RANDOM CURVES ON SURFACES INDUCED FROM THE LAPLACIAN DETERMINANT

BY ADRIEN KASSEL[1] AND RICHARD KENYON[2]

ETH Zürich and Brown University

We define natural probability measures on finite multicurves (finite collections of pairwise disjoint simple closed curves) on curved surfaces. These measures arise as universal scaling limits of probability measures on cycle-rooted spanning forests (CRSFs) on graphs embedded on a surface with a Riemannian metric, in the limit as the mesh size tends to zero. These in turn are defined from the Laplacian determinant and depend on the choice of a unitary connection on the surface.

Wilson's algorithm for generating spanning trees on a graph generalizes to a cycle-popping algorithm for generating CRSFs for a general family of weights on the cycles. We use this to sample the above measures. The sampling algorithm, which relates these measures to the loop-erased random walk, is also used to prove tightness of the sequence of measures, a key step in the proof of their convergence.

We set the framework for the study of these probability measures and their scaling limits and state some of their properties.

1. Introduction. Classical statistical mechanics deals with systems of large numbers of particles interacting through local forces. These systems are naturally defined on Euclidean spaces, so that the notions of scaling limit and scale invariance make sense. In this work, we define a large family of statistical mechanical systems on curved spaces: curved surfaces with possibly nontrivial topology, where the curvature and topology play both a local and global role in the underlying probability measure. By scaling limit in such a context, we mean that as the system size grows we shrink the discretization parameter so that the metric properties of the underlying surface remain constant.

A fundamental property of our scaling limits is that they are *universal*: they are independent of the details of the discrete approximating sequence. In other words, they are natural, parameter-free systems on the curved surface itself, depending only on its geometry and topology.
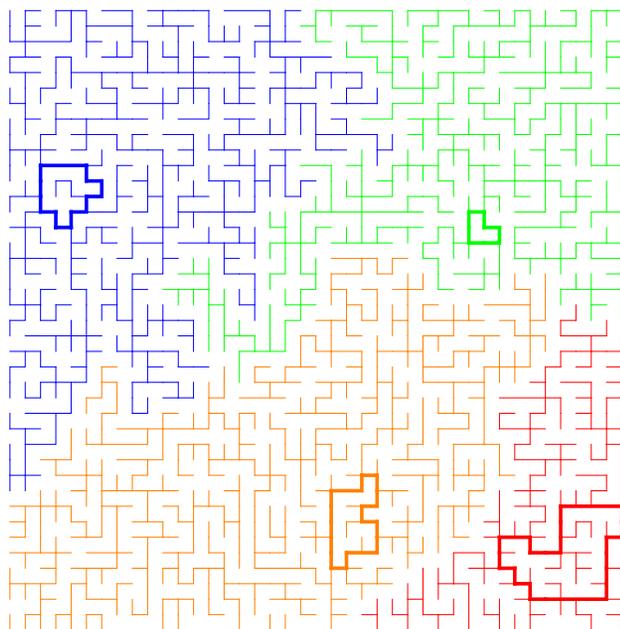
FIG. 1. *A CRSF on the* 40×40 *square grid*; *each connected component is in a different color and the cycles are in bold*.

Let us detail the systems we consider. A *cycle-rooted spanning forest* (CRSF) on a graph $\mathcal{G}$ is a subgraph each of whose components contains a unique cycle, or equivalently, contains as many vertices as edges; see Figure 1. A *cycle-rooted spanning tree* (CRST) is a connected CRSF.

Natural probability measures on CRSFs arising from the determinant of the graph Laplacian were introduced in [15]: the probability of a CRSF is proportional to the product over its cycles of a certain function of the cycle, depending on the holonomy of a discrete $\mathbb{C}^*$- or $SL_2(\mathbb{C})$-connection. The interest of these measures is that they can give to a cycle a weight which is a function of its shape. Furthermore, these measures are determinantal viewed as point processes on the set of edges.

We study here the *scaling limits* of these measures: their limits on a sequence of finer and finer graphs approximating a surface with a fixed Riemannian metric. By *surface*, we will mean here an oriented smooth surface with a Riemannian metric. By *approximation* of a surface Σ, we will mean a sequence of graphs $(\mathcal{G}_n)$ geodesically embedded on Σ and *conformally approximating* it in a sense defined below (essentially, the simple random walk on $\mathcal{G}_n$ converges to Brownian motion on Σ). We endow the space of multiloops on Σ with a natural topology and show the weak convergence of the above probability measures on multiloops on $\mathcal{G}_n$.

We may informally summarize our main statement as follows.

THEOREM (see Theorem 20).   *For any Riemannian surface, there is a natural probability measure on finite collections of disjoint simple closed curves drawn on it (which are fractal, locally resembling the Schramm–Loewner evolution* SLE_2*). This measure is the universal scaling limit of natural discrete measures on CRSF loops defined on graphs conformally approximating the surface. When the surface has nontrivial topology and its metric is flat, these curves are noncontractible. In the case of the flat Euclidean disk, the probability measure is degenerate but a measure on single curves is obtained by a limiting procedure.*

We study scaling limits in two different settings: a topological setting and a geometrical setting.

In the topological setting, we consider only the conformal class of the metric on Σ. Define a *noncontractible* CRSF to be a CRSF with no contractible cycles. Let $(\mathcal{G}_n)$ be an approximating sequence and μ^n_{nonc} be the uniform measure on non-contractible CRSFs of $\mathcal{G}_n$. We show (Theorem 18) that the cycle process $\mathbb{P}^n_{\mathrm{nonc}}$ of the μ^n_{nonc}-random CRSF on $\mathcal{G}_n$ converges to a random loop process $\mathbb{P}_{\mathrm{nonc}}$ on Σ, *independent* of the approximating sequence $\mathcal{G}_n$. The limit only depends on the conformal class of Σ, in the following sense. Let $z_1, \ldots, z_k$ be distinct points of Σ. For any isotopy class of sets of pairwise disjoint simple loops $\{\gamma_1, \ldots, \gamma_m\}$ of $\Sigma \setminus \{z_1, \ldots, z_k\}$, the probability that a random noncontractible CRSF on $\mathcal{G}_n$ has m cycles, and these are isotopic to the γ_i, has a probability converging as $n \to \infty$ to a limit independent of the approximating sequence $\mathcal{G}_n$. This refines the result of [16] who showed that (for the dimer model, which is closely related to the CRSF model via Temperley's bijection [17]) the distribution of the homotopy classes of the cycles in $\pi_1(\Sigma)$ has a conformally invariant limit when Σ is a planar domain. A related (infinite) measure μ_{-2} on simple loops on Σ was recently constructed in a very different manner by Benoist and Dubedat in [4]. The measure μ_{-2} restricted to noncontractible loops of the annulus (and normalized to be a probability measure) is the same as our measure $\mathbb{P}_{\mathrm{nonc}}$ conditioned to have one loop. When extended to all surfaces in such a way that a "conformal restriction" property is satisfied, the measure μ_{-2} was conjectured to exist by Kontsevich and Suhov. Further relations between our measures and μ_{-2} will be considered in a forthcoming paper.

In the geometrical setting, we take into account the metric on Σ, and in particular its curvature. Let $(\mathcal{G}_n)$ be a sequence of finite graphs conformally approximating Σ. Associated to this data is a discrete connection on a complex line bundle over $\mathcal{G}_n$ arising from the Levi-Civita connection on the tangent bundle on Σ. It is defined up to gauge equivalence by the property that the holonomy around a loop is $e^{i\theta}$ where θ is the enclosed curvature. From this connection Φ^n_{LC}, we construct a natural probability measure μ^n_{LC} on CRSFs on $\mathcal{G}_n$: each CRSF has a probability proportional to the product over its cycles of $2 - 2\cos\theta$, where θ is the curvature enclosed. The corresponding loop process $\mathbb{P}^n_{\mathrm{LC}}$ is shown (Theorem 20) to converge

to a probability measure $\mathbb{P}_{LC}$ on multicurves on Σ, independent of the approximating sequence.

When the surface Σ is contractible, we define another measure $\mu_{LC^0}^n$ which is in some sense more natural. This measure is a limit when $\varepsilon \to 0$ of the CRSF connection measures for the connection with curvature $e^{i\varepsilon\theta}$, a limit which was introduced in [15]. This yields a measure on cycle-rooted spanning trees (CRSTs, i.e., CRSFs with one component) with weight proportional to θ^2 where θ is the enclosed curvature. We show (Theorem 20) that the loop measure $\mathbb{P}_{LC^0}^n$ converges to a measure $\mathbb{P}_{LC^0}$ on simple closed curves on Σ, again independent of the approximating sequence.

We give a "cycle-popping" algorithm (Theorem 1) for rapid exact sampling from the above measures (as well as more general measures), generalizing the well-known cycle-popping algorithm of Wilson [23] for generating uniform spanning trees. One simply runs Wilson's algorithm, and when a cycle is created, flip a coin (with bias depending on the cycle weight) to decide whether to keep it or not.

We use this sampling algorithm to sample approximations of the above scaling limits. In all three cases, the cycle-popping algorithm is an essential part of the convergence argument: it is used to show tightness of the sequence of measures (Section 4.4).

Another essential result shows that there are almost surely a finite number of components in a random CRSF, and the scaling limits of the loops are nondegenerate, in the sense that they do not shrink to points in the limit. This is accomplished by computing the universal limit of the probability of having no loops and by exploring in a Markovian way via the algorithm the surface with positive probability of creating macroscopic loops at each step. This implies a super-exponential tail for the number of loops which excludes the fact of having microscopic loops since otherwise their number would be infinite by the weak large scale dependence of the process.

We give samples from the measures μ_{LC} and μ_{LC^0} for the round sphere (Figure 2), a saddle surface (Figure 3, left) and a compact disk in the Poincaré plane (Figure 3, right), and for the measure μ_{nonc} on a flat torus (Figure 4) and planar domains (Figure 5). For μ_{LC}, these are conditional samples, conditioned on having only loops with area (curvature) bounded by $\pi/2$; our sampling algorithm does not work without this condition (see, however, [10] where it is shown how to sample from any determinantal process with Hermitian kernel, of which μ_{LC} is one).

The paper is organized as follows. In Section 2, we introduce the sampling algorithm and prove its correctness. In Section 3, we introduce the probability measures on CRSFs on graphs on surfaces and show how they are exactly sampled by the algorithm. In Section 4, we show that the probability measures on loops that these induce converge to probability measures on the space of multiloops of the surface (this section contains the proof of our main statement). Section 5 enumerates some of the properties of the measures on loops considered in the paper. The paper concludes with a list of open questions in Section 6.

2. A general sampling algorithm. An *oriented* CRSF is a CRSF in which each cycle has a chosen orientation. A measure on CRSFs induces a measure on oriented CRSFs by giving each cycle an independent $1/2 - 1/2$-chosen orientation, and a measure on oriented CRSFs induces one on CRSFs by forgetting the orientation.

Let $\mathcal{G} = (V, E)$ be a finite graph with vertex and edge sets V and E, respectively, and $c : E \to \mathbb{R}_{>0}$ a positive function on the edges which we call the *conductance*. Let α be a function which assigns to each oriented simple loop γ in $\mathcal{G}$ a positive weight $\alpha(\gamma) \in [0, 1]$. We allow loops γ consisting of two edges (a backtrack: an edge which is immediately traversed in the reverse direction). These functions c, α define a probability measure $\mu = \mu_{c,\alpha}$ on oriented CRSFs, giving an oriented CRSF Γ a probability proportional to $\prod_{e \in \Gamma} c(e) \prod_{\text{cycles } \gamma \subset \Gamma} \alpha(\gamma)$. We describe an algorithm to sample an oriented CRSF according to the measure μ.

We note that this sampling algorithm requires $\alpha \in [0, 1]$; it will not work without modification for larger α. In the special case where $\alpha = 1$ and $c = 1$, the algorithm samples according to the uniform measure on oriented CRSFs.

Let us describe a cycle-popping procedure, named $P[w, \Gamma]$, which takes as arguments w a vertex and Γ an oriented subgraph of $\mathcal{G}$ not containing w, and outputs another oriented subgraph of $\mathcal{G}$ containing Γ and w. The procedure is the following: start at vertex w and perform a simple random walk (with each step proportional to the conductances) until it first reaches a vertex v which either belongs to Γ or is the first self-intersection of its path.

- If v is in Γ, then replace Γ by the union of Γ and the oriented path just traced by the random walk.
- If v is the first self-intersection, let γ denote the oriented cycle thus obtained, and sample a $\{0, 1\}$-Bernoulli random variable with success probability $\alpha(\gamma)$.
  - If the outcome is 1, then replace Γ by the union of Γ and the oriented path just traced by the random walk.
  - If the outcome is 0, erase the cycle that was just closed and continue to perform the random walk from v until it reaches Γ or self-intersects, in which case repeat the above instructions.

The algorithm, called $\mathcal{A}$, is then the following: start with Γ empty and w an arbitrary vertex, and perform $P[w, \Gamma]$. If the output Γ' is not a CRSF, take a new vertex $w' \notin \Gamma'$ and perform $P[w', \Gamma']$, and so on until the output contains all vertices of $\mathcal{G}$. Note that the output of $\mathcal{A}$ is an oriented CRSF and that we forget the information about the order of construction of the cycles.

THEOREM 1. *If $\alpha(\gamma) > 0$ for some γ, then the algorithm $\mathcal{A}$ terminates and its output is an oriented CRSF, sampled according to the measure μ.*

REMARK 2. Note that if in the above definition of $\mathcal{A}$ one starts with Γ equal to S, a distinguished set of vertices in $\mathcal{G}$, then the algorithm samples an oriented *essential CRSF*; see definition in Section 3.3 below.

PROOF OF THEOREM 1.   Following the proof of Wilson's algorithm [23], we construct an equivalent description, denoted $\mathcal{A}'$, of the algorithm $\mathcal{A}$.

Let us define $\mathcal{A}'$ in the following way. Consider, over each vertex $v \in V$, an infinite sequence $X^{(v)}$ of i.i.d. random variables $X^{(v)} = (X_1^{(v)}, X_2^{(v)}, \ldots)$, each distributed as a random neighbor of v according to the conductance measure, that is, for each $i \geq 1$ and each neighbor w of v, we have

$$\mathbb{P}(X_i^{(v)} = w) = \frac{c(vw)}{\sum_{v' \sim v} c(vv')}.$$

We represent $X^{(v)}$ as an infinite stack of cards, with only $X_1^{(v)}$ being visible at the top of the stack.

We draw an edge from each vertex v to the neighbor shown on the top of the stack $X^{(v)}$; the oriented graph thus seen is an oriented CRSF (with possible loops of length 2). This is our initial CRSF. We now describe a step by step random popping algorithm of the cycles. Note that at each step, the graph that we see remains an oriented CRSF. Here is the algorithm: For each cycle γ encountered in the current CRSF, pop it with probability proportional to $1 - \alpha(\gamma)$; when a cycle is popped off, the top card on the stacks for each of its vertices is discarded. When a cycle is "kept," its cards are fixed and can no longer be removed. The algorithm stops once the cycles that remain have been all "kept" in a Bernoulli trial. It is easy to see that the order in which the cycles are popped is not relevant.

Note that this algorithm terminates since there is at least one cycle γ with positive weight $\alpha(\gamma)$, because eventually, with probability 1, all the cycles present at one step will have been previously kept (the argument is the same as in Wilson's proof).

Since the cards of the stacks are distributed as the steps of a conductance-biased random walk, we see that algorithm $\mathcal{A}'$ has the same output in distribution as algorithm $\mathcal{A}$. In order to compute the output distribution of $\mathcal{A}$, we will therefore use algorithm $\mathcal{A}'$.

Let us compute the probability that a given oriented CRSF Γ is obtained as an output of algorithm $\mathcal{A}'$. Let $\gamma_1, \ldots, \gamma_k$ be the cycles of Γ. The CRSF Γ is obtained as an output if and only if there exists a finite sequence of oriented cycles $C_1, \ldots, C_m$ such that these cycles are popped, and after removing them, the cards that appear correspond to Γ, and there are k successful trials for Bernoullis with success probability $\alpha(\gamma_i)$.

By independence of the cards in the stacks, the last CRSF considered is independent of the cycles that were popped. Therefore, for any oriented CRSF Γ, we have

$$\mathbb{P}(\Gamma) = \sum_{\mathcal{C} = \{C_1, \ldots, C_m\}} \mathbb{P}(\Gamma | \text{pop } \mathcal{C}) \mathbb{P}(\text{pop } \mathcal{C})$$

$$= \sum_{\mathcal{C}} \mathbb{P}(\text{pop } \mathcal{C}, \text{ and } \Gamma \text{ occurs underneath and is kept})$$

$$= \sum_{\mathcal{C}} \mathbb{P}(\text{pop } \mathcal{C}) \prod_{e \in \Gamma} \mathbb{P}(e) \prod_{i=1}^{k} \alpha(\gamma_i)$$

$$= \left(\sum_{\mathcal{C}} \mathbb{P}(\text{pop } \mathcal{C}) \right) \prod_{e \in \Gamma} \mathbb{P}(e) \prod_{i=1}^{k} \alpha(\gamma_i)$$

which we see is proportional to the weight of Γ. $\square$

To sample a nonoriented CRSF according to a measure which assigns a CRSF Γ a weight proportional to $\prod_{e \in \Gamma} c(e) \prod_{\gamma \subset \Gamma} \alpha(\gamma)$, where the product is over nonoriented cycles γ, and α is a function invariant under orientation, it suffices to have $\alpha \in [0, 2]$, perform algorithm $\mathcal{A}$ for the measure $\mu_{c,\alpha/2}$, and forget the orientation in the resulting oriented CRSF. In particular, we obtain the uniform measure on nonoriented CRSFs with the choice $c = \alpha = 1$.

There is a variant of the previous algorithm to sample an oriented CRSF according to measure $\mu_{c,\alpha}$ conditional on having a single loop: multiply all the loop weights by a small constant ε. Then perform $\mathcal{A}$; if ε is small there will typically be a single loop (if not, start over).

Let N be the total number of vertices of the graph. The running time of the algorithm is bounded by the time to obtain the first loop [which is bounded by $O(N^2)$ if $\alpha \geq O(1/N^2)$] plus the running time of Wilson's algorithm, that is $O(N(\log N)^2)$ (Wilson's algorithm has a running time bounded by the cover time [23] which is linear up to a logarithmic correction [3]). The running time is at least linear. Extreme cases correspond to extreme values of α: for $\alpha = 1$ (uniform measure on CRSFs), the running time is linear; for α close to zero (like in the conditional measure described in the previous paragraph), the running time is large, at least $O(1/\sup \alpha)$.

3. Natural probability measures on CRSFs. The most natural probability measure on CRSFs on a finite unweighted graph is the uniform measure. If the edges are weighted with a real positive conductance function, then in this setting it is natural to give a CRSF a probability proportional to the product of its edge weights. We call this the *background measure*. There are, however, other natural probability measures that can be constructed from connections on bundles and that are meaningful for graphs embedded in surfaces.

3.1. *Connections.* Let $\mathcal{G} = (V, E)$ be a finite graph. A *vector bundle on $\mathcal{G}$* is a copy W_v of some fixed complex vector space W associated to each vertex $v \in V$. The *total space* of the bundle is the direct sum $\mathcal{W} = \bigoplus_{v \in V} W_v$. A *unitary connection* Φ on $\mathcal{W}$ is the data consisting of, for each oriented edge $e = vv'$, a unitary complex linear map $\varphi_{vv'} : W_v \to W_{v'}$ such that $\varphi_{v'v} = \varphi_{vv'}^{-1}$. The map $\varphi_{vv'}$ is referred to as the *parallel transport* from v to v'. We say that two connections Φ, Φ' are *gauge equivalent* if there exist unitary $\psi_v : W_v \to W_v$ such that

$\psi_{v'}\varphi_{vv'} = \varphi'_{vv'}\psi_v$, that is, Φ' is obtained from Φ by changing the basis of each space W_v by a unitary transformation. In this paper, we deal uniquely with vector bundles with $W = \mathbb{C}$ (line bundles) or $\mathbb{C}^2$, and $U(1)$- or SU(2)-connections, respectively.

Let $c : E \to \mathbb{R}_{>0}$ be a conductance function. We let Δ_Φ be the associated Laplacian acting on $f \in \mathcal{W}$ defined, for each vertex v, by

$$\Delta_\Phi(f)(v) = \sum_{v' \sim v} c(vv')(f(v) - \varphi_{v'v} f(v')),$$

where the sum is over all neighbors v' of v.

When $\mathcal{G}$ is geodesically embedded in a surface Σ with a Riemannian metric (i.e., embedded in such a way that edges are geodesic segments), there is a natural connection $\Phi = \Phi_\nabla$ on $\mathcal{G}$ arising from any unitary connection ∇ on a vector bundle on Σ: we define for each vertex v the space W_v to be the fiber over v; the ∇-parallel transports along edges e of $\mathcal{G}$ define the parallel transports φ_e, and thus the connection Φ.

The product of parallel transports along a closed path is called the *holonomy* of the connection along the path. For flat connections, it is also called the *monodromy*.

3.2. *Laplacian determinant and measures.*

THEOREM 3 ([8, 15]). *For a graph with unitary connection Φ on a line bundle, we have*

(1) $$\det(\Delta_\Phi) = \sum_{\text{CRSFs}} \prod_{\text{edges}} c(e) \prod_{\text{cycles}} (2 - 2\cos\theta),$$

where $e^{i\theta}$ is the holonomy of the connection around the cycle, for any choice of its orientation.

Associated to Φ is a probability measure μ_Φ on CRSFs, where the probability of a CRSF is proportional to $\prod_{\text{edges}} c(e) \prod_{\text{cycles}} (2 - 2\cos\theta)$. This measure exists as long as there is at least one cycle with $\theta \neq 0 \mod 2\pi$.

See Theorem 15 below for a generalization to $\mathbb{C}^2$-bundles with SU(2)-connection, where the weight $2 - 2\cos\theta$ is replaced by $2 - \operatorname{Tr} w$, with w denoting the holonomy of the connection around the cycle. Note that for an element $w \in$ SU(2) we have $2 - \operatorname{Tr} w = 2 - 2\cos\theta$ where $e^{\pm i\theta}$ are the eigenvalues of w. One can treat a line bundle connection with parallel transports $\varphi_e = e^{i\theta_e}$ as a special case of a SU(2)-connection with parallel transports which are diagonal matrices

$$\begin{pmatrix} e^{i\theta_e} & 0 \\ 0 & e^{-i\theta_e} \end{pmatrix}.$$

The measures on SU(2)-connections are used to analyze the measures of primary interest μ_{nonc}, μ_{LC} and μ_{LC^0} we discuss below.

3.3. *Graphs with wired boundary.* Let $B \subset \mathcal{G}$ be a subset of vertices which we consider to be the *wired boundary* of $\mathcal{G}$ or simply *boundary*. An *essential CRSF* on a graph with wired boundary is a subgraph, each of whose components is either a unicycle not containing any boundary vertex, or a tree containing a single boundary vertex. For a graph with connection and boundary, we define Δ_Φ to be the associated Laplacian acting on $f \in W$ defined, for each vertex $v \in \mathcal{G} \setminus B$, by

$$\Delta_\Phi(f)(v) = \sum_{v' \sim v} c(vv')(f(v) - \varphi_{v'v} f(v')),$$

where the sum is over all neighbors $v' \in \mathcal{G}$ of v (including neighbors in B). In the natural basis Δ_Φ is a submatrix (indexed by $\mathcal{G} \setminus B$) of the full Laplacian on $\mathcal{G}$. The analog of Theorem 3 above holds (see [15]) where the sum is over essential CRSFs.

3.4. *Flat connections.* A connection is *flat* if it has trivial holonomy around any contractible cycle. Suppose that $\mathcal{G}$ is geodesically embedded on a nonsimply connected surface Σ with flat connection ∇. Let $\Phi = \Phi_\nabla$ be the associated connection on $\mathcal{G}$. The associated measure μ_Φ gives zero weight to contractible cycles, and thus is supported on noncontractible CRSFs.

Let μ_{nonc} be the background measure on noncontractible CRSFs on $\mathcal{G}$ (giving a CRSF a probability proportional to the product of its edge weights, that is, ignoring any connection). The measure μ_Φ has density $\prod_{\gamma \subset \Gamma}(2 - \text{Tr}(w_\gamma))$ with respect to μ_{nonc}.

Although μ_{nonc} cannot itself be written as a connection measure μ_Φ for some flat connection Φ, we can use the μ_Φ to study μ_{nonc}; see Lemma 16 below.

3.5. *Graphs embedded on a curved surface.*

3.5.1. *The Levi-Civita measure.* Suppose that $\mathcal{G}$ is geodesically embedded on a Riemannian surface Σ. There is a natural complex line bundle on Σ, the tangent bundle. Take ∇ to be the Levi-Civita connection on the tangent bundle associated to a metric g on Σ. Define μ_{LC} to be the associated probability measure. It gives a CRSF a probability proportional to

$$\prod_e c(e) \prod_{\gamma \subset \Gamma}(2 - 2\cos\theta_\gamma),$$

where, by the Gauss–Bonnet theorem, θ_γ is the Gaussian curvature enclosed by γ. (If γ is not contractible it is the "net turning angle" of γ.)

3.5.2. *The CRST measure.* When Σ is contractible, there is another measure μ_{LC^0} we can associate to this situation, introduced in [15]. It is supported on CRSTs (one-component CRSFs) of $\mathcal{G}$. Let $\Phi = \{e^{i\theta_e}\}_{e \in E}$ be the parallel transports on $\mathcal{G}$ defined from ∇, and for real t let $\Phi_t = \{e^{it\theta_e}\}_{e \in E}$; these are well defined by

contractibility of Σ. Let μ_{LC^0} be the limit as $t \to 0$ of the measures μ_{Φ_t}. Since loop weights are going to zero, there will be only one loop remaining in the limit $t \to 0$, so the limit is a CRST. In μ_{LC^0}, each CRST Γ has a weight proportional to (the product of the edge weights times) θ_γ^2, the square of the curvature enclosed by the unique cycle γ of Γ.

Let $Z_{\mathrm{LC}^0} = \sum_{\mathrm{CRSTs}} \theta_\gamma^2 \prod_e c(e)$ be the partition function of μ_{LC^0}. By Theorem 3, we have

$$Z_{\mathrm{LC}^0} = \lim_{t \to 0} t^{-2} \det \Delta_{\Phi_t}.$$

An explicit computation of this limit yields the following. Let $\kappa = \kappa(\mathcal{G})$ be the weighted sum of spanning trees of $\mathcal{G}$.

LEMMA 4. *We have* $Z_{\mathrm{LC}^0} = \kappa(\mathcal{G}) \langle \Theta, (I - T)\Theta \rangle$, *where* $\langle \cdot, \cdot \rangle$ *is the usual scalar product in the space of* 1-*forms,* T *is the transfer current* (*defined below*), *and* Θ *is the one-form on edges giving the angle* θ *of the connection.*

This lemma also appears as Lemma 2 in [14], page 14, but we include the proof here for self-containedness.

The same statement (and proof) applies more generally to the $t \to 0$ limit starting from any $U(1)$-connection, not necessarily the Levi-Civita connection. However, we will not have use for these other measures here.

PROOF OF LEMMA 4.   The map T is the orthogonal projection onto $\mathrm{Im}(d)$, which is the orthocomplement of the space of $\mathrm{Ker}(d^*)$. By a (generalization to varying conductances of a) result of Biggs ([5], Proposition 7.3) the projection onto $\mathrm{Ker}(d^*)$ is given by

$$\kappa^{-1} \sum_{\text{spanning trees } t} wt(t) \sum_{e \notin t} c_e \pi_{C(e,t)},$$

where $wt(t) = \prod_{e' \in t} c_{e'}$, $C(e,t)$ is the cycle of $t \cup e$ and $\pi_{C(e,t)}$ is the projection onto this cycle. Hence,

$$T = I - \kappa^{-1} \sum_{\text{spanning trees } t} wt(t) \sum_{e \notin t} c_e \pi_{C(e,t)}.$$

We then have

$$\kappa \langle \Theta, (I - T)\Theta \rangle = \kappa \left\langle \Theta, \kappa^{-1} \sum_t wt(t) \sum_{e \notin t} c_e \pi_{C(e,t)} \Theta \right\rangle$$

$$= \sum_t wt(t) \sum_{e \notin t} c_e \langle \Theta, \pi_{C(e,t)} \Theta \rangle$$

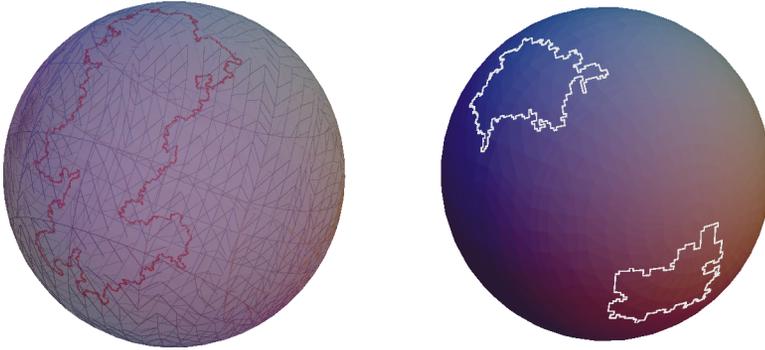$$= \sum_t wt(t) \sum_{e \notin t} c_e \frac{\theta_C^2}{|C(e,t)|}$$

FIG. 2. *A μ_{LC^0}-random CRSF on the sphere with its round metric*; *a μ_{LC}-random CRSF on the sphere with two components* (*a rare event*).

$$= \sum_{\text{CRSTs } u} wt(u)\theta_C^2,$$

where we used $\langle w, \pi_v w \rangle = \frac{\langle w, v \rangle^2}{\langle v, v \rangle}$ in the second from last equality.   □

3.6. *Exact sampling.* The measures μ_{nonc}, μ_{LC^0}, and μ_{LC} can be sampled using our generalization of Wilson's algorithm as follows.

The measure μ_{nonc} is sampled by using a function α which assigns a loop weight 0 if it is contractible and 1 otherwise.

For μ_{LC^0}, we set $\alpha(\gamma) = \varepsilon\theta_\gamma^2$ for small ε. For small enough ε, there will typically be only one loop (if there is more than one, start over).

We can sample μ_{LC} only in the case where the absolute value of the curvature θ enclosed by any curve does not exceed $\pi/2$. Indeed, in that case, we will always have $2 - 2\cos\theta \in [0, 2]$ which is necessary to sample.

See Figures 2, 3, 4 which are obtained by using fine conformal approximations to the underlying surfaces (only the cycles of the CRSFs are drawn). Figures 2, right and 3, right, are samples conditional on enclosing curvature less than $\pi/2$.

In order to sample the unconditional measures, one can use a general algorithm for determinantal processes with Hermitian kernels [10] which is slower.

Figure 5 shows a sample of μ_{nonc} on a multiply-connected planar domain.

4. Scaling limits for graphs on surfaces. The measures $\mu_{\text{nonc}}, \mu_{LC}, \mu_{LC^0}$ induce measures on sets of cycles on $\mathcal{G}$, by forgetting the rest of the CRSF. We use notation $\mathbb{P}_{\text{nonc}}, \mathbb{P}_{LC}, \mathbb{P}_{LC^0}$, respectively, for these cycle measures.

In this section, we prove our main statement (convergence of these cycle measures) in the following way. We first review in 4.1 the conformal approximation setup in which we consider the scaling limit. Then in 4.2, we define the probability space in which convergence takes place, namely the metric space of loops up to time-reparameterization. A main ingredient is obtained in 4.3 where we show that
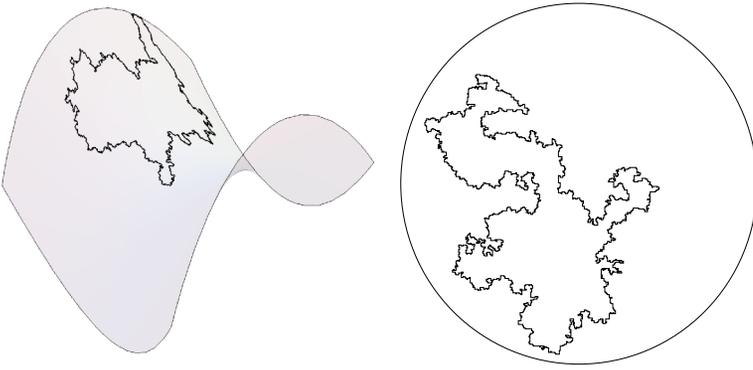
FIG. 3. *A μ_{LC}-random CRSF on the saddle surface $z = x^2 - y^2$ (left), and the hyperbolic plane* (*right*).

the probability that there be no loop in the (wired) CRSF measure has a nontrivial limit. This is instrumental in 4.4 to show that, in the limit, the number of loops remains finite and that the loops are necessarily macroscopic. Combined with earlier technical results on LERW by other authors, this implies tightness of the sequence of measures. We then conclude by Prokhorov's theorem, showing the convergence of a determining class of "cylindrical" events defined in 4.5. This is done first in 4.6 in the case of the measure $\mathbb{P}_{nonc}$ using the representation theory of the surface's
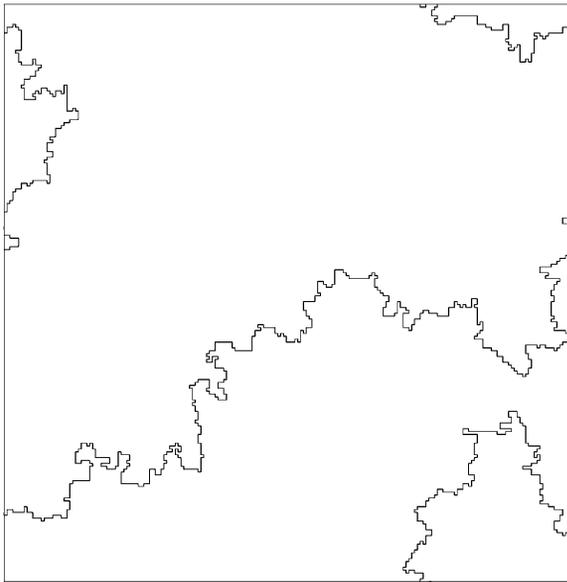

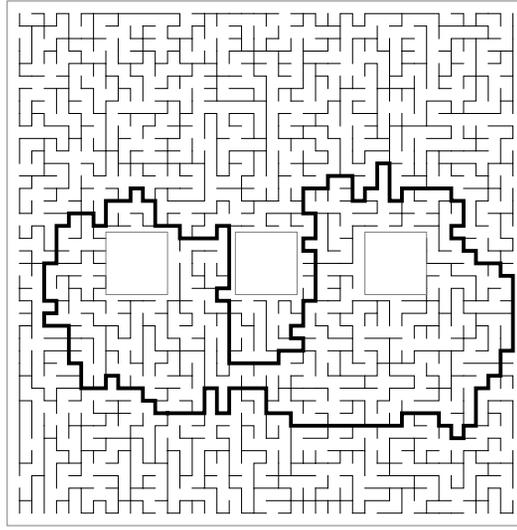FIG. 4. *A μ_{nonc}-random CRSF on the flat torus (obtained from the unit square by identifying opposite sides).*

FIG. 5.   *A μ_{nonc}-random CRSF on a punctured disk conditioned on having a particular homotopy type.*

fundamental group, then extended to flat connections, and finally used to prove the main result, Theorem 20 in 4.7. The last Section 4.8 presents some applications of the main theorem.

4.1. *Conformal approximation.*   Let $(\mathcal{G}_n)$ be a sequence of (edge-weighted) graphs geodesically embedded in Σ with mesh size (longest edge length) going to zero.

There are a number of equivalent definitions of the notion of conformal approximation of Σ by the sequence $(\mathcal{G}_n)$. Perhaps the easiest is to say that conductance-weighted random walk on $\mathcal{G}_n$ converges to the Brownian motion on Σ, up to reparameterization. Another approach, more computationally useful, uses the transfer impedance.

Let $R^n(e, e')$ be the transfer impedance of two oriented edges $e = xy$ and $e' = x'y'$ in $\mathcal{G}_n$. This is defined as the potential drop across e' when one unit of current enters at e_- and leaves at e_+ (the endpoints of e), when the graph is viewed as an electrical network with conductance c. In terms of the Green function G^n one has

$$R^n(e, e') = G^n(e_+, e'_+) - G^n(e_+, e'_-) - G^n(e_-, e'_+) + G^n(e_-, e'_-).$$

A related quantity is $T(e, e') = c(e')R(e, e')$ the *transfer current*: the current across e' when one unit of current enters at e_- and leaves at e_+.

The function $R^n(e, e')$, as a function of either e or e', is a discrete one-form on $\mathcal{G}_n$ (function on oriented edges which changes sign under change of orientation). We say that *conformal approximation holds* if the mesh size tends to zero as $n \rightarrow$

∞ and, when e_1, e_2 are not within $o(1)$ of each other,

$$(2) \qquad R^n(e_1, e_2) = \big(D_{e_1} D_{e_2} g(z_1, z_2)\big)\ell(e_1)\ell(e_2) + o\big(\ell(e_1)\ell(e_2)\big),$$

where g is the continuous Green function, $\ell(e_i)$ is the edge length, and the notation $D_{e_i} g$ represents the directional derivative in the direction of the edge e_i applied to the variable z_i, that is,

$$D_{e_1} g(z_1, z_2) = \lim_{\delta \to 0} \frac{g(z_1 + \delta u_1, z_2) - g(z_1, z_2)}{\delta},$$

where u_1 is the unit displacement in the direction of e_1, and similarly for D_{e_2} for the second variable z_2.

If we sum the transfer impedance for e_1 varying along a path $\gamma_{a,b}$ from vertex a to b, and e_2 on $\gamma_{c,d}$ from c to d, we find that

$$\sum_{e_1 \in \gamma_{a,b}} \sum_{e_2 \in \gamma_{c,d}} G^n\big(e_1^-, e_2^-\big) - G^n\big(e_1^-, e_2^+\big) - G^n\big(e_1^+, e_2^-\big) + G^n\big(e_1^+, e_2^+\big)$$

$$(3)$$
$$= g(a, c) - g(b, c) - g(a, d) + g(b, d) + o(1).$$

This quantity represents the change in potential from c to d when one unit of current enters at a and exits at b. As a function $F(c)$ of c, this is the unique harmonic function with appropriate logarithmic singularities at a and b, and which is zero at d. In particular, the level curves of F are equipotentials and, taking equipotentials near a and b (which are close to circles), one can thus compute the resistance between a small circle around a and a small circle around b. In this way, using convergence of the transfer impedance, one can show that the main notions of potential theory including the Poisson kernel, Cauchy kernel, holomorphic functions, *etc.* all behave well under conformal approximation.

As an example of a conformally approximating sequence, standard arguments show that the square grid in $\mathbb{C}$, scaled by ε, conformally approximates any planar domain in $\mathbb{C}$ as $\varepsilon \to 0$. Thus, for a simply connected domain in Σ we can pull back a fine square grid in $\mathbb{C}$ under a conformal map from a domain in $\mathbb{C}$ to Σ. More generally, the (almost-)equally-spaced real and imaginary curves of a holomorphic quadratic differential $\phi(z)\,dz^2$ on Σ give a graph structure on Σ, with unit conductances, conformally approximating the surface. Other examples from Poisson point processes are given in [9].

4.2. *The probability space of simple multiloops.* For a graph $\mathcal{G}_n$, let $\mathbb{P}^n$ be one of the measures on loops discussed above. We can see the probability measures $\mathbb{P}^n$ for different ns as living on the same probability space Ω that we now describe.

A *multiloop* in Σ is a finite union of simple loops (ignoring parameterization) with disjoint images, that is, an injective map from the union of k copies of the unit circle $\mathbb{T}$ to Σ, for some $k \geq 1$ (and modulo reparameterizations). The space of single simple loops is a metric space: the distance is defined by

$$d(f, g) = \inf_{\alpha} \sup_{t \in S^1} \big\| f(t) - g\big(\alpha(t)\big) \big\|,$$

where the infimum is over all reparameterizations α. In other words, two loops are close if they can be parameterized so that their images are close for all parameter values t. One can extend this distance to the case of multiloops by taking the infimum over all permutations of loops, of the max of the pairwise distances (and defining the distance between Ω_k and $\Omega_{k'}$ for $k \neq k'$ to be infinite). With this distance Ω is a topological metric space. It is a disjoint union $\Omega = \bigcup_{k=1}^{\infty} \Omega_k$ where Ω_k consists of multiloops with k components.

This space is not complete: it is easy to construct Cauchy sequences that shrink to a point or to nonsimple loops. However, it is separable: take all finite multiloops lying on fine lattice approximations of the surface. This is a countable family of loops which is dense in Ω.

The set of cycles of a CRSF on $\mathcal{G}_n$ defines an element of Ω.

A *finite lamination* on a surface is an isotopy class of a multiloop. For any points $x_1, \ldots, x_m \in \Sigma$, and small $\delta > 0$, let B_i be a ball around x_i of some radius less than δ, and consider the finite laminations in $\Sigma \setminus \{B_1 \cup \cdots \cup B_m\}$. Any multiloop γ which avoids the balls B_i defines a lamination $[\gamma]$. For any of these laminations L, we consider the event

$$E_{B_1, \ldots, B_m; L} = \{\gamma \in \Omega | [\gamma] = L\}.$$

We call these sets *cylindrical events* and consider the σ-field $\mathcal{B}$ on Ω generated by these events.

LEMMA 5. *$\mathcal{B}$ contains the Borel sets in Ω.*

PROOF. We just prove this for one loop, that is, for Ω_1. The proof is easily extended to the general case.

Let c be a smooth simple closed curve in Σ and for some small $\varepsilon > 0$ let $U_\varepsilon(c)$ be its ε-neighborhood. Consider the event E that the random curve γ maps into $U_\varepsilon(c)$, winding once around the annulus with, say, the positive orientation. These types of events generate the Borel sets.

Let $x_1, x_2, \ldots$ be a sequence of points dense in the boundary of $U_\varepsilon(c)$. Let B_i be a ball around x_i of radius δ/i. Let $E_{\delta;n}$ be the cylinder event that the curve γ in $\Sigma \setminus \{B_1 \cup \cdots \cup B_m\}$ separates the points $\{x_1, \ldots, x_n\}$ on the two boundary components of $U_\varepsilon(c)$, that is, is consistent with γ winding once around $U_\varepsilon(c)$. The event E is contained in the intersection over n of the $E_{\delta;n}$. In fact, we have $E = \cap_n E_{0;n} = \bigcup_{\delta \to 0} \cap_n E_{\delta;n}$: any continuous simple loop which separates the points lies in the interior of the annulus and winds once around. This can be seen as follows. First of all, any curve in $\cap_n E_{0;n}$ is contained in $U_\varepsilon(c)$: otherwise, we could find some x_j lying on the wrong side of the curve since the family $(x_n)_{n \geq 1}$ is dense in the boundary which would yield a contradiction. Second, the curve cannot be contractible since this would contradict the fact that it separates the points. Hence, it winds once around the hole of the annulus. Its orientation is necessarily positive. □

LEMMA 6. *The set of cylindrical events forms a determining class for $\mathcal{B}$.*

PROOF. This class is stable under finite intersection and generates the sigma-field. □

4.3. *The probability of no loops.* Recall that on a graph with boundary, an essential CRSF is a subgraph, each of whose components is either a unicycle not containing any boundary vertex, or a tree containing a single boundary vertex. For a graph with boundary and flat connection Φ, the probability that a μ_Φ-random essential CRSF has no loops tends to 1 as the holonomy tends to the identity, since each noncontractible cycle has weight $2 - 2\cos\theta$ which tends to zero.

For μ_{LC}, the following theorem describes a similar result. For a graph embedded on a surface Σ with boundary, we define the boundary of the graph to be the set of vertices adjacent to the boundary of Σ. Let $\{\mathcal{G}^n\}_{n=1,2,\ldots}$ be a sequence of graphs conformally approximating Σ, where $\mathcal{G}^n$ has mesh size at most $1/n$.

The intuition behind the following theorem is that one may express the quantities as functionals of the random walk loop soup and observe that the random walk loop soup converges to the Brownian one; see, for example, [19].

THEOREM 7. *For any compact surface with nonempty boundary and any smooth unitary connection ∇ on a line bundle on Σ, the probability $\mathbb{P}^n$(no loops) converges to a universal limit $p(\Sigma)$.*

PROOF. Let Δ_{Id} be the Laplacian on $\mathcal{G}^n$ for the trivial connection. We have

$$\mathbb{P}^n(\text{no loops}) = \frac{\det \Delta_{\mathrm{Id}}}{\det \Delta_\Phi}.$$

Write $\Delta_\Phi = D(I - P_\Phi)$ where D is the diagonal "degree" (sum of weights) matrix. Then

$$-\log \mathbb{P}^n(\text{no loops}) = -\mathrm{Tr}\log(I - P_{\mathrm{Id}}) + \mathrm{Tr}\log(I - P_\Phi)$$

(4)
$$= \sum_{k=1}^{\infty} \frac{\mathrm{Tr}\, P_{\mathrm{Id}}^k}{k} - \frac{\mathrm{Tr}\, P_\Phi^k}{k}$$

$$= \sum_{\text{loops}\,\ell} \mathbb{P}(\ell)(1 - \cos\theta_\ell),$$

where the sum is over unrooted loops ℓ, and where $\mathbb{P}(\ell)$ is the probability of ℓ (for the weighted random walk started at some vertex of ℓ) and $e^{i\theta_\ell}$ is the holonomy of the loop ℓ.

We divide this sum into three parts: tiny loops (consisting of at most $\varepsilon\sqrt{n}$ steps), small loops (consisting of at most εn^2 steps), and large loops (at least εn^2 steps).

For tiny loops of k steps, the area enclosed is at most $O(k^2/n^2)$ by the isoperimetric inequality. Since θ_ℓ is at most a constant times the enclosed area, the contribution for tiny loops is bounded by

$$\sum_{k=1}^{\varepsilon\sqrt{n}} \frac{\mathrm{Tr}(P^k)}{k} O\big((k^2/n^2)^2\big).$$

Since P is substochastic, $\mathrm{Tr}(P^k) = O(n^2)$, and the total contribution is $O(\varepsilon^2)$.

For small loops of length k, we argue that they enclose signed area $O(k/n^2)$. We use a small generalization of the following result of Wehn, [22]: a two-dimensional simple random walk of length k, conditioned to return to the origin, encloses a signed area (that is, the integral of $x\,dy$) of order $O(N)$, that is, has standard deviation $O(N)$. The argument is as follows. The signed area is obtained from an integral $\int x\,dy$ along the loop [where (x, y) are local orthogonal coordinates with origin at the starting point]. Each step makes an essentially independent contribution to the integral (the only dependence being the condition that the loop is closed after k steps). By grouping several steps at a time, the mean contribution for each group is zero, whereas the variance is of order k/n^4, since $x = O(\sqrt{k}/n)$ and $dy = O(1/n)$. Summing the variance over the k steps, one gets a total variance $O(k^2/n^4)$. Using the fact that a random walk of length k returns to its starting point with probability $O(1/k)$, the contribution for small loops is

$$\sum_{k=1}^{\varepsilon n^2} \frac{\mathrm{Tr}(P^k)}{k} O\big(k^2/n^4\big) = \sum_{k=1}^{\varepsilon n^2} O\left(\frac{n^2}{k^2}\right) O\big(k^2/n^4\big) = O(\varepsilon).$$

By [19], the sum over large loops converges to the analogous continuous loop measure. This is because long loops can be approximated with Brownian excursions. Along such an excursion the parallel transport is approximated by the Brownian parallel transport.[3] This Brownian parallel transport is a universal quantity in the sense of being independent of the underlying graph, only depending on the underlying smooth connection. □

Suppose now ∇ is close to the identity: that is, for some small $c > 0$, the integral of the absolute value of the curvature of ∇ is less than c, and the holonomies w of ∇ on a fixed cycle basis satisfy $|w - 1| < c$.

---

[3]The *Brownian parallel transport* is defined as a limit of the parallel transport along piecewise geodesic approximations to the Brownian motion: for a Brownian path $B(t)_{t\in[0,1]}$, take, for example, for N large the piecewise geodesic path connecting the points $B(j/N)$ and $B((j+1)/N)$ for $j \in \{0, N-1\}$. Itô showed in [11] that almost surely the parallel transports along the piecewise geodesic path converge in the limit $N \to \infty$ to a quantity independent of the approximating discrete path.

COROLLARY 8. *For ∇ close to the identity in the above sense, the probability of no loops satisfies*

$$\mathbb{P}^n(\text{no loops}) = 1 + o(1),$$

where the error is independent of n.

PROOF. In the proof of the above theorem, with a total curvature bound c, the term $1 - \cos\theta_\ell$ is of order $c^2 O(\theta_\ell^2)$. Thus, the ratio of determinants, and hence the probability of no loops, is $1 - O(c^2)$. □

4.4. *Asymptotic size of loops and tightness of the measures.*

4.4.1. *Macroscopic loops.* For $\mathbb{P}^n_{\text{nonc}}$, the loops necessarily are macroscopic since they are noncontractible. In the curved case for the measures $\mathbb{P}^n_{\text{LC}}$ and $\mathbb{P}^n_{\text{LC}^0}$, we show that there are, with positive probability, macroscopic loops.

THEOREM 9. *With positive probability $\mathbb{P}^n_{\text{LC}}$ and $\mathbb{P}^n_{\text{LC}^0}$ contain a macroscopic loop, that is, for sufficiently small $\varepsilon > 0$ the probability that there is a loop with diameter $\geq \varepsilon$ does not tend to zero with n.*

PROOF. We can suppose the surface is a disk: if not, take a point of nonzero curvature on the surface and a disk around it, small enough so that the curvature is roughly constant on the disk. Wiring the boundary of this disk, and making the disk boundary part of the surface boundary (effectively cutting the surface apart along the disk boundary) *decreases* the probability of finding a macroscopic loop in this disk, since the loop-erased walk from any interior point will halt sooner. So once we prove the theorem for a disk with nonzero curvature we are done.

We will prove the theorem for $\mathbb{P}^n_{\text{LC}^0}$. Since we are assuming at this point that the surface is a disk, this will also suffice for $\mathbb{P}^n_{\text{LC}}$ by the following argument. Choose the disk small enough so that the probability of no loops is at least $1/2$. Then for each loop discovered by the algorithm, the conditional probability of having no further loops is at least $1/2$ (since each new loop discovered decreases the probability of further loops). Thus, the number of loops is smaller than a geometric random variable of rate $1/2$, and the probability of 1 loop strictly dominates the probability of more than one loop. The change of loop weight between $\mathbb{P}^n_{\text{LC}}$ and $\mathbb{P}^n_{\text{LC}^0}$ tends to 1 with small curvature θ, so these are absolutely continuous with a Radon–Nikodym derivative independent of n.

Now consider the case of $\mathbb{P}^n_{\text{LC}^0}$. For this, consider first the $n \times n$ grid H_n scaled by $1/n$ to the unit square $[0, 1]^2$. Let $z_1 \neq z_2 \in (0, 1)^2$ and let f_1, f_2 be faces of the grid close to z_1, z_2, respectively. In [13], it is shown that the probability $\mathbb{P}(f_1, f_2)$ that f_1, f_2 are enclosed in the cycle of a uniform CRST of H_n is $\text{const}/n^2 + o(1/n^2)$. Let E be the event that f_1, f_2 are enclosed in the cycle. On this event,

the area of the enclosing cycle is with high probability $\geq \varepsilon n^2$ for some $\varepsilon > 0$ (see below, where it is shown that the loops are absolutely continuous with SLE_2). Then

$$(5) \qquad \mathbb{P}^n_{LC^0}(E) = \frac{\sum_E \text{Area}^2}{\sum_{CRSTs} \text{Area}^2} \geq \frac{\varepsilon^2 n^4 \sum_E 1}{Cn^2 \sum_{CRSTs} 1} = \varepsilon^2 C' n^2 \mathbb{P}(f_1, f_2)$$

for constants C, C'. Here, the denominator of the central inequality follows from the result of [13] that the second moment of the area of a uniform CRST is of order n^2 times a constant. The right-hand side of (5) is bounded below by a positive constant independently of n.

A similar argument holds for any sequence of graphs $(\mathcal{G}_n)$ conformally approximating a curved disk Σ, and near a point where the curvature of the metric is nonzero: In a small neighborhood U of such a point, the curvature is approximately constant and the transfer impedance for $\varepsilon \Phi$ will thus agree to first order with that in a similar small neighborhood of a point in H_n and far from the boundary of H_n. Since $\mathbb{P}(f_1, f_2)$ is determined by the Green's function on the dual graph [13], for $f_1, f_2 \in U$ we have a similar estimation of $\mathbb{P}^n_{LC^0}(E)$ as in the case for H_n above. $\qquad \square$

4.4.2. *Number of loops.* We show here that the number of loops has super-exponential decay for both $\mathbb{P}^n_{\text{nonc}}$ and $\mathbb{P}^n_{LC}$.

LEMMA 10. *For any sequence $(\mathcal{G}_n)_{n \geq 1}$ of graphs conformally approximating a compact Riemannian surface Σ, possibly with boundary, we have the following. For any $0 < \xi < 1$, there exists N and K such that for all $n \geq N$ and for all $k \geq K$, we have*

$$\mathbb{P}^n(\text{there are at least } k \text{ loops}) \leq \xi^k.$$

PROOF. Each loop created during the performance of the cycle-popping algorithm either disconnects the surface or decreases the rank of the first homology. By the Markov property (Section 5.1), the law of the conditional CRSF is obtained by independently sampling in each of the connected components.

For $\mathbb{P}^n_{LC}$, suppose that we have created k closed curves. Each of the complementary components is either planar or non-planar. Each new loop found has a positive probability of being macroscopic, that is, either reduces the rank of the first homology or removes definite area from both resulting components, by Theorem 9. Thus, for k sufficiently large, the curvature enclosed c on either side will be small, hence with definite probability $1 - p = 1 - o(1)$ the two sides will contain no further loops, by Corollary 8. By taking k large enough, and mesh size $1/n$ small enough, the curvature enclosed c can be taken as small as needed such that $p < \xi$.

Thus, the probability of an extra loop eventually decays exponentially with rate less than ξ.

A similar argument works for $\mathbb{P}_{\text{nonc}}$. $\quad \square$

As an example, the distribution of loops for $\mathbb{P}_{\text{nonc}}$ with free boundary conditions was computed for an annulus using the standard square grid approximation in [15]. For a cylinder of aspect ratio (height to circumference) τ, the probability generating function of the number of loops is an elliptic function

$$\mathbb{E}_\tau^{\text{free}}(X^{\#\text{ loops}}) = X \prod_{j=1}^{\infty} \frac{X + 2\cosh(\pi j/\tau) - 2}{2\cosh(\pi j/\tau) - 1},$$

which can be checked to have a super-exponential tail. For wired boundary conditions, and by planar duality (since the dual of an essential incompressible CRSF is a CRSF with free boundary conditions) the distribution is $\mathbb{E}_\tau^{\text{wired}}(X^{\#\text{ loops}}) = X^{-1}\mathbb{E}_\tau^{\text{free}}(X^{\#\text{ loops}})$, which naturally also has a super-exponential tail.

4.4.3. *Microscopic loops.* The following theorem shows that loops of $\mathbb{P}_{\text{LC}}^n$ do not shrink to points as $n \to \infty$.

THEOREM 11. *Any subsequential limit of $\mathbb{P}_{\text{LC}}^n$ as $n \to \infty$ is supported on Ω, that is, as $\alpha \to 0$*

(6) $$\limsup_{n\to\infty} \mathbb{P}_{\text{LC}}^n(\text{there is a loop of area} \leq \alpha) = o(1).$$

PROOF. Let us argue by contradiction. If we suppose that (6) is false, it means that there exists $p > 0$ and arbitrarily small values of $\delta > 0$ such that for n large enough, we have

$$\mathbb{P}_{\text{LC}}^n(\text{there is a loop with area} \leq \delta^2) > p.$$

A small-area loop has a small diameter; conditioning on this loop in particular does not change the transfer impedance operator far from that loop. In fact from Lemma 12 below, the loops are absolutely continuous with respect to the loop-erased walk; as such their probability of intersecting a small disk tends to zero with the disk's diameter, since the same is true for the loop-erased walk. Thus, the conditioning has negligible effect.

So we can expect to find many loops: there is a $0 < q < p$ such that, dividing Σ into regions of diameter $\sqrt{\delta}$, the probability that there is a loop in each of these $\Theta(1/\delta)$ regions is on the order of $q^{1/\delta}$. Thus,

$$\mathbb{P}_{\text{LC}}^n(\text{there are } 1/\delta \text{ loops with area} \leq \delta^2) > q^{1/\delta}.$$

However, by Lemma 10 above, there exists $0 < \xi < q$ such that for any k and n large enough, we have

$$\mathbb{P}_{\text{LC}}^n(\text{there are} \geq k \text{ loops}) \leq \xi^k.$$

By taking $k = 1/\delta$, we obtain that for arbitrarily small values of $\delta > 0$, there is a large enough n such that

$$0 < q^{1/\delta} < \mathbb{P}_{\text{LC}}^n(\text{there are } 1/\delta \text{ loops with area} \leq \delta^2) \leq \xi^{1/\delta}.$$

This yields a contradiction since the right-hand side of the last equation tends to zero faster than the left-hand side when $\delta \to 0$. $\square$

4.4.4. *Resampling and tightness.* We will show tightness of the sequence of measures $\mathbb{P}^n_{\mathrm{nonc}}, \mathbb{P}^n_{\mathrm{LC}^0}$ and $\mathbb{P}^n_{\mathrm{LC}}$. This will yield the existence of subsequential limits.

We show in fact that the scaling limits of the macroscopic loops are absolutely continuous with respect to SLE_2. For this, we establish the link to the loop-erased random walk (LERW).

LEMMA 12 (Link to LERW). *Let η be a simple path from vertices η^+ to η^-. The law of the cycle γ of a uniform CRST, conditional on the fact that it contains η, is (η followed by) the LERW from η^+ to η^- with wired boundary conditions along η.*

For the measures $\mathbb{P}_{\mathrm{LC}}, \mathbb{P}_{\mathrm{LC}^0}$ or $\mathbb{P}_{\mathrm{nonc}}$, conditional on all other components, the law of γ is asymptotically that of (η followed by) the LERW from η^+ to η^- with wired boundary conditions along η, biased by the cycle weight of γ.

PROOF. Let γ be a loop in a $\mathbb{P}^n_{\mathrm{LC}^0}$-, $\mathbb{P}^n_{\mathrm{LC}}$- or $\mathbb{P}^n_{\mathrm{nonc}}$-random CRST, and $a, b \in \gamma$ distinct points on it. Let $\gamma[a, b]$ be the part of γ counterclockwise between a and b, and $\gamma[b, a]$ the complementary part. If we erase $\gamma[a, b]$, we can define a new loop γ' by taking a LERW from a to b in the domain defined by $\Sigma \cup \gamma[b, a]$, with wired boundary conditions on $\gamma[b, a]$, and conditioning on the LERW to end at b. The union of this LERW and $\gamma[b, a]$ is the simple closed curve γ'. By the sampling algorithm, this curve γ' is absolutely continuous with $\gamma[a, b]$, with Radon–Nikodym derivative given by the ratio of the cycle weights. If a and b are close to each other, the LERW from a to b will with high probability not exit a small ball around $[a, b]$. Hence, the cycle weight of γ' will be close to that of γ.

Wilson's algorithm thus shows that this is a fair sample of $\mathbb{P}$ conditioned on $\gamma[b, a]$. Thus, the LERW from a to b with the appropriate boundary conditions is absolutely continuous with respect to $\gamma[a, b]$, with a bound independent of mesh size $1/n$. □

This proves that in the scaling limit, for any converging subsequence, the loops are locally absolutely continuous with respect to the scaling limit of LERW, which was shown in [18] to be SLE_2.

In particular the scaling limit, for any converging subsequence, is supported on simple curves (recall that SLE_2 curves are simple, and note that this is a local property). This tightness property of LERW was proved in [1, 2, 18, 21].

THEOREM 13. *The sequences $(\mathbb{P}^n_{\mathrm{nonc}})_{n \geq 0}$ and $(\mathbb{P}^n_{\mathrm{LC}})_{n \geq 0}$ are tight on Ω.*

The above argument also shows that the sequence $(\mathbb{P}^n_{\mathrm{LC}^0})_{n \geq 0}$ is tight, provided we allow for the possibility that the curve shrinks to a point. We need to add to Ω_1 a copy of Σ whose points represent the constant maps of S^1 to that point. The metric naturally extends to this augmented space Ω_1^*, so that the limit of a

sequence of curves shrinking to a point is the constant curve at that point. Let Ω_1^* be this augmented space.

THEOREM 14.   *The sequence* $(\mathbb{P}_{LC^0}^n)_{n \geq 0}$ *is tight on* Ω_1^*.

4.5. *Probabilities of cylindrical events.*

4.5.1. *Flat connections.*   Let Σ be a compact surface with $b > 0$ boundary components. Let $x_1, \ldots, x_k \in \Sigma$. Let $\mathcal{M}$ be the space of flat SU(2)-connections modulo gauge transformations on $\Sigma \setminus \{x_1, \ldots, x_k\}$. This space is compact. Such a flat connection is determined uniquely by a homomorphism from $\pi_1(\Sigma \setminus \{x_1, \ldots, x_k\})$ into SU(2) modulo conjugation.

Provided $b + k > 0$, the fundamental group $\pi_1(\Sigma \setminus \{x_1, \ldots, x_k\})$ is a free group F_m on $m = 2g + k - 1 + b$ letters, where g is the genus and b the number of boundary components of Σ. Thus, a flat connection is determined by m arbitrary elements of SU(2), one for each generator of π_1.

There is no canonical basis for $\pi_1(\Sigma \setminus \{x_1, \ldots, x_k\})$, and hence for $\mathcal{M}$. For any choice of a basis, we consider the measure on $\mathcal{M}$ obtained by the image of the Haar measure on SU(2)^m. It can be seen, using a theorem of Nielsen (see [20], Chapter 1), that the measure is independent of that choice of basis. (This theorem states that one can go from one basis to another in a free group by a sequence of elementary moves. It is easy to check that these moves preserve the Haar measure.) We let ν be this measure and call it the canonical Haar measure on $\mathcal{M}$.

4.5.2. *Trivalent graph.*   A useful device (see [7]) is to define a trivalent graph H in Σ (unrelated to $\mathcal{G}$) with a single boundary component of Σ or x_i in each face, so that H is a deformation retract of $\Sigma \setminus \{x_1, \ldots, x_k\}$. We can think of Σ as a ribbon graph structure on H.

Recall that a *finite lamination L* of Σ is an isotopy class of a finite number of disjoint noncontractible simple closed curves. A lamination L retracts to a "multicurve with multiplicity" on H; L is determined by, for each edge of H, a nonnegative integer giving the number of strands of a minimal representative of L retracted onto that edge. These integers satisfy the conditions that at each vertex of H the sum of the three integers is even and the three integers satisfy the triangle inequality; see Figure 6. Moreover, any set of integers satisfying these two conditions arises from a unique lamination.

We define a partial order on laminations: $L \leq L'$ if the integers on edges of H associated to L are all less than or equal to those of L'. We define the *complexity* $n(L)$ of L to be the sum of these integers.

An SU(2)-connection on H determines a flat connection on Σ. After gauge transformation, one can take the SU(2)-connection to be the identity on all edges of a spanning tree of H.
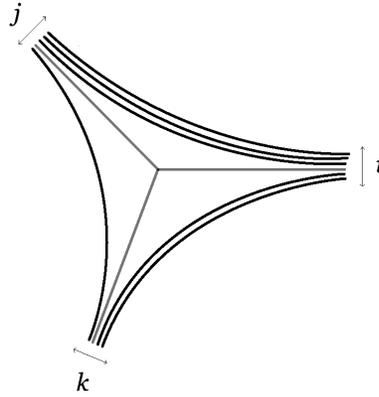
FIG. 6.   *Strands of the lamination at a vertex of the trivalent graph; the integers i, j, k (here $5, 4, 3$) satisfy $i + j + k = 0 \bmod 2$ and the triangle inequality.*

4.5.3. *Integrals over the space of flat connections.*   Given a finite weighted graph $\mathcal{G}$ embedded in Σ, and $\Phi \in \mathcal{M}$ define

$$(7) \qquad Z(\Phi) = \sum_{\text{CRSFs}} \prod_{\text{edges}} c(e) \prod_{\text{cycles } \gamma} \big(2 - \mathrm{Tr}(\omega_\gamma)\big),$$

where ω_γ is the holonomy of the connection Φ around the cycle γ. The function Z is a real-valued function on $\mathcal{M}$.

We denote

$$Z_0 = \sum_{\text{nonc. CRSFs}} \prod_{\text{edges}} c(e),$$

the partition function for all noncontractible CRSFs, without cycle weight. [This is *not* the same as $Z(\mathrm{Id})$, which is zero.]

For a flat connection Φ, we may rewrite (7) as

$$Z(\Phi) = \sum_L X_L T_L(\Phi),$$

where the sum is over finite laminations L, where X_L is the conductance-weighted sum of CRSFs whose cycles are isotopic to L, and

$$T_L(\Phi) = \prod_{\gamma \in L} \big(2 - \mathrm{Tr}(\omega_\gamma)\big)$$

is a real-valued function on $\mathcal{M}$.

Fock and Goncharov [7] proved that, seen as real-valued functions on $\mathcal{M}$, the functions $S_L = \prod_{\gamma \in L} \mathrm{Tr}(\omega_\gamma)$ are linearly independent and generate the vector space of regular (polynomial) functions on $\mathcal{M}$, when L runs through all finite laminations. Hence, the functions $T_L = \prod_{\gamma \in L}(2 - \mathrm{Tr}(\omega_\gamma))$ are also linearly independent, generate the same vector space and, when the bases are ordered by

increasing number of cycles, the change of basis matrix $M_{S,T}$ is an invertible infinite triangular matrix.

Choose an ordering of the T_L consistent with the partial order on laminations defined above. Let $\mathbf{P} = \{P_L(\Phi)\}$ be the Gram–Schmidt orthonormalization of the $\mathbf{T} = \{T_L(\Phi)\}$ with respect to this ordering and with respect to the inner product $\langle f, g \rangle = \int_{\mathcal{M}} fg \, d\nu$. Let $\mathbf{A} = (A_{L,L'})$ be the infinite lower-triangular matrix such that $\mathbf{P} = \mathbf{A}T$.

Recall the linear operator Δ_Φ on the total space $\mathbb{C}^{2|V|}$ of the $\mathbb{C}^2$-bundle on $\mathcal{G}$.

THEOREM 15 ([15]).   *We have*

$$Z(\Phi) = \sqrt{\det(\Delta_\Phi)}.$$

One can extract the coefficients of any desired lamination L as follows.

LEMMA 16.   *For any cylindrical event E_L, we have*

$$\mu_{\mathrm{nonc}}(E_L) = \sum_{L' \geq L} A_{L',L} \int_{\mathcal{M}} \frac{Z(\Phi)}{Z_0} P_{L'}(\Phi) \, d\nu.$$

This sum is finite for any finite graph.

PROOF OF LEMMA 16.   The probability of E_L is X_L/Z_0. Write

$$Z(\Phi) = \sum_L X_L T_L = \mathbf{X}^t \mathbf{T}(\Phi) = \mathbf{X}^t \mathbf{A}^{-1} \mathbf{P}(\Phi).$$

Since $\mathbf{P}$ is orthonormal, we have $\int_{\mathcal{M}} \mathbf{P}^t \mathbf{P} \, d\nu = \mathrm{Id}$. Hence,

$$\int_{\mathcal{M}} Z\mathbf{P}^t \mathbf{A} \, d\nu = \int_{\mathcal{M}} \mathbf{X}^t \mathbf{A}^{-1} \mathbf{P}(\Phi) \mathbf{P}^t(\Phi) \mathbf{A} \, d\nu$$

$$= \mathbf{X}^t \mathbf{A}^{-1} \left(\int_{\mathcal{M}} \mathbf{P}(\Phi) \mathbf{P}^t(\Phi) \, d\nu \right) \mathbf{A} = \mathbf{X}^t.$$

Hence,

$$\mathbf{X} = \int_{\mathcal{M}} Z\mathbf{A}^t \mathbf{P} \, d\nu.$$

Each entry X_L is the integral

$$X_L = \sum_{L' \geq L} A_{L',L} \int_{\mathcal{M}} Z(\Phi) P_{L'}(\Phi) \, d\nu.$$

Dividing by Z_0 we obtain the result.   □

4.6. *Convergence in the flat case.* We first consider Φ to be a flat connection. Associated to this is a measure μ_Φ^n on noncontractible CRSFs. Since μ_Φ^n has a density (independent of the graph) with respect to μ_{nonc}^n, it suffices to show that this latter converges.

The main tool is the following convergence result. Let $x_1, \ldots, x_k$ be points of Σ and B_j a small ball around x_j. Let Φ be a flat connection on $\Sigma \setminus \{B_1 \cup \cdots \cup B_k\}$.

THEOREM 17. *There exists a function $H \in L^2(\mathcal{M}^2)$ depending only on the conformal type of the surface $\Sigma' = \Sigma \setminus \{B_1 \cup \cdots \cup B_k\}$ such that for any Φ' not gauge-equivalent to the identity, we have*

$$\frac{Z(\Phi)}{Z(\Phi')} \to H(\Phi, \Phi').$$

There exists a bounded function F on $\mathcal{M}$ depending only on the conformal type of the surface $\Sigma' = \Sigma \setminus \{B_1 \cup \cdots \cup B_k\}$ such that

$$\frac{Z(\Phi)}{Z_0} \to F(\Phi).$$

PROOF. The first statement is proved in the same way as Theorem 7.

For the second statement, let $Z_0 = \sum_L X_L$ be the total number of CRSFs. By the first statement, $\frac{Z_0}{Z_\Phi}$ converges for any $\Phi \neq 1$. Now consider its inverse.

We write

(8) $$\frac{Z_\Phi}{Z_0} = \sum_L \frac{X_L}{Z_0} T_L.$$

First, $T_L \leq e^{O(n(L))}$ because $2 - \text{Tr}(\omega)$ is uniformly bounded by a constant over $\mathcal{M}$. Second, by Lemma 10, $X_K \leq e^{-cn(K)} Z_0$ for arbitrarily small $c > 0$. Since there are at most $n(K)^M$ laminations L with complexity $n(L) = n(K)$, where M is the number of edges of H, the sum (8) is bounded by a convergent series. $\quad\square$

THEOREM 18. *Let Σ be a compact nonsimply connected Riemann surface. There is a conformally invariant measure $\mathbb{P}_{\text{nonc}}$ on $(\Omega, \mathcal{B})$ supported on noncontractible multicurves, such that for any sequence $(\mathcal{G}_n)_{n\geq 1}$ of graphs conformally approximating Σ, the measures $\mathbb{P}_{\text{nonc}}^n$ on noncontractible CRSFs of $\mathcal{G}_n$ converge to $\mathbb{P}_{\text{nonc}}$.*

The main result of [16] shows that the homotopy classes on Σ of the noncontractible loops have a conformally invariant limit distribution.

PROOF OF THEOREM 18. Take points $z_1, \ldots, z_k$ in Σ, take $\delta > 0$ small, and for each i let B_i be the ball of small radius δ around z_i. Let $\mathcal{G}_B^n = \mathcal{G}^n \setminus \{B_1 \cup \cdots \cup B_k\}$ and $\mathbb{P}_B^n$ the associated measure on multiloops of noncontractible CRSFs on $\mathcal{G}_n$ whose loops stay in $\mathcal{G}_B^n$.

Consider L a finite lamination in $\Sigma \setminus \{B_1 \cup \cdots \cup B_k\}$, which has no peripheral curves (no curves isotopic to one of the boundary curves ∂B_i). It can also be thought of as a lamination of $\mathcal{G}_B^n$. Let E_L be the event that a CRSF of $\mathcal{G}_B^n$ has lamination L.

Up to errors uniform in n and tending to zero with δ,

$$\mathbb{P}^n(E_L | \text{no peripheral curves}) = \mathbb{P}_B^n(E_L | \text{no peripheral curves}).$$

[Here, the term on the left is equal to $\mathbb{P}^n(E_L)$ because the connection is flat.] This follows from the sampling algorithm since removing one or more very small disks does not change the distribution of the LERW away from those disks.

We need to show that $\lim_{n \to \infty} \mathbb{P}_B^n(E_L)$ exists and depends only on the conformal type of the domain $\Sigma \setminus \{B_1 \cup \cdots \cup B_k\}$.

By Lemma 16, the probability $\mathbb{P}_B^n(E_L)$ is given by a sum over $L' \geq L$ of integrals over $\mathcal{M}$ of $Z(\Phi)/Z_0$ times a function $P_{L'}(\Phi)$ independent of n. By Corollary 17, the integrand $Z(\Phi)/Z_0$ converges and is bounded independently of n. Thus, by bounded convergence, for each L' the integral $\langle \frac{Z(\Phi)}{Z_0}, P_{L'} \rangle = \int_{\mathcal{M}} \frac{Z(\Phi)}{Z_0} P_{L'} \, d\nu$ converges.

We need to show that the sum (weighted by the coefficient $A_{L,L'}$) over L' converges. We have

$$(9) \qquad \langle Z(\Phi), P_{L'} \rangle = \left\langle \sum_K X_K T_K, P_{L'} \right\rangle = \left\langle \sum_{K \geq L'} X_K T_K, P_{L'} \right\rangle$$

since $\langle T_K, P_{L'} \rangle = 0$ unless $K \geq L'$. We also have $|\langle T_K, P_{L'} \rangle| \leq 4^{|K|} = e^{O(n(K))}$ by the Cauchy–Schwarz inequality and because $|K| = \Theta(n(K))$.

Furthermore, by the Gram–Schmidt process, we see that $A_{L,L'}$ is growing at most exponentially in $n(L')$. That is, there exists $\xi > 0$ such that for any L', we have $|A_{L,L'}| \leq \xi^{n(L)}$.

Using the above bounds, the sum (9) is bounded by $e^{-c'n(L)} Z_0$ for some arbitrarily small $c' > 0$ and summing over L' and weighting by $A_{L,L'}$ gives a convergent sum.

We now use a classical convergence argument [6]. We showed that the probabilities of any cylindrical event converge. The cylindrical events form a family of sets which is stable under finite intersection. Since it generates the σ-field $\mathcal{B}$, it is a determining class, that is, if two probability measures on $(\Omega, \mathcal{B})$ coincide on all cylindrical events, then they are equal.

Since we furthermore have tightness by Theorem 13, the sequence of probability measures $\mathbb{P}^n$ admits subsequential limits by Prokhorov's theorem. More precisely, for any subsequence, there exists a sub-subsequence which converges. Since its value on the cylindrical events is known, there is only one possible limit. Let us call it $\mathbb{P}_{\text{nonc}}$. Now, since Ω is a metric space, this implies that the sequence $\mathbb{P}_{\text{nonc}}^n$ converges weakly to $\mathbb{P}_{\text{nonc}}$. $\square$

As a corollary, we obtain the convergence of the measures μ_Φ^n for any flat connection Φ.

COROLLARY 19.   *For any flat unitary connection* Φ, *there exists a probability measure* $\mathbb{P}_\Phi$ *on* $(\Omega, \mathcal{B})$ *such that*

$$\mathbb{P}_\Phi^n \to \mathbb{P}_\Phi$$

in the sense of weak convergence.

PROOF.   It suffices to show the convergence of this sequence of measures on finite intersections of balls in Ω of small radius because this is a determining class for $\mathcal{B}$. For any curve γ, there is a small radius $r > 0$ such that its tubular r-neighborhood retracts onto γ. Any curve in this r-neighborhood and winding once around is isotopic to γ. On any such neighborhood the density $\prod_{\gamma \subset L}(2 - \omega_\gamma - \omega_\gamma^{-1})$ is constant, hence the convergence follows by Theorem 18. $\qquad\square$

4.7. *Convergence in the curved case* (*proof of main statement*).   The measures $\mu_{\mathrm{LC}}, \mu_{\mathrm{LC}^0}$ converge in the following sense.

THEOREM 20.   *There exist probability measures* $\mathbb{P}_{\mathrm{LC}}, \mathbb{P}_{\mathrm{LC}^0}$ *on* $(\Omega, \mathcal{B})$ *and* $(\Omega_1^*, \mathcal{B})$ *respectively such that for any sequence* $(\mathcal{G}_n)_{n \geq 1}$ *of graphs, geodesically embedded on* Σ *and conformally approximating* Σ, *the sequences of probability measures* $\mathbb{P}_{\mathrm{LC}}^n$ *and* $\mathbb{P}_{\mathrm{LC}^0}^n$ *converge weakly toward respectively* $\mathbb{P}_{\mathrm{LC}}$ *and* $\mathbb{P}_{\mathrm{LC}^0}$.

PROOF.   Let us approximate Σ by a polygonal surface Σ_ε, that is, with a surface which is flat except for conical singularities. A standard way to do this is to take a fine triangulation of the surface (with triangles of diameter at most ε and whose angles are bounded from below), and replace each triangle with the Euclidean triangle with the same edge lengths. As $\varepsilon \to 0$ the conformal structure on Σ_ε converges to that of Σ. [Indeed, there is a homeomorphism ψ_ε from Σ_ε to Σ which is $(1 + o(1))$-bi-Lipschitz and, therefore, $(1 + o(1))$-quasiconformal.]

Any graph embedded on Σ or Σ_ε can be embedded on Σ_ε or Σ using ψ_ε^{-1} or ψ_ε; furthermore a graph conformally close to Σ_ε will have image on Σ conformally close to Σ, and vice versa.

Let $z_1, \ldots, z_k$ be the vertices of Σ_ε. The Levi-Civita connection on Σ_ε is a flat connection on $\Sigma_\varepsilon \setminus \{z_1, \ldots, z_k\}$ and approximates the Levi-Civita connection on Σ, in the sense that for small ε the curvature enclosed by any loop (chosen independently of the triangulation) is close for both connections. Restricting to $\mathcal{G}_n$, this shows that $\mathbb{P}_{\mathrm{LC},\varepsilon}^n$ is close to $\mathbb{P}_{\mathrm{LC}}^n$ [since cylinder events have measures which are within $o(1)$ of each other].

By Corollary 19, if we fix Σ_ε and take $\mathcal{G}_n$ embedded on Σ_ε and conformally approximating it as $n \to \infty$, the measures $\mathbb{P}_{\mathrm{LC},\varepsilon}^n$ converge as $n \to \infty$ to a limit $\mathbb{P}_{\mathrm{LC},\varepsilon}$. Similarly, for $\mathbb{P}_{\mathrm{LC}^0,\varepsilon}^n$.

For any $\varepsilon > 0$, the probabilities of the cylindrical events (away from the singularities) are determined by the function F of Lemma 17. These are expressed as integrals of Green's function on paths avoiding the singularities. The addition of a singularity modifies the Green's function, and hence the probability, by a negligible amount.

By a diagonal argument, we can take $n \to \infty$ and then $\varepsilon \to 0$ and conclude that $\mathbb{P}_{\mathrm{LC}}^n$ and $\mathbb{P}_{\mathrm{LC}^0}^n$ converge weakly as $n \to \infty$. $\quad\square$

From experimental simulations, the probability of getting two or more loops in μ_{LC} on the round sphere is on the order of one percent. Hence, measures μ_{LC} and μ_{LC^0} are close for the total variation distance.

The measure $\mathbb{P}_{\mathrm{LC}^0}$ is a limit of $\mathbb{P}_{\mathrm{LC}}$ when the metric is scaled by a factor $t \to 0$. Hence, it is not obvious that Theorem 11 implies that $\mathbb{P}_{\mathrm{LC}^0}$ also is supported on macroscopic loops. However, we conjecture it to be true.

CONJECTURE 21. $\mathbb{P}_{\mathrm{LC}^0}$ is supported on Ω_1 that is

$$\mathbb{P}_{\mathrm{LC}^0}(\text{the area of the loop is zero}) = 0.$$

4.8. *A comment and two applications.* The convergence result of probability measures $\mu_{c,\alpha}$ on CRSFs is actually more general and can be adapted for a wide range of functions α on the cycles, not necessarily coming from connections. This is due to the fact that the crucial convergence argument is made for the uniform measure on noncontractible CRSFs (Theorem 18). We now state two applications of the previous theorems.

As an application of Theorem 20, we obtain a result on higher moments of the area of a uniform CRST, also mentioned in the paper [13]. Let A denote the combinatorial area (number of faces) of the cycle of a cycle-rooted spanning tree of $\mathcal{G}_n$. Let $\theta = A/n^2$ be the Euclidean area of the cycle. We denote by $\mathbb{E}_{\mathrm{unif}}^n$ the expectation with respect to the uniform measure on CRSTs on $\mathcal{G}_n$.

COROLLARY 22. *For $k \geq 2$, there exists $a_k(D) > 0$ such that*

$$\mathbb{E}_{\mathrm{unif}}^n(A^k) = a_k(D)n^{2k-2}(1 + o(1)).$$

PROOF. Let $k \geq 2$. We have

$$\mathbb{E}_{\mathrm{unif}}^n(A^k) = \mathbb{E}_{\mathrm{LC}^0}^n(A^{k-2})\mathbb{E}_{\mathrm{unif}}^n(A^2)$$

$$= n^{2k-4}\mathbb{E}_{\mathrm{LC}^0}^n(\theta^{k-2})\mathbb{E}_{\mathrm{unif}}^n(A^2)$$

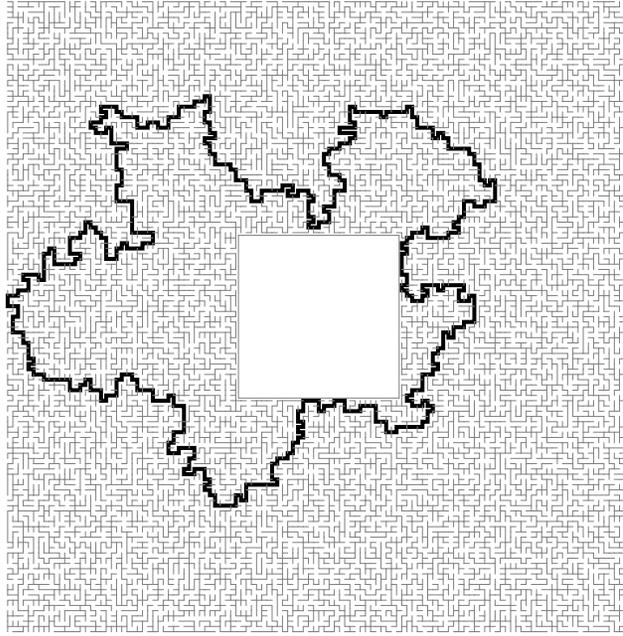$$= n^{2k-2}C(D)|D|\mathbb{E}_{\mathrm{LC}^0}^n(\theta^{k-2})(1 + o(1)),$$

FIG. 7.   *A uniform noncontractible CRST on an annulus in the square grid.*

where the last equality follows from Theorem 6 of [13] [here $C(D)$ is a constant proportional to the torsional rigidity of the domain] and the weak convergence of $\mathbb{P}^n_{\mathrm{LC}^0}$ to $\mathbb{P}_{\mathrm{LC}^0}$. Since θ is bounded and for $\mathbb{P}_{\mathrm{LC}^0}$ is with positive probability nonzero by Theorem 9, the limit $\mathbb{E}_{\mathrm{LC}^0}(\theta^{k-2})$ is a positive real. The corollary is proved by taking $a_k(D) = C(D)|D|\mathbb{E}_{\mathrm{LC}^0}(\theta^{k-2})$.   □

As an application of Theorem 18, let us state the following corollary.

COROLLARY 23.   *Consider a uniform spanning forest on an annulus-graph wired on its boundary. Then the simple closed curve separating the two connected components has a conformally invariant limit which is given by the measure* $\mathbb{P}_{\mathrm{nonc}}$ *conditional on having only one component.*

PROOF.   This follows from the fact that the dual of a wired essential forest on the annulus is a uniform noncontractible CRST on the annulus with free boundary conditions. The measure is thus given by μ^n_{nonc} conditional on having one loop. The convergence follows from Theorem 18.   □

Figure 7 shows a sample of this interface between the two tree components of a wired uniform spanning forest on the annulus.

5. Properties of the measures. In this section, we mention a few properties of the measures on CRSFs on surfaces we have been considering.

5.1. *Markov property.* CRSFs on surfaces with general cycle weights satisfy the following *spatial Markov property*. Consider a graph $\mathcal{G}$ embedded in a compact oriented surface Σ, possibly with boundary. Let $\alpha : \Omega_1 \to \mathbb{R}_{>0}$ be any positive weight function on the cycles of $\mathcal{G}$.

Let $\mathbf{\Gamma}$ be the random CRSF on $\mathcal{G}$. Let $\{\gamma\} = \{\gamma_1, \ldots, \gamma_k\}$ be a family of its cycles. These cycles separate Σ in a number of connected components $\Sigma_1, \ldots, \Sigma_r$. For $i = 1, \ldots, r$, denote by $\mathcal{G}_i$ the intersection of $\mathcal{G}$ and the closure of Σ_i (i.e., Σ_i along with its boundary) and by $\partial \mathcal{G}_i$ the boundary cycles.

PROPOSITION 24. *Conditional on $\{\gamma\} \subset \mathbf{\Gamma}$, the random CRSF $\mathbf{\Gamma}$ is equal in distribution to*

$$\{\gamma\} \bigsqcup_{i=1}^{r} \mathbf{\Gamma}_i.$$

PROOF. A CRSF Γ on $\mathcal{G}$ which contains $\{\gamma\}$ is the union $\{\gamma\} \bigsqcup_{i=1}^{r} \Gamma_i$ of $\{\gamma\}$ with essential CRSFs Γ_i on each one of the connected component Σ_i of $\Sigma \setminus \{\gamma\}$. The proposition then follows directly from the cycle popping algorithm. □

5.2. *Restriction property.* Let $D_1 \subset D$ be two planar Jordan domains. Let $\mathcal{G}$ be a finite graph approximation of D. Let Φ be a connection on a line bundle over $\mathcal{G}$. We denote $\det \Delta^D(\Phi)$ the line bundle Laplacian on $\mathcal{G}$ with connection Φ. For any subset S of the set of vertices of $\mathcal{G}$, we denote by $\det \Delta_S^D(\Phi)$ the line bundle Laplacian with Dirichlet boundary conditions on S; see [15].

The measure on multicurves that stay in D_1 is absolutely continuous with respect to the measure on D. The Radon–Nikodym derivative is given by a cross-ratio of determinants of the Laplacian with different boundary conditions.

LEMMA 25. *For any finite set of simple nonintersecting curves $\{\gamma\} \subset D_1$, we have*

$$\frac{\mu_\Phi^{D_1}(\{\gamma\})}{\mu_\Phi^{D}(\{\gamma\})} = \left(\frac{\det \Delta^{D_1}(\Phi)}{\det \Delta_{\{\gamma\}}^{D_1}(\Phi)} \right) \Big/ \left(\frac{\det \Delta^{D}(\Phi)}{\det \Delta_{\{\gamma\}}^{D}(\Phi)} \right).$$

PROOF. The proof follows by direct computation using the Forman–Kenyon matrix tree theorem. □

5.3. *Stochastic domination.* Recall from Section 2 the notation $\mu_{c,\alpha}$ for measures that assign a CRSF Γ a probability proportional to $\prod_{e \in \Gamma} c(e) \prod_{\gamma \subset \Gamma} \alpha(\gamma)$, where each $\alpha(\gamma) \in [0, 1]$.

LEMMA 26.   *Let $S_1 \subset S_2$ be two subgraphs of $\mathcal{G}$. Let $\mathbb{P}_1$ and $\mathbb{P}_2$ be essential CRSF measures $\mu_{c,\alpha}$ with Dirichlet boundary conditions on S_1 and S_2, respectively. For any curve $\gamma \in \mathcal{G} \setminus S_2$, we have*

$$\mathbb{P}_1(\gamma) \geq \mathbb{P}_2(\gamma).$$

PROOF.   This follows from the cycle-popping algorithm, as follows. We couple the cycle-popping algorithms for $\mathcal{G}_1$ and $\mathcal{G}_2$, starting with identical stacks of cards at each vertex. When a cycle is found for $\mathcal{G}_1$ which is also a cycle for $\mathcal{G}_2$, keep them both or discard them both according to the coin toss, increasing S_1, S_2 appropriately. If a cycle is found for $\mathcal{G}_1$ which is not a cycle of $\mathcal{G}_2$ (i.e., defines a LERW connected to the current S_2), keep that LERW in $\mathcal{G}_2$ (and, in addition, every part of the cycle which connects to the current boundary S_2) and toss a coin to determine whether or not to keep the cycle in $\mathcal{G}_1$. For either outcome of the coin toss it is still true that S_2, the current union of boundaries of $\mathcal{G}_2$, contains S_1, the current union of boundaries of $\mathcal{G}_1$. Continue until cycle γ is found for $\mathcal{G}_1$; at that point it will be kept in $\mathcal{G}_2$ if and only if it does not intersect the current S_2.   □

Note that in the case the measures come from a line bundle connection Φ, this implies

$$\frac{\det \Delta_{S_1 \cup \gamma}(\Phi)}{\det \Delta_{S_1}(\Phi)} \geq \frac{\det \Delta_{S_2 \cup \gamma}(\Phi)}{\det \Delta_{S_2}(\Phi)},$$

which is a nontrivial potential theoretic consideration (which can be translated in terms of Dirichlet-to-Neumann map).

6. Questions.

1. Can the measures $\mathbb{P}_{\text{nonc}}$, $\mathbb{P}_{\text{LC}^0}$, and $\mathbb{P}_{\text{LC}}$ be defined directly instead of via limits of CRSF measures? For example, via a stochastic differential equation, like a variant of SLE_2 defined on Riemann surfaces subject to some potential depending on the metric?

2. On the round sphere for the measure μ_{LC^0}, can one use the connection Laplacian to say more about the shape of the cycle, as is done in [13] in the flat case?

3. What is the function $x \mapsto \mathbb{P}_{\text{LC}^0}$ (x is enclosed by the loop) for the unit disk?

4. What can be said about the Gaussian Free Field associated to the line bundle Laplacian? Are our loop models related to this GFF? In particular, for a choice of an infinite curvature, we expect to obtain loops at all scales.

5. What is the right framework for the study of $\mathbb{P}_{\text{LC}}$? What is the distribution of the number of loops?

6. Is there a probabilistic interpretation for the coefficients of the triangular matrix **A** defined in Section 4.5.3?

7. What is the scaling limit of CRSFs in higher dimension?

8. What can be said about the scaling limit of measures μ_α on CRSFs for weight functions α not coming from a connection?

9. Can our loop measures be used to study the scaling limit of waves of avalanches in the sandpile model in the scaling limit? See [12] for a definition of waves in this model.

REFERENCES

[1] AIZENMAN, M. and BURCHARD, A. (1999). Hölder regularity and dimension bounds for random curves. *Duke Math. J.* **99** 419–453. MR1712629

[2] AIZENMAN, M., BURCHARD, A., NEWMAN, C. M. and WILSON, D. B. (1999). Scaling limits for minimal and random spanning trees in two dimensions. *Random Structures Algorithms* **15** 319–367. MR1716768

[3] ALDOUS, D. J. (1983). On the time taken by random walks on finite groups to visit every state. *Z. Wahrsch. Verw. Gebiete* **62** 361–374. MR0688644

[4] BENOIST, S. and DUBÉDAT, J. (2016). An SLE_2 loop measure. *Ann. Inst. Henri Poincaré Probab. Stat.* To appear. Available at arXiv:1405.7880.

[5] BIGGS, N. (1997). Algebraic potential theory on graphs. *Bull. Lond. Math. Soc.* **29** 641–682. MR1468054

[6] BILLINGSLEY, P. (1999). *Convergence of Probability Measures*, 2nd ed. Wiley, New York. MR1700749

[7] FOCK, V. and GONCHAROV, A. (2006). Moduli spaces of local systems and higher Teichmüller theory. *Publ. Math. Inst. Hautes Études Sci.* **103** 1–211. MR2233852

[8] FORMAN, R. (1993). Determinants of Laplacians on graphs. *Topology* **32** 35–46. MR1204404

[9] GINÉ, E. and KOLTCHINSKII, V. (2006). Empirical graph Laplacian approximation of Laplace–Beltrami operators: Large sample results. In *High Dimensional Probability*. *Institute of Mathematical Statistics Lecture Notes—Monograph Series* **51** 238–259. IMS, Beachwood, OH. MR2387773

[10] HOUGH, J. B., KRISHNAPUR, M., PERES, Y. and VIRÁG, B. (2006). Determinantal processes and independence. *Probab. Surv.* **3** 206–229. MR2216966

[11] ITÔ, K. (1963). The Brownian motion and tensor fields on Riemannian manifold. In *Proc. Internat. Congr. Mathematicians* (*Stockholm*, 1962) 536–539. Inst. Mittag-Leffler, Djursholm. MR0176500

[12] IVASHKEVICH, E. V., KTITAREV, D. V. and PRIEZZHEV, V. B. (1994). Waves of topplings in an Abelian sandpile. *Phys. A* **209** 347–360.

[13] KASSEL, A., KENYON, R. and WU, W. (2015). Random two-component spanning forests. *Ann. Inst. Henri Poincaré Probab. Stat.* **51** 1457–1464. MR3414453

[14] KASSEL, A. and WU, W. (2015). Transfer current and pattern fields in spanning trees. *Probab. Theory Related Fields* **163** 89–121. MR3405614

[15] KENYON, R. (2011). Spanning forests and the vector bundle Laplacian. *Ann. Probab.* **39** 1983–2017. MR2884879

[16] KENYON, R. (2014). Conformal invariance of loops in the double-dimer model. *Comm. Math. Phys.* **326** 477–497. MR3165463

[17] KENYON, R. W., PROPP, J. G. and WILSON, D. B. (2000). Trees and matchings. *Electron. J. Combin.* **7** Research Paper 25, 34 pp. (electronic). MR1756162

[18] LAWLER, G. F., SCHRAMM, O. and WERNER, W. (2004). Conformal invariance of planar loop-erased random walks and uniform spanning trees. *Ann. Probab.* **32** 939–995. MR2044671

[19] LAWLER, G. F. and TRUJILLO FERRERAS, J. A. (2007). Random walk loop soup. *Trans. Amer. Math. Soc.* **359** 767–787 (electronic). MR2255196

[20] LYNDON, R. C. and SCHUPP, P. E. (2001). *Combinatorial Group Theory*. Springer, Berlin. MR1812024

[21] SCHRAMM, O. (2000). Scaling limits of loop-erased random walks and uniform spanning trees. *Israel J. Math.* **118** 221–288. MR1776084

[22] WEHN, D. (1962). Probabilities on Lie groups. *Proc. Natl. Acad. Sci. USA* **48** 791–795. MR0153042

[23] WILSON, D. B. (1996). Generating random spanning trees more quickly than the cover time. In *Proceedings of the Twenty-Eighth Annual ACM Symposium on the Theory of Computing* (*Philadelphia*, *PA*, 1996) 296–303. ACM, New York. MR1427525

DEPARTEMENT MATHEMATIK
ETH ZÜRICH
RÄMISTRASSE 101
8092 ZÜRICH
SWITZERLAND
E-MAIL: adrien.kassel@math.ethz.ch

MATHEMATICS DEPARTMENT
BROWN UNIVERSITY
151 THAYER ST.
PROVIDENCE, RHODE ISLAND 02912
USA
E-MAIL: rkenyon@math.brown.edu