle^* = \frac{\sum_{\alpha_0\in\mathcal{B}_0} \mathbb{E}\left[f(\eta_0)e^{2\pi i(\alpha_0,\eta_0)}\right]}{\sum_{\alpha_0\in\mathcal{B}_0} \mathbb{E}\left[e^{2\pi i(\alpha_0,\eta_0)}\right]}.$$

This means that heuristically we can treat $\frac{d\mu_0}{d\gamma_0} = \sum_{\alpha_0)} e^{2\pi i(\alpha_0,\eta)}$ as the density of the discrete model with respect to the continuous model at any point $\eta \in \mathcal{U}_0$. (So at integer points, (α_0, η) is always an integer, so this is like a sum of delta masses.) We can think of $\mathcal{A}_0$ as the "closest harmonic approximation," but in that framework it's hard to see why integer-valued forms would be sent to integer-valued forms.

Lemma 92

Sample the random element $\eta_\varepsilon \sim \gamma_\varepsilon$, and let $\eta_{0,\varepsilon}, \eta_{1,\varepsilon}$ be its projections onto the orthogonal subspaces $\mathcal{U}_0$ and $\mathcal{V}_0$. Then $\eta_{0,\varepsilon}$ and $\eta_{1,\varepsilon}$ are independent 1-forms with covariance matrices βV_ε (recall that $V_\varepsilon = (\Pi(\delta d + \varepsilon))^{-1}$) and $\beta\varepsilon^{-1} I$ on $\mathcal{U}_0$ and $\mathcal{V}_0$. Then as $\varepsilon \to 0$, we have $\eta_{0,\varepsilon}$ converging in distribution to $\eta_0 \sim \gamma_0$.

Proof. Since we have Gaussian forms, we just need to calculate variances. Fix $(\alpha, \alpha') \in \mathcal{U}$; let Π_0 be the projection to $\mathcal{U}_0$, and define $\alpha_0 = \Pi_0\alpha$, $\alpha_1 = (I - \Pi_0)\alpha$, and similarly define α_0', α_1'. Now we compute the linear combinations (using that Π_0 is self-adjoint)

$$\begin{aligned}\mathbb{E}\left[((\alpha, \eta_{0,\varepsilon}) + (\alpha', \eta_{1,\varepsilon}))^2\right] &= \mathbb{E}\left[((\alpha, \Pi_0\eta_\varepsilon) + (\alpha', (I - \Pi_0)\eta_\varepsilon))^2\right] \\ &= \mathbb{E}\left[((\alpha_0, \eta_\varepsilon) + (\alpha_1', \eta_\varepsilon))^2\right] \\ &= \mathbb{E}\left[(\alpha_0 + \alpha_1', \eta_\varepsilon)^2\right]\end{aligned}$$

Now $\alpha_1' \in \mathcal{V}_0$, so $d\alpha_1' = 0$, and therefore $\Pi(\delta d + \varepsilon)\alpha_1' = \varepsilon\alpha_1'$. Thus applying the inverse yields $V_\varepsilon\alpha_1' = \varepsilon^{-1}\alpha_1'$; using this and orthogonality simplifies the expectation to

$$((\alpha_0 + \alpha_1'), V_\varepsilon(\alpha_0 + \alpha_1')) = (\alpha_0, V_\varepsilon\alpha_0) + \varepsilon^{-1}||\alpha_1'||^2.$$

And now for any Gaussian vector supported on a subspace of $\mathbb{R}^n$, $(\alpha, x)^2$ for any $\alpha \in \mathbb{R}^n$ will be exactly the quadratic form $\alpha^T V\alpha$ for V the covariance on that subspace. So this covariance formula gives us exactly what we are claiming. □

Lemma 93

There is some constant $c_0 > 0$, possibly dependent on Λ, such that $||d\alpha|| \geq c_0||\alpha||$ for all $\alpha \in \mathcal{A}_0$.

We proved a similar result for $\mathcal{U}_0$, but $\mathcal{A}_0$ is a discrete set and we're taking specifically the minimum-norm elements so things are a bit more complicated.

Proof. For $\alpha \in \mathcal{A}_0$, consider its projection $\alpha_1 = \Pi_0 \alpha$, which is now real-valued, and $\alpha_2 = (I - \Pi_0)\alpha$. Since $d\alpha_2 = 0$ and α_2 is supported on $*\Lambda$ (this is true for any form supported only on the interior edges), there exists a 1-form κ supported on $*\Lambda$ with $\alpha_2 = d\kappa$ by the Poincaré lemma. We thus get an integer approximation $\eta(e) = \lfloor \kappa(e) \rfloor$; let $\theta = \alpha - d\eta$. We know that $\alpha - d\kappa$ would be α_1, and θ is an integer-valued approximation of that quantity. We know that $\theta \in \mathcal{A}$ and $\theta \sim \alpha$, so $||\theta|| \geq ||\alpha||$ since we took the minimum norm element. Now

$$\begin{aligned} ||\theta - \alpha_1|| &= ||\alpha - d\eta - \alpha_1|| \\ &= ||\alpha_2 - d\eta|| \\ &= ||d\kappa - d\eta|| \\ &\leq C \end{aligned}$$

for some constant C because η and κ differ by at most 1 on each edge. So $||\theta|| \leq ||\alpha_1|| + C$, and since $\alpha_1 \in \mathcal{U}_0$ we have $||\alpha_1|| \leq C_1 ||d\alpha_1||$ because we've already proven the analogous result on the full subspace. And $||d\alpha|| = ||d\alpha_1||$ because $\alpha \sim \alpha_1$, so putting this together we get

$$||\alpha|| \leq ||\theta|| \leq C_1 ||d\alpha|| + C,$$

and furthermore $||\alpha||^2$ and $||d\alpha||^2$ are always nonnegative integers, and if $||d\alpha|| = 0$ then $||\alpha|| = 0$ since α must be specifically 0 (the minimum-norm element in that equivalence class). So in fact we have the bound $||\alpha|| \leq (C_1 + C)||d_\alpha||$, as desired. □

Lemma 94

Let $f : \mathcal{A} \to \mathbb{C}$ be a function with subexponential growth, meaning that $|f(\alpha)| \leq C_1 e^{C_2 ||\alpha||}$, which is constant on equivalence classes. Then

$$\langle f \rangle^* = \lim_{\varepsilon \to 0} \langle f \rangle_\varepsilon^*.$$

The plan is to use this result and then relate $\langle f \rangle^*$ to some result with γ_ε, and then take the limit of that as $\varepsilon \to 0$.

Proof. In this proof, $C, C_0, C_1, \cdots$ will be arbitrary positive constants depending only on β and Λ. For any $\alpha \in \mathcal{A}_0$, we have $[\alpha] = \alpha + \mathcal{B}$ for $\mathcal{B} = \{\eta \in \mathcal{A} : d\eta = 0\}$. For any $\eta \in \mathcal{B}$, define $\alpha' = \alpha + \eta$ and notice that (because α, α' are equivalent)

$$||d\alpha'||^2 + \varepsilon ||\alpha'||^2 = ||d\alpha||^2 + \varepsilon ||\alpha||^2 + \varepsilon ||\eta||^2 + 2\varepsilon(\alpha, \eta).$$

The sum in the numerator we are interested in looks like

$$\begin{aligned} \sum_{\alpha \in \mathcal{A}} f(\alpha) \exp\left(-\frac{1}{2\beta} ||d\alpha||^2 + \varepsilon ||\alpha||^2 \right) &= \sum_{\alpha \in \mathcal{A}_0} \sum_{\eta \in \mathcal{B}} f(\alpha + \eta) \exp\left(-\frac{1}{2\beta} (||d\alpha||^2 + \varepsilon ||\alpha||^2 + \varepsilon ||\eta||^2 + 2\varepsilon(\alpha, \eta) \right) \\ &= \sum_{\alpha \in \mathcal{A}_0} \sum_{\eta \in \mathcal{B}} f(\alpha) \exp\left(-\frac{1}{2\beta} (||d\alpha||^2 + \varepsilon ||\alpha||^2 + \varepsilon ||\eta||^2 + 2\varepsilon(\alpha, \eta) \right). \end{aligned}$$

Remember Λ_0^* is the set of interior vertices of Λ^*. Define a function which separates out the ε factor

$$g_\varepsilon(\alpha) = \varepsilon^{\frac{1}{2}|\lambda_0^*|} \sum_{\eta \in \mathcal{B}} \exp\left(-\frac{\varepsilon}{2\beta} (||\alpha||^2 + ||\eta||^2 + (2\alpha, \eta) \right),$$

so that the expectations in the discrete model are related by

$$\langle f \rangle_\varepsilon^* = \frac{\langle f g_\varepsilon \rangle^*}{\langle g_\varepsilon \rangle^*},$$

and we want to show that as $\varepsilon \to 0$ this just approaches $\langle f \rangle^*$ on the right-hand side.

Now **we claim that** $\eta \in \mathcal{B}$ if and only if $\eta = d\theta$ for some unique 0-form θ supported in the interior vertices Λ_0^*. Indeed, for any such θ we know that the 1-form $d\theta$ vanishes outside E_0^* (we only get nonzero edges if one of the adjacent vertices is in the interior), meaning that indeed $d\theta \in \mathcal{B}$ because $dd\theta = 0$. On the other hand, if $\eta \in \mathcal{B}$, then there exists some $\mathbb{Z}$-valued 0-form θ with $\eta = d\theta$. And $\eta(e) = 0$ for all $e \notin E^*$ and also for $e \in \partial E^*$, so θ must be constant on the boundary and outside and finitely supported, hence zero on all of those points. This proves the claim; uniqueness is because $\eta = d\theta = d\theta' \implies d(\theta - \theta') = 0$, so $\theta = \theta'$ for 0-forms because we're finitely supporetd.

So therefore we can write $g_\varepsilon(\alpha)$ by replacing η with $d\theta$:

$$g_\varepsilon(\alpha) = \varepsilon^{\frac{1}{2}|\lambda_0^*|} \sum_{\theta \in \mathbb{Z}^{\Lambda_0^*}} \exp\left(-\frac{\varepsilon}{2\beta}(||\alpha||^2 + ||d\theta||^2 + 2(\alpha, d\theta)\right).$$

We can now further simplify this by taking a rescaled lattice

$$= \varepsilon^{\frac{1}{2}|\lambda_0^*|} \sum_{\theta \in \sqrt{\varepsilon}\mathbb{Z}^{\Lambda_0^*}} \exp\left(-\frac{1}{2\beta}(\varepsilon||\alpha||^2 + ||d\theta||^2 + 2\sqrt{\varepsilon}(\alpha, d\theta)\right).$$

But as $\varepsilon \to 0$, **we claim** this actually converges to an integral where there is no involvement of α by Riemann sum approximation (since those terms go away). Thus $g_\varepsilon(\alpha)$ converges to a constant and thus we have $\langle f \rangle_\varepsilon^* \to \langle f \rangle^*$. What's left is just to verify that the conditions of the dominated convergence theorem actually apply, and that's why we proved all of the various technical details we did before.

Indeed, define the summand as $h_\varepsilon(\alpha, \theta)$ and now extend it to all real forms by saying that for $\theta \in \sqrt{\varepsilon}\mathbb{Z}^{\Lambda_0^*}$, define B_{θ_0} to be the cube with center θ_0 and sidelength $\sqrt{\varepsilon}$ (this partitions $\mathbb{R}^{\Lambda_0^*}$). We can then define the piecewise constant function $h_\varepsilon(\alpha, \theta) = h_\varepsilon(\alpha, \theta_0)$ for any $\theta \in B_{\theta_0}$. The sum of interest is then $\int_{\mathbb{R}^{\Lambda_0}} h_\varepsilon(\alpha, \theta) d\theta$. We want to say that we can apply Fubini's theorem to get

$$\langle f g_\varepsilon \rangle^* = \int_{\mathcal{A}_0} f(\alpha) g_\varepsilon(\alpha) d\mu_0(\alpha) = \int_{\mathcal{A}_0} \int_{\mathbb{R}^{\Lambda_0^*}} f(\alpha) h_\varepsilon(\alpha, \theta) d\theta d\mu_0(\alpha).$$

(On its own, h wouldn't cause any problems because it is always nonnegative, but f is complex-valued.) Indeed, for any fixed $\varepsilon > 0$ and $\theta_0 \in \sqrt{\varepsilon}\mathbb{Z}^{\Lambda_0^*}$, we can take any $\theta \in B_{\theta_0}$ in its surrounding cube so that

$$\begin{aligned} h_\varepsilon(\alpha, \theta) = h_\varepsilon(\alpha, \theta_0) &= \exp\left(-\frac{1}{2\beta}\left(\varepsilon||\alpha||^2 + ||d\theta_0||^2 + 2\sqrt{\varepsilon}(\alpha, d\theta_0)\right)\right) \\ &\leq \exp\left(-\frac{1}{2\beta}\left(0 + C||\theta_0||^2 + 2\sqrt{\varepsilon}(\alpha, d\theta_0)\right)\right) \\ &\leq \exp\left(-\frac{C}{2\beta}||\theta_0||^2 - \frac{\sqrt{\varepsilon}}{\beta}(\delta\alpha, \theta_0)\right), \end{aligned}$$

where in the second line we used that by compactness, there is some $C > 0$ such that $||d\theta||^2 \geq c||\theta||^2$ (since $d\theta = 0 \implies \theta = 0$). And now $||\theta - \theta_0|| \leq C\sqrt{\varepsilon}$ implies $||\theta||^2 \leq 2||\theta - \theta_0||^2 + 2||\theta_0||^2 \leq 2||\theta_0||^2 + C\varepsilon$, and also $|(\delta\alpha, \theta_0) - (\delta\alpha, \theta)| \leq C\sqrt{\varepsilon}||\alpha||$. This means we can substitute in to replace θ_0 by θ in our upper bound, yielding

$$h_\varepsilon(\alpha, \theta) \leq C_1 e^{C_2\varepsilon||\alpha|| + C_3\varepsilon} \exp\left(-\frac{C}{\beta}||\theta||^2 - \frac{\sqrt{\varepsilon}}{\beta}(\delta\alpha, \theta)\right).$$

We want a uniform bound over all $\varepsilon \in (0, \varepsilon_0)$ (for ε_0 to be chosen later), and the trick is that we can of course bound

this as

$$h_\varepsilon(\alpha,\theta) \le 2C_1 e^{C_2\varepsilon||\alpha||+C_3\varepsilon}\exp\left(-\frac{C}{\beta}||\theta||^2\right)\cosh\left(\frac{\sqrt{\varepsilon}}{\beta}(\delta\alpha,\theta)\right)$$
$$\le 2C_1 e^{C_2\varepsilon_0||\alpha||+C_3\varepsilon_0}\exp\left(-\frac{C}{\beta}||\theta||^2\right)\cosh\left(\frac{\sqrt{\varepsilon_0}}{\beta}(\delta\alpha,\theta)\right).$$

If we now integrate this whole quantity $H(\alpha,\theta)$, we want to show we get something absolutely convergent to apply the dominated convergence theorem, and we handle the resulting Gaussian integral: we end up finding that

$$\int H(\alpha,\theta)d\theta = C_1 e^{C_2\varepsilon_0||\alpha||+C_3\varepsilon_0+C_4\varepsilon_0||\delta\alpha||^2} \le C_1 e^{C_2\varepsilon_0||\alpha||+C_3\varepsilon_0+C_5\varepsilon_0||\alpha||^2}$$

since we have a Gaussian integral with independent components. Taking $\varepsilon_0 \le 1$ without loss of generality, if we do the double integral we get

$$\iint |f(\alpha)|H(\alpha,\theta)d\theta\, d\mu_0(\alpha) \le \sum_{\alpha\in\mathcal{A}_0}\int C_1 e^{C_2||\alpha||+C_3+C_5\varepsilon_0||\alpha||^2}C_6 e^{-C_7||d\alpha||^2}$$
$$\le \sum_{\alpha\in\mathcal{A}_0}\int C_1 e^{C_2||\alpha||+C_3+C_5\varepsilon_0||\alpha||^2}C_6 e^{-C_8||\alpha||^2}$$

so for ε_0 small enough that $C_5\varepsilon_0$ is smaller than C_8, this is finite and we can move the limit inside the integral; $H(\alpha,\theta)$ has no dependence on ε so it indeed shows the finiteness we wanted. □

We'll use this to prove Proposition 91 next time (so far, we haven't actually made any explicit references to this $\mathcal{B}_0$ yet).

18 December 1, 2025

We'll continue with the 4D $U(1)$ deconfinement proof we did before break – for these last two days, we'll take some little things for granted so that we can get to the full proof involving renormalization. Our end goal is to show the perimeter law $\langle W_\ell\rangle \ge C_1 e^{-C_2\text{length}(\ell)}$ for large enough β, and the first step has been to move to a dual lattice $*\Lambda$ – in fact the more important lattice to think about is the "slightly enlarged by $\frac{1}{2}$" Λ^*. Let's again remind ourselves of notation: defining E^* to be all edges of this larger lattice, ∂E^* the boundary edges, and E_0^* the internal edges, we have the spaces $\mathcal{A} = \mathbb{Z}^{E_0^*}$ and $\mathcal{U} = \mathbb{R}^{E_0^*}$ (the 1-forms supported on the internal edges). Defining the equivalence relation $\alpha \sim \alpha'$ if $d(\alpha-\alpha') = 0$, we can decompose $\mathcal{U} = \mathcal{U}_0 \oplus \mathcal{U}_0^\perp$ for $\mathcal{U}_0^\perp = V_0$ the set of α in $\mathcal{U}$ with $d\alpha = 0$, and where we have the explicit characterization

$$\mathcal{U}_0 = \{\alpha\in\mathcal{U} : \delta\alpha(x) = 0 \text{ for all } x\in\Lambda_0^* = \Lambda^*\setminus\partial\Lambda^*\}.$$

It turns out we can also have the equivalent characterization

$$\mathcal{U}_0 = \{\alpha\in\mathcal{U} : \alpha = \Pi\delta\tau \text{ for some } \tau\},$$

where recall that Π is the orthogonal projection onto $\mathcal{U}$ (this is just setting the edges outside to zero). The subtle detail is that $\delta\alpha$ only vanishes on the **internal** vertices even though α is supported on the edges of E_0^*, so we can't just use the ordinary Poincaré lemma. So now $\mathcal{U}_0$ contains one element of each equivalence class, and it's also the smallest-norm element of the class.

Meanwhile on the integer-valued side, we similarly defined $\mathcal{A}_0$ to be a subset of $\mathcal{A}$ of one smallest-norm element per equivalence class. The dual model we care about for our problem is a probability measure μ_0 on the countable set $\mathcal{A}_0$ (or equivalently we could think of it as one on equivalence classes) with probability mass proportional to $\exp\left(-\frac{1}{2\beta}||d\alpha||^2\right)$. (Our gauge group in the dual model is now $\mathbb{Z}$ with addition, but we need to go to the quotient space of orbits to get summability because we now don't have compactness.) Recall that we had the operator $D_\ell : \mathcal{A}_0 \to (0,\infty)$ defined by

$$D_\ell(\alpha) = \exp\left(-\frac{1}{\beta}\sum_{p\in\Sigma} *d\alpha(p) - \frac{|\Sigma|}{2\beta}\right)$$

where Σ is the set of plaquettes enclosed by the rectangular loop ℓ, and where $*d\alpha(p) = d\alpha(*p)$ (since the dual of a plaquette is a plaquette in four dimensions). Visually, remember that Σ is a collection of nicely aligned rectangles, but $*p$ is taking each of those and spinning it around to the other dimensions around the midpoint – they are not forming a perfect tiling. And what we proved earlier on is that

$$\langle W_\ell\rangle = \langle D_\ell\rangle^*;$$

in particular this proves that the Wilson loop is always positive, even though the Wilson loop is a complex-valued random variable.

Remark 95. *Notice that we don't necessarily have uniqueness of the infinite-volume limit even in the primal limit (subsequential limits exist but they may not agree), so this dual lattice construction only really makes sense for finite boxes. And that's also why we're obtaining all of our estimates for finite boxes and then showing that they don't depend on the actual boxes, and in fact the value of $\langle W_\ell\rangle$ may not even be known to be unique in the limit. What our inequalities are then saying is that for any such subsequential limit and any Λ we must have $\langle W_\ell\rangle$ satisfying that inequality.*

So now we want to compare to the continuum setup: we define a Gaussian measure γ_ε on the real-valued $\mathcal{U}$ now with density proportional to $\exp\left(-\frac{1}{2\beta}||d\alpha||^2 + \varepsilon||\alpha||^2\right)$ (which is nondegenerate for any positive $\varepsilon > 0$). Letting γ_ε be randomly drawn from γ_ε and taking its projections $\eta_{0,\varepsilon}, \eta_{1,\varepsilon}$ onto $\mathcal{U}_0, \mathcal{V}_0$ respectively, we proved last time that $\eta_{0,\varepsilon}$ and $\eta_{1,\varepsilon}$ are independent and $\eta_{1,\varepsilon}$ is mean zero and has covariance matrix $\beta\varepsilon^{-1}I$ (which makes sense, since for $\varepsilon = 0$ it should be completely spread out and not well-defined). Meanwhile $\eta_{0,\varepsilon}$ converges in distribution to η_0 following the Gaussian measure γ_0 on $\mathcal{U}_0$ with density proportional to $\exp\left(-\frac{1}{2\beta}||d\alpha||^2\right)$. We'll now compare μ_0 (discrete) with γ_0 (continuum) and try to get the theorem out of that.

To do so, we defined μ_ε to have probability mass function $\exp\left(-\frac{1}{2\beta}(||d\alpha||^2 + \varepsilon||\alpha||^2\right)$, which is a discrete analog of γ_ε. Letting $\langle f\rangle_\varepsilon^*$ be $\int f\,d\mu_\varepsilon$ and $\langle f\rangle^* = \int f\,d\mu_0$, we proved last time that if f is a function of subexponential growth and is constant on equivalence classes, then we can obtain $\langle f\rangle^*$ as the limit of $\langle f\rangle_\varepsilon^*$ (this is the analog of $\gamma_{0,\varepsilon}$ converging to γ_0 in distribution, but it requires more work because we were on the discrete set). The main idea is that if ε is very small, then the measure μ_ε is almost uniform on each equivalence class. But μ_ε is somehow nicer because it's on all of $\mathcal{A}$, while $\mathcal{A}_0$ is not such a nice object.

So the first crucial step of our comparison procedure is a rigorous version of the heuristic identity $d\mu_0(\eta) \propto \left(\sum_{\alpha\in\mathcal{A}\cap\mathcal{U}_0} e^{2\pi i(\alpha,\eta)}\right) d\gamma_0(\eta)$ (which is saying that we have a density with respect to the continuum model). This is basically the statement of Proposition 91 above: for any function of the form $e^{(\tau,\alpha)}$ for $\tau \in \mathcal{U}_0$, we have

$$\langle f\rangle^* = \frac{\sum_{\alpha_0\in\mathcal{B}_0}\mathbb{E}\left[f(\eta_0)e^{2\pi i(\alpha_0,\eta_0)}\right]}{\sum_{\alpha_0\in\mathcal{B}_0}\mathbb{E}\left[e^{2\pi i(\alpha_0,\eta_0)}\right]}.$$

with both sums absolutely convergent. **The key observation** compared to last time is that $\mathcal{B}_0 = \mathcal{A}_0 \cap \mathcal{U}_0$ is actually $\mathcal{A}\cap\mathcal{U}_0$, since if we take anything in $\mathcal{A}\cap\mathcal{U}_0$ it must be the unique smallest-norm element in its equivalence class, so

it is in $\mathcal{A}_0$. (Of course this still doesn't explain why $\mathcal{A} \cap \mathcal{U}_0$ has elements other than the all-zero form, but the lemma says that it must be.)

Proof of Proposition 91. We know that $\langle f \rangle^* = \lim_{\varepsilon \to 0} \langle f \rangle_\varepsilon^*$. Letting m_ε be the probability mass function of μ_ε, we know that $f m_\varepsilon$ is a Schwartz function (because m_ε is quadratically decaying in the exponent) and so by the Poisson summation formula

$$\begin{aligned}\langle f \rangle_\varepsilon^* &= \sum_{\alpha \in \mathcal{A}} (f m_\varepsilon)(\alpha) \\ &= \sum_{\alpha \in \mathcal{A}} \widehat{f m_\varepsilon}(2\pi\alpha) \\ &= \frac{1}{Z_\varepsilon} \sum_{\alpha \in \mathcal{A}} \int_{\mathcal{U}} f(\kappa) e^{2\pi i (\alpha, \kappa)} \exp\left(-\frac{1}{2\beta} ||d\kappa||^2 + \varepsilon ||\kappa||^2 \right) d\kappa.\end{aligned}$$

Similarly taking $f = 1$ we get a formula for the normalizing constant, so in fact

$$\langle f \rangle_\varepsilon^* = \frac{\sum_{\alpha \in \mathcal{A}} \mathbb{E}[f(\eta_\varepsilon) e^{2\pi i (\alpha, \eta_\varepsilon)}]}{\sum_{\alpha \in \mathcal{A}} \mathbb{E}[e^{2\pi i (\alpha, \eta_\varepsilon)}]}.$$

Now taking $\varepsilon \to 0$, we claim that basically the terms go to zero unless $\alpha \in \mathcal{A} \cap \mathcal{U}_0$ so we end up needing to restrict our sum. We know that $\eta_\varepsilon = \eta_{0,\varepsilon} + \eta_{1,\varepsilon}$ is a sum of two independent parts, and f is constant on equivalence classes because $e^{(\tau, \alpha)} = e^{(\tau, \alpha')}$ for any $\alpha \sim \alpha'$ (since $\alpha' - \alpha$ is in the orthogonal complement of $\mathcal{U}_0 \ni \tau$). So in fact

$$\begin{aligned}\mathbb{E}[f(\eta_\varepsilon) e^{2\pi i (\alpha, \eta_\varepsilon)}] &= \mathbb{E}[f(\eta_{0,\varepsilon}) e^{2\pi i (\alpha, \eta_{0,\varepsilon})} e^{2\pi i (\alpha, \eta_{1,\varepsilon})}] \\ &= \mathbb{E}[f(\eta_{0,\varepsilon}) e^{2\pi i (\alpha, \eta_{0,\varepsilon})}] \mathbb{E}[e^{2\pi i (\alpha, \eta_{1,\varepsilon})}] \\ &= \mathbb{E}[f(\eta_{0,\varepsilon}) e^{2\pi i (\alpha, \eta_{0,\varepsilon})}] \exp\left(-\frac{2\pi^2 \beta}{\varepsilon} ||\alpha_1||^2 \right),\end{aligned}$$

where $\alpha = \alpha_0 + \alpha_1$ under the decomposition into the subspaces U_0, V_0 and where $(\alpha, \eta_{1,\varepsilon}) = (\alpha_1, \eta_{1,\varepsilon})$. But now as $\varepsilon \to 0$ this whole thing converges to $\mathbb{E}[f(\eta_0) e^{2\pi i (\alpha, \eta_0)}]$ if $\alpha_1 = 0$ and zero otherwise, so we must have $\alpha \in \mathcal{U}_0$. What's left is to use the dominated convergence to show that it's okay to take the limit within the sum, and we won't go through the details of that here. □

So we're getting closer to a point where we can replace our expectations with something smooth, which is going back to something we did earlier in the class. We choose a sequence of positive real numbers $\{z_q\}_{q \geq 1}$ with $\sum_q \frac{2}{z_q} = 1$ and let $\mathcal{Q} = \{1, 2, \cdots, \}^{E_0^*}$. For $q \in \mathcal{Q}$ we then let $c(q) = \prod_{e \in E_0^*} \frac{2}{z_{q(e)}}$; analogously we can do a finite approximation $\mathcal{Q}_N = \{1, \cdots, n\}^{E_0^*}$ and define $a_N = \sum_{q=1}^N \frac{2}{z_q}$. We can then define the function ρ_N on $\mathcal{U}$ via

$$\rho_N(\alpha) = \sum_{q \in \mathcal{Q}_N} c(q) \prod_{e \in E_0^*} \left(1 + a_N z_{q(e)} \cos(2\pi i \alpha(e) q(e)) \right).$$

Lemma 96

For $\tau \in \mathcal{U}_0$ we define the function $f(\alpha) = e^{\tau(\alpha)}$, and let η_ε have law γ_ε and decomposition $\eta_{0,\varepsilon} + \eta_{1,\varepsilon}$. Our expectation of interest then satisfies

$$\langle f \rangle^* = \lim_{N \to \infty} \lim_{\varepsilon \to 0} \frac{\mathbb{E}[f(\eta_{0,\varepsilon}) \rho_N(\eta_\varepsilon)]}{\mathbb{E}[\rho_N(\eta_\varepsilon)]}.$$

This is almost what we proved before but with the ε and N interchanged, and we're phrasing things in this way because we do want to make use of the $\eta_{0,\varepsilon}$s (we'll see soon).

Proof. Defining $\mathcal{A}_N = \{-N, \cdots, N\}^{E_0^*}$, we know from what we've already proven that

$$\langle f \rangle^* = \lim_{N\to\infty} \frac{\sum_{\alpha\in\mathcal{A}_N\cap\mathcal{U}_0} \mathbb{E}[f(\eta_0)e^{2\pi i(\alpha,\eta_0)}]}{\sum_{\alpha\in\mathcal{A}_N\cap\mathcal{U}_0} \mathbb{E}[e^{2\pi i(\alpha,\eta_0)}]}$$

because we know that our previous numerator and denominator with the full $\mathcal{A}$s converged. But $\mathcal{A}_N$ are each finite sets (and nonempty for large enough N), so we can in fact take the ε limit for any fixed N; this means

$$\langle f \rangle^* = \lim_{N\to\infty} \lim_{\varepsilon\to 0} \frac{\sum_{\alpha\in\mathcal{A}_N} \mathbb{E}[f(\eta_0)e^{2\pi i(\alpha,\eta_\varepsilon)}]}{\sum_{\alpha\in\mathcal{A}_N} \mathbb{E}[e^{2\pi i(\alpha,\eta_\varepsilon)}]}.$$

Crucially here (by the same logic as before), we don't need to take intersections with $\mathcal{U}_0$ anymore since as $\varepsilon \to 0$ the terms that aren't in $\mathcal{U}_0$ are going away anyway, so we can actually sum over all terms in $\mathcal{A}_N$. But this is now nice because we can write it as an E_0^*-fold sum over $\{-N, \cdots, N\}$, and then the remaining calculation to get to the form of ρ_N has already been previously done. □

So the density of μ_0 now takes $\rho_N(\eta_\varepsilon)$, then sends $\varepsilon \to 0$, and then takes $N \to \infty$. We're going to keep reinterpreting this density step by step further, since ρ_N is still very spiky. Remember that for each (integer-valued) $\alpha \in \mathcal{A}$ we define the support $E(\alpha) = \{e \in E_0^* : \alpha(e) \neq 0\}$ and say that $\mathcal{E} \subseteq \mathcal{A}$ form an ensemble if $E(\alpha) \cap E(\alpha') = \varnothing$ for all disjoint $\alpha, \alpha' \in \mathcal{E}$ and a k-ensemble if $\operatorname{dist}(E(\alpha), E(\alpha')) \geq 2^{k/2}$. (In particular, 1-ensembles mean that plaquettes cannot contribute to different forms simultaneously.) We previously proved this result in Lemma 90, but we'll rearrange the quantifiers in a slightly different way which is nicer for us:

Lemma 97

For all $q \in \mathcal{Q}$, there is a finite set $\Gamma(q)$ which indexes strictly positive coefficients $\{c_\gamma\}_{\gamma\in\Gamma(q)}$ and 1-ensembles $\{\mathcal{E}_\gamma\}_{\gamma\in\Gamma(q)}$, such that for any $N \geq 1$, there are strictly positive coefficients $\{K_\gamma^N(\xi)\}_{\gamma\in\Gamma(q),\eta\in\mathcal{E}_\gamma}$ so that for all $\alpha \in \mathcal{U}$ we have

$$\prod_{e\in E_0^*} \Big(1 + a_N z_{q(e)} \cos(2\pi q(e)\alpha(e))\Big) = \sum_{\gamma\in\Gamma(q)} c_\gamma \prod_{\xi\in\mathcal{E}_\gamma} \Big(1 + K_\gamma^N(\xi)\cos(2\pi(\alpha,\xi))\Big),$$

and furthermore we have that $K_\gamma^N(\xi) \leq 3^{N_1(\xi)} \prod_{e\in E(\xi)} z_{q(e)}$ for $N_1(\eta)$ the number of edges $e \in E_0^*$ of distance at most 1 from $E(\eta)$.

Remember that our proof involved putting together ensembles together by replacing products with finite sums of products using a strange cosine identity. We want to consider products of 1-forms in $\mathcal{U}_0$, so we want to be able to simplify the expectation when ξ is drawn from a Gaussian measure γ_0; that's easier to do on the right-hand side than on the left-hand side. (We'll do some kind of conditioning on certain edges of α in a renormalization step, and it'll be done in a way where the replacement of K_γ^N becomes very small instead of very big.)

So now without loss of generality we can assume all $\Gamma(q)$s are disjoint (since they're just indexing sets) and let $\Gamma = \bigcup_{q\in\mathcal{Q}} \Gamma(q)$ and $\Gamma_N = \bigcup_{q\in\mathcal{Q}_N} \Gamma(q)$. We get the following corollary:

Corollary 98

Our density can be written as

$$\rho_N(\alpha) = \sum_{\gamma \in \Gamma_N} c_\gamma \prod_{\xi \in \mathcal{E}_\gamma} \left(1 + K_\gamma^N(\xi) \cos(2\pi(\alpha, \xi))\right).$$

Of course, as written it's not clear that this is positive yet, but our goal will be to chip away at that and end up with something positive. We can replace ρ_N by the following simpler object: let $\mathcal{E}_\gamma^0 = \mathcal{E}_\gamma \cap \mathcal{U}_0$ (remember $\mathcal{E}_\gamma$ is always some collection of integer-valued 1-forms), and define analogously

$$\rho_N^0(\alpha) = \sum_{\gamma \in \Gamma_N} c_\gamma \prod_{\xi \in \mathcal{E}_\gamma^0} \left(1 + K_\gamma^N(\xi) \cos(2\pi(\alpha, \xi))\right).$$

The point will be to show that we can actually replace ρ_N by ρ_N^0, and at that point we can take away the ε:

Lemma 99

We have that

$$\langle f \rangle^* = \lim_{N \to \infty} \frac{\mathbb{E}[f(\eta_0)\rho_N^0(\eta_0)]}{\mathbb{E}[\rho_N^0(\eta_0)]}.$$

What's nice is that even if ρ_N^0 is not that much nicer than ρ_N, the next step we take will be able to be carried out with this modified density, and we'll see that next time! It'll basically come down to a Girsanov-type "completing the square," where the D_ℓ factor will be used to do a change of variable to express the numerator and denominator as certain shifts. But that factor will be exactly what is needed to give us the perimeter law, and the ratio of ρ_N with the shift will be lower bounded accordingly.

19 December 3, 2025

Last time, we showed that the dual expectation $\langle D_\ell \rangle^*$ can be written in the form $\lim_{N\to\infty} \lim_{\varepsilon \to 0} \frac{\mathbb{E}[D_\ell(\eta_{0,\varepsilon})\rho_N(\eta_\varepsilon)]}{\mathbb{E}[\rho_N(\eta_\varepsilon)]}$, where $\eta_0 \sim \gamma_0$ is a Gaussian 1-form (on the subspace $\mathcal{U}_0$ with density proportional to $\exp\left(-\frac{1}{2\beta}||d\alpha||^2\right)$) and ρ_N is some positive linear combination of $\prod(1 + K_\gamma^N(\xi)\cos(2\pi(\alpha, \xi)))$ for various 1-ensembles (meaning these ensembles are sufficiently separated from each other). But $\rho_N(\alpha)$ can take both positive and negative values and the K^Ns can be very large, so we don't have a guarantee on the sign with this.

The lemma we claimed last time is that ρ_N can be replaced by ρ_N^0, in which each $\mathcal{E}_\gamma$ is replaced by $\mathcal{E}_\gamma \cap \mathcal{U}_0$.

Proof of Lemma 99. For any function $f : \mathcal{U} \to \mathbb{C}$ which is constant on equivalence classes, we have (with the same notation as last time)

$$\begin{aligned}
\mathbb{E}[f(\eta_{0,\varepsilon})\rho_N(\eta_\varepsilon)] &= \sum_{\gamma \in \Gamma_N} c_\gamma \mathbb{E}\left[f(\eta_{0,\varepsilon}) \prod_{\xi \in \mathcal{E}_\gamma} (1 + K_\gamma^N(\xi)\cos(2\pi(\eta_\varepsilon, \xi)))\right] \\
&= \sum_{\gamma \in \Gamma_N} c_\gamma \sum_{\mathcal{E} \subset \mathcal{E}_\gamma} \left(\prod_{\xi \in \mathcal{E}} K_\gamma^N(\xi)\right) \mathbb{E}\left[f(\eta_{0,\varepsilon}) \prod_{\xi \in \mathcal{E}} \cos(2\pi(\eta_\varepsilon, \xi))\right]
\end{aligned}$$

by expanding out the product by distributivity. Now using $\cos(x) = \frac{e^{ix}+e^{-ix}}{2}$ and again expanding out, we end up with

$$\sum_{\gamma\in\Gamma_N} c_\gamma \sum_{\mathcal{E}\subset\mathcal{E}_\gamma} \left(\prod_{\xi\in\mathcal{E}} K_\gamma^N(\xi)\right) \frac{1}{2^{|\mathcal{E}|}} \sum_{\sigma\in\{-1,1\}^{\mathcal{E}}} \mathbb{E}\left[f(\eta_{0,\varepsilon}) \exp\left(2\pi i \sum_{\xi\in\mathcal{E}} \sigma_\xi(\eta_\varepsilon,\xi)\right)\right]$$
$$= \sum_{\gamma\in\Gamma_N} c_\gamma \sum_{\mathcal{E}\subset\mathcal{E}_\gamma} \left(\prod_{\xi\in\mathcal{E}} K_\gamma^N(\xi)\right) \frac{1}{2^{|\mathcal{E}|}} \sum_{\sigma\in\{-1,1\}^{\mathcal{E}}} \mathbb{E}\left[f(\eta_{0,\varepsilon}) \exp\left(2\pi i(\eta_\varepsilon,\alpha_\sigma)\right)\right]$$

where $\alpha_\sigma = \sum_{\xi\in\mathcal{E}} \sigma_\xi \xi$. We want to show that $\mathcal{E}_\gamma$ can be replaced with $\mathcal{E}_\gamma \cap \mathcal{U}_0$, so now the claim is that if $\mathcal{E}\setminus\mathcal{U}_0$ is nonempty then this whole contribution is actually zero. Indeed, if $\kappa\in\mathcal{E}\setminus\mathcal{U}_0$, then (by our alternate characterization of $\mathcal{U}_0$) there is some $x\in\Lambda_0^*$ such that $\delta\kappa(x)=0$. Therefore there must be some $e\in E_0^*$ incident to x with $\kappa(e)\neq 0$, but that implies that no other edge incident to x is in $E(\xi)$ for $\xi\in\mathcal{E}, \xi\neq\kappa$ (by the ensemble property). So

$$\delta\alpha_\sigma(x) = \sigma_\kappa \delta(\kappa(x)) \neq 0,$$

which means we must have $\alpha_\sigma \notin \mathcal{U}_0$ for all possible signs σ. But if we're not in $\mathcal{U}_0$, then as $\varepsilon\to 0$ we have

$$\frac{1}{2^{|\mathcal{E}|}} \sum_{\sigma\in\{-1,1\}^{\mathcal{E}}} \mathbb{E}\left[f(\eta_{0,\varepsilon}) \exp\left(2\pi i(\eta_\varepsilon,\alpha_\sigma)\right)\right] \to 0 \text{ as } \varepsilon\to 0$$

because η_ε breaks up into $\eta_{0,\varepsilon}$ and $\eta_{1,\varepsilon}$, and the huge-variance part of $\eta_{1,\varepsilon}$ makes the expectation vanish in the limit if α_σ always has a nonzero component in that direction. Thus we can indeed take intersections with $\mathcal{U}_0$ and get the desired result. □

We can think of this as saying that we've now evaluated the limit $\varepsilon\to 0$ and found which terms contribute in that limit, and now we have

$$\langle D_\ell\rangle^* = \lim_{N\to\infty} \frac{\mathbb{E}[D_\ell(\eta_0)\rho_N^0(\eta_0)]}{\mathbb{E}[\rho_N^0(\eta_0)]}, \quad D_\ell(\alpha) = \exp\left(\frac{1}{\beta}\sum_{p\in\Sigma} d\alpha(*p) - \frac{|\Sigma|}{2\beta}\right).$$

We can write this in a more convenient form: instead of Σ being the set of plaquettes enclosed by our loop, we write down the 2-form σ on $*\mathbb{Z}^4$ with $\sigma(p)=1$ if $p\in *\Sigma$ and $\sigma(p)=0$ otherwise. Then by a direct computation we have

$$D_\ell(\alpha) = \exp\left(\frac{1}{\beta}(d\alpha,\sigma) - \frac{1}{2\beta}||\sigma||^2\right),$$

and $(d\alpha,\sigma) = (\alpha,\delta\sigma)$ is actually an observable of the form we've been studying, since $\delta\sigma\in\mathcal{U}_0$. (Indeed we just need our loop to not touch the boundary, so then $\delta\sigma$ is a 1-form and nonzero only inside with δ of it equal to zero, so it satisfies our condition of being of the form $f(\alpha)=e^{(\alpha,\tau)}$ for $\tau\in\mathcal{U}_0$.) So we can apply our earlier results, and so what we get is that

$$\langle D_\ell\rangle^* = \lim_{N\to\infty} \frac{\mathbb{E}\left[\exp\left(\frac{1}{\beta}(d\eta_0,\sigma) - \frac{1}{2\beta}||\sigma||^2\right)\rho_N^0(\eta_0)\right]}{\mathbb{E}[\rho_N^0(\eta_0)]},$$

so we basically have a linear tilting factor to our Gaussian measure. In such cases we can do a change of variables (completing the square) and get the following:

Lemma 100

Let $\tau = -\delta\Delta_{\Lambda^*}^{-1}\sigma$ and $\kappa = -\Pi_{\Lambda^*}\delta d\Delta_{\Lambda^*}^{-1}\sigma$. Then we have

$$\langle D_\ell\rangle^* = \exp\left(\frac{1}{2\beta}||d\tau - \Pi_{\Lambda^*}d\tau||^2 - \frac{1}{2\beta}||\kappa||^2\right)\lim_{N\to\infty}\frac{\mathbb{E}[\rho_N^0(\eta_0+\tau)]}{\mathbb{E}[\rho_N^0(\eta_0)]},$$

where Δ_{Λ^*} is the **Hodge Laplacian** with zero boundary conditions on Λ^* (explained below).

Explicitly, we define $\Delta = \delta d + d\delta$ to be an operator on real-valued k-forms, and for any hypercube Λ^* we have a projection operator Π_{Λ^*} on k-forms defined by only keeping the interior k-cells

$$\Pi_{\Lambda^*}\alpha(c) = \begin{cases}\alpha(c) & \text{if all vertices of } c \text{ are in } \Lambda^* \text{ but not all in the boundary of } \Lambda^*,\\ 0 & \text{otherwise.}\end{cases}$$

Then $\Delta_{\Lambda^*} = \Pi_{\Lambda^*}\Delta$. This object is actually not so complicated: if we're working with 0-forms, then this just reduces to the usual Laplacian as we understand it. And if $f = \sum_{1\le i_1<cdots<i_k\le n} f_{i_1,\cdots,i_k}dx_{i_1}\wedge\cdots\wedge dx_{i_k}$, then actually we just apply Δ component-wise. Thus inverting it is also easy because we just need to invert the usual discrete Laplacian using the discrete Green's function. And thus here τ is just some 1-form determined by the loop, but it's weird because it's not supported on some set of edges just around the loop because of that Green's function term – it can be nonzero throughout the hypercube.

Letting $\mathcal{U}(\Lambda^*)$ be the range of Π_{Λ^*} for some fixed k, we have the set of all k-forms that are nonzero on any non-interior k-cells of Λ^*, and we have the following fact:

Lemma 101

The map $\Delta_{\Lambda^*}:\mathcal{U}(\Lambda^*)\to\mathcal{U}(\Lambda^*)$ is a linear bijection.

Proof. $\mathcal{U}(\Lambda^*)$ is a finite-dimensional vector space and the operator maps it into itself, so we just need to prove injectivity. If $\Delta_{\Lambda^*}(\alpha) = 0$ for some $\alpha\in\mathcal{U}(\Lambda^*)$, then

$$(\alpha,\Delta_{\Lambda^*}\alpha) = 0 \implies (\alpha,\Pi_{\Lambda^*}\Delta\alpha) = 0 \implies (\alpha,\Delta\alpha) = 0,$$

but this means $(\alpha,(\delta d+d\delta)\alpha) = 0$ and so $||d\alpha||^2+||\delta\alpha||^2 = 0$ by summation by parts; this implies that $\alpha = d\tau = \delta\xi$ for some τ,ξ, and thus $(\alpha,\alpha) = (d\tau,\delta\xi) = (\tau,\delta\delta\xi) = 0$, and therefore $\alpha = 0$. So injectivity proves bijectivity. □

Our point now is to do a renormalization transform to understand the ratio, and this is a particularly instructive idea. We have

$$\rho_N^0(\eta_0) = \sum_{\gamma\in\Gamma_N} c_\gamma\prod_{\xi\in\mathcal{E}_\gamma^0}(1+K_N^\gamma(\xi)\cos(2\pi(\eta_0,\xi))),$$

and so

$$\rho_N^0(\eta_0+\tau) = \sum_{\gamma\in\Gamma_N} c_\gamma\prod_{\xi\in\mathcal{E}_\gamma^0}(1+K_N^\gamma(\xi)\cos(2\pi(\eta_0,\xi)+\theta_\xi)),\quad \theta_\xi = 2\pi(\tau,\xi).$$

Let's understand what happens with these general phase terms θ_ξ. We'll consider a renormalization

$$\xi\to\overline{\xi},\quad K_N^\gamma(\xi)\to R_N^\gamma(\xi),$$

so that for any fixed γ we have the **exact equality**

$$\mathbb{E}\left[\prod_{\xi\in\mathcal{E}_\gamma^0}(1+K_N^\gamma(\xi)\cos(2\pi(\eta_0,\xi)+\theta_\xi))\right]=\mathbb{E}\left[\prod_{\xi\in\mathcal{E}_\gamma^0}(1+R_N^\gamma(\xi)\cos(2\pi(\eta_0,\bar{\xi})+\theta_\xi))\right]$$

so that $R_N^\gamma(\xi)$ is now small so that everything is always positive inside the product (and thus we can apply things like Jensen's inequality). Indeed,

$$\begin{aligned}\mathbb{E}\left[\prod_{\xi\in\mathcal{E}_\gamma^0}(1+K_N^\gamma(\xi)\cos(2\pi(\eta_0,\xi)+\theta_\xi))\right]&=\lim_{\varepsilon\to 0}\mathbb{E}\left[\prod_{\xi\in\mathcal{E}_\gamma^0}(1+K_N^\gamma(\xi)\cos(2\pi(\eta_{0,\varepsilon},\xi)+\theta_\xi))\right]\\&=\lim_{\varepsilon\to 0}\mathbb{E}\left[\prod_{\xi\in\mathcal{E}_\gamma^0}(1+K_N^\gamma(\xi)\cos(2\pi(\eta_\varepsilon,\xi)+\theta_\xi))\right]\end{aligned}$$

first because $\eta_{0,\varepsilon}$ converges in distribution to η_0, and then because $\xi\in\mathcal{U}_0$. If we now define the 1-form $\overline{\eta}_\varepsilon$ via

$$\overline{\eta}_\varepsilon(e)=\frac{1}{6}\sum_{\substack{p\in P^*\\ p\ni e}}\sum_{e'\in p,e'\neq e}s_{e,e'}\eta_\varepsilon(e'),$$

where for any (positively oriented) edges $e\neq e'$ in the same plaquette (there's always a unique one), we define $s_{e,e'}=-1$ if e,e' have the same orientation when traversing p as a loop and 1 otherwise. Notice that $\overline{\eta}_\varepsilon(e)$ has no dependence on $\eta_\varepsilon(e)$ itself, only on the neighboring edges; the point of this definition is the following:

Lemma 102

Because η_ε has density $\exp\left(-\frac{1}{2\beta}||d\alpha||^2+\varepsilon||\alpha||^2\right)$, we have that conditionally on $\{\eta_\varepsilon(e')\}_{e'\in E_0^*\setminus\{e\}}$, we have

$$\eta_\varepsilon(e)\sim N\left(\frac{6\overline{\eta}_\varepsilon(e)}{6+\varepsilon},\frac{\beta}{6+\varepsilon}\right).$$

Thus when we evaluate our expected value of interest, take any $\gamma\in\Gamma_N$ and any $\xi\in\mathcal{E}_\gamma^0$. Call e,e' neighbors if $\text{dist}(e,e')<\sqrt{2}$ (this can only happen within a single collection); the maximum degree of this graph is at most some universal constant C because we're in 4 dimensions, and thus the resulting graph is $(C+1)$-colorable by a greedy algorithm. Therefore there is some subset $\mathcal{B}_\xi\subseteq E(\xi)$, such that for all distinct $e,e'\in\mathcal{B}_\xi$, we have $\text{dist}(e,e')\geq\sqrt{2}$ and also

$$\sum_{e\in\mathcal{B}_\xi}\xi(e)^2\geq\frac{1}{c+1}||\xi||^2$$

(by picking the best possible color for this sum-of-squares quantity). We will therefore define

$$\bar{\xi}(e)=0\text{ if }e\in\mathcal{B}_\xi,\quad\bar{\xi}(e)=\frac{1}{6}\sum_{\substack{p\in P^*\\ p\ni e}}\sum_{\substack{e'\in p\\ e'\in\mathcal{B}_\xi}}s_{e,e'}\xi(e')\text{ otherwise.}$$

Notice that this is a pretty inefficient way to write the sum – even though we have a double sum, there can only actually ever be at most one nonzero contribution total because all of the edges in $\mathcal{B}_\xi$ are supposed to be at least $\sqrt{2}$ apart, and that's not true for disjoint edges in any plaquettes containing e. And we will also define

$$R_\gamma^N(\xi)=K_\gamma^N(\xi)\exp\left(-\frac{\pi^2\beta}{3}\sum_{e\in\mathcal{B}_\xi}\xi(e)^2\right).$$

We'll see soon that plugging in these renormalized values exactly gives us the desired identity above, and furthermore this R is small because we previously proved the bound $K_\gamma^N(\xi) \leq 3^{N_1(\xi)} \prod_{e \in E(\xi)} z_{|\xi(e)|}$ (where we previously chose our qs so that $\sum \frac{2}{z_q} = 1$, and **now we make the specific choice** $z_q = e^{\beta_0 q^2}$ for the exact value of β_0 which makes this sum true); thus we can prove that in fact

$$K_\gamma^N(\xi) \leq e^{\beta_1 ||\xi||^2}$$

for some β_1 depending on β_0. So for β chosen large enough we can make $R_\gamma^N(\xi)$ small, since we're multiplying by something at most $\exp\left(-\frac{\pi^2\beta}{3} \cdot \frac{1}{c+1}||\xi||^2\right)$.

So to show why this renormalization identity is true, the point is that **conditioned on everything outside** $\mathcal{B}B_\xi$ it's not that hard to work this out: we're interested in

$$\mathbb{E}\left[\prod_{\xi \in \mathcal{E}_\gamma^0} (1 + K_N^\gamma(\xi)\cos(2\pi(\eta_\varepsilon, \xi) + \theta_\xi))\right] = \sum_{\mathcal{E} \subseteq \mathcal{E}_\gamma^0} \left(\prod_{\xi \in \mathcal{E}} K_\gamma^N(\xi)\right) \left(2^{-|\mathcal{E}|} \sum_{\sigma \in \{-1,1\}^{\mathcal{E}}} \mathbb{E}\left[\exp\left(i \sum_{\xi \in \mathcal{E}} (2\pi\sigma_\xi(\eta_\varepsilon, \xi) - \sigma_\xi\theta_\xi\right)\right]\right)$$

by doing the same expansion over all subsets. Now for a nonempty $\mathcal{E} \subseteq \mathcal{E}_\gamma^0$, we can evaluate the innermost sum by defining (this is a disjoint union)

$$\mathcal{B}_\mathcal{E} = \bigcup_{\xi \in \mathcal{E}} \mathcal{B}_\xi.$$

Note that $e, e' \in \mathcal{B}_\mathcal{E}$ are always at least $\sqrt{2}$ apart, either by the 1-ensemble property or the definition of individual $\mathcal{B}_\xi$s. Thus if we take this inner expectation (ignoring the phase term)

$$\mathbb{E}\left[\exp\left(i \sum_{\xi \in \mathcal{E}} (2\pi\sigma_\xi(\eta_\varepsilon, \xi)\right)\right] = \mathbb{E}\left[\exp\left(i \sum_{\xi \in \mathcal{E}} \sum_{e \in E(\xi)} 2\pi\sigma_\xi \eta_\varepsilon(e)\xi(e)\right)\right],$$

which is a double sum, we can only take the terms with e corresponding to $\mathcal{B}_\mathcal{E}$, and so we can consider the **conditional expectation**

$$\mathbb{E}\left[\exp\left(i \sum_{\xi \in \mathcal{E}} \sum_{e \in \mathcal{B}_\mathcal{E}} 2\pi\sigma_\xi \eta_\varepsilon(e)\xi(e)\right) \middle| \{\eta_\varepsilon(e)\}_{e \in \mathcal{B}_\mathcal{E}^\xi}\right]$$

and plug that in instead by the tower law. But in this conditional expectation, we know that the $\eta_\varepsilon(e)$s for e within $\mathcal{B}_\mathcal{E}$ become conditionally independent and we can evaluate it exactly: we get that it evaluates to

$$\exp\left(\sum_{\xi \in \mathcal{E}} \sum_{e \in \mathcal{B}_\mathcal{E}} \left(\frac{12\pi i \sigma_\xi \overline{\eta}_\varepsilon(e)\xi(e)}{6+\varepsilon} - \frac{2\pi^2\beta\xi(e)^2}{6+\varepsilon}\right)\right).$$

So we get some terms that suppress our large constants K^N, and furthermore the way we defined things we exactly get $(\overline{\eta}_\varepsilon, \xi) = (\eta_\varepsilon, \overline{\xi})$. So now we can trace our steps back with $\overline{\xi}$ instead of ξ, and that gives us the desired identity with the second fraction terms suppressing the factors of K_γ^N.

So in summary, expanding out the cosines and then taking conditional expectations lets us integrate out some coordinates with respect to the Gaussian measure. In the end, things seem like they depend on the full η_ε, but they don't actually depend on the values in $\mathcal{B}_\mathcal{E}$ because $\overline{\xi}$ is zero at the edges in $\mathcal{B}_\mathcal{E}$. That gets us from the "weird" oscillatory ρ_N^0 to a nice positive quantity.