

Dimers on Riemann surfaces I: Temperleyan forests

Nathanaël Berestycki, Benoit Laslier, and Gourab Ray

Abstract. This is the first article in a series of two papers in which we study the Temperleyan dimer model on an arbitrary bounded Riemann surface of finite topological type. The end goal of both papers is to prove the convergence of height fluctuations to a universal and conformally invariant scaling limit. In this part, we show that the dimer model on the Temperleyan superposition of a graph embedded on the surface and its dual is well posed, provided that we remove an appropriate number of punctures. We further show that the resulting dimer configuration is in bijection with an object which we call Temperleyan forest, whose law is characterised in terms of a certain topological condition. Finally, we discuss the relation between height differences and Temperleyan forest and prove that the convergence of the latter (which is the subject of the second paper in this series) implies the existence of a conformally invariant scaling limit of height fluctuations.

Contents

1. Introduction	363
2. Background and setup	371
3. Temperley's bijection on Riemann surfaces	385
4. Winding and height function	393
5. Local coupling	407
6. Convergence of height function and forms	416
A. Geometry of spines	438
References	442

1. Introduction

1.1. Background and main result

This paper is the first part of a series of two papers in which we show universality and conformal invariance of the Temperleyan dimer model on Riemann surfaces.

Mathematics Subject Classification 2020: 60B10 (primary); 60C05 (secondary).

Keywords: dimers, CRSF, loop soup, loop-erased random walk, Temperley.

Given a finite graph G (with an even number of vertices) and nonnegative weights $w = (w_e)_{e \in E}$ on the edges, a dimer configuration is a perfect matching of the vertices along the edges of G . The dimer model is the random perfect matching obtained by sampling a dimer configuration \mathbf{m} on G with probability proportional to the product of the weights of the edges in the matching \mathbf{m} :

$$\mathbb{P}(\mathbf{m}) = \frac{\prod_{e \in \mathbf{m}} w_e}{Z}, \tag{1.1}$$

where Z is the partition function. If all the weights are equal to 1, \mathbb{P} is simply a uniformly picked perfect matching. Following Thurston, a particularly convenient way to encode the dimer model on a *planar* bipartite graph is through a notion of height function. Informally, it is a real-valued function defined on the faces of the graph, which is uniquely determined by the dimer configuration. This turns the dimer model into a model of discrete random surfaces. A central question in the study of the dimer model is to describe the fluctuations around the mean of the height function. It turns out that the height function of the dimer model is extremely sensitive to boundary conditions (see, for example, the quote of Kasteleyn at the beginning of [14]). In the so-called *liquid phase*, the Kenyon–Okounkov conjecture predicts that the fluctuations of the height function around its mean are asymptotically given by a *Gaussian free field* in a certain nontrivial *conformal gauge* (see [27, Section 2.3] for details). Despite the remarkable progress on this conjecture over the last decade, this remains an outstanding problem in all but a handful of cases: see [4, 19, 26, 34, 35] for examples of models where this is already shown, and see [7, 37] for an introduction to the Gaussian free field. We emphasise that, given the extreme sensitivity of the model to its boundary conditions, the occurrence of a universal limit such as the Gaussian free field is to some extent very surprising.

The goal of this paper is to describe the fluctuations of the height function for certain graphs embedded on Riemann surfaces of finite topological type, i.e., with finitely many handles and holes. Our assumptions on these graphs ensure that the dimer model is liquid on the entire Riemann surface (as opposed to having frozen or gaseous regions, say as in the Aztec diamond) and the conformal gauge is simply the identity; we call the resulting dimer model *Temperleyan*, and we say that the sequence of graphs $G^{\#\delta}$ is a Temperleyan discretisation of the surface. Even though this assumption simplifies the problem, conformal invariance in the scaling limit is by no means obvious; for instance, in the simply connected setting and for the square lattice, this is in fact precisely the content of Kenyon’s landmark paper [25].

Before we give a simplified statement of the main theorem, we point out that, in the context of a Riemann surface M , the height function becomes a height one-form $h^{\#\delta}$: that is, the gradient $\nabla h^{\#\delta}$ is well defined, but this gradient does not derive itself from a single-valued function on the surface. Instead, it can be viewed as the

gradient of a (single-valued) function defined on the universal cover of the surface. For concreteness, we will denote by $h^{\#\delta}$ the associated function on the universal cover of the surface. Readers who prefer it can however equivalently think of $h^{\#\delta}$ as the random one-form on M associated to the dimer configuration. (See Sections 2.5 and 6.1 for precise definitions.) The function $h^{\#\delta}$ is the central object of this sequence of papers, and we still refer to it as the height function.

Our main theorem, stated informally below, confirms that the dimer model admits a universal scaling limit on Riemann surfaces. This scaling limit is furthermore conformally invariant.

Theorem 1.1. *Let M denote a bounded Riemann surface, possibly with a boundary, which is not the sphere or a simply connected domain. Let $G^{\#\delta}$ be a sequence of Temperleyan discretisations of M satisfying the invariance principle and a crossing estimate for random walk described in Section 2.4. Let $h^{\#\delta}$ denote the height function of the dimer model on $G^{\#\delta}$. For all smooth compactly supported test functions f on the universal cover of M ,*

$$\int (h^{\#\delta} - \mathbb{E}(h^{\#\delta})) f$$

converges as $\delta \rightarrow 0$ in law and the limit is conformally invariant and independent of the sequence $G^{\#\delta}$.

The precise assumptions on the surface M will be given in Section 2.1; the meaning of Temperleyan discretisations and our precise assumptions on the sequence $G^{\#\delta}$ can be found in Section 2.4. Finally, note that the smooth test function f is taken on the universal cover \tilde{M} of M so that the integration takes place on \tilde{M} (equipped, e.g., with Lebesgue measure). See Sections 2.5 and 6.1 for details. We refer to Theorem 6.1 for the complete version of the theorem.

Although we cannot find what the limiting object is, as an artefact of our proof, we can show that the limit is nontrivial. Indeed, we prove later in Proposition 5.10 that the limiting field restricted to a small enough neighbourhood of M away from the punctures can be coupled to the restriction of a Gaussian free field in a slightly bigger neighbourhood, with positive probability; in particular, this field is nontrivial. We conjecture that the limiting field is a *compactified Gaussian free field* with appropriate parameters depending only on the surface M and the positions of the punctures (x_1, \dots, x_k) on M .

1.2. Relation with previous work

One of the first rigorous, works on the dimer model on Riemann surfaces is the inspiring paper by Boutillier and de Tilière [9], who were motivated by nonrigorous, physics predictions based on the Coulomb gas formalism (see, e.g., [17] and

the references in [9]). They described the topological part of the fluctuations in the dimer model on the honeycomb lattice on the torus. This was complemented some years later by the fundamental paper of Dubédat [19], in which the full law of the limiting height fluctuations was described for isoradial graphs on the flat torus, and identified as a so-called compactified Gaussian free field. Dubédat also conjectured a greater form of universality beyond the isoradial setting; such a universality statement was subsequently proved by Dubédat and Gheissari [20] but only for the topological part of fluctuations on the torus. Note that in particular our Theorem 1.1 settles this conjecture of Dubédat.

We point out that all techniques in the paper mentioned above, based on Kasteleyn theory (in contrast with our approach), are very specific to the case of the torus and cannot easily be adapted to general Riemann surfaces. Cimasoni [10] shows that, for general surfaces of genus g (and no boundary), the partition function of the dimer model becomes a weighted (in fact, signed) sum of 2^{2g} determinants of Kasteleyn matrices in which the orientation has been reversed along some of the $2g$ cycles forming a basis of the first homology group of the surface. The signs in front of the determinants themselves form a nontrivial algebraic topological invariant of the surface which Cimasoni, remarkably, was able to identify in terms of the so-called Arf invariant. (In the case of the torus and the honeycomb lattice, this goes back to the work of Boutillier and de Tilière [9].) This illustrates the difficulty in generalising the Kasteleyn approach to general surfaces. See also [1, 11, 12, 15] for related results.

The above difficulty illustrates the well-known fact that the dimer models become “less integrable” as the genus of the surface increases, or at any rate integrability becomes harder to exploit. In this paper, we will therefore follow the general strategy of our previous work [4], which relies instead on Temperley’s bijection between the dimer model and uniform spanning trees (in fact, see the overview of the proof in Section 1.3 for a discussion of this strategy and the relation between the two papers). The increased robustness of this method accounts in large part for the fact that we are able to tackle generic Riemann surfaces. (In contrast with methods based on Kasteleyn theory or interacting particle systems.) This is not to say that the difficulties caused by the increased complexity of the surface vanish entirely, but they manifest themselves in a different manner, which boils down to the following. As we will describe in more details in Section 2, a Temperleyan discretisation of the surface is the graph resulting from the superposition of a graph Γ embedded on the surface and its dual. The resulting graph is bipartite, with its black vertices being given by the vertices and faces of Γ and the white vertices being given by the edges of Γ . However, a straightforward calculation using Euler’s formula shows that this does not admit a dimer covering unless we remove $2g + b - 2 = |\chi|$ many white vertices from the graph (where g is the genus and b the number of boundary components, and χ is the Euler characteristic); these can be thought of as *punctures* in the surface, or monomers. Note

that the number of punctures increases with the complexity of the surface, and when the surface is a torus or an annulus, no such puncture needs to be removed. Roughly speaking, we will see that the locally tree-like random subgraphs appearing when we generalise Temperley's bijection, which we will call *Temperleyan self-dual pair*, have one topological constraint for each puncture. In the case of a torus or an annulus where there is no puncture, we will see that the Temperleyan self-dual pair boils down to the simpler notion of cycle-rooted spanning forest studied by Kassel and Kenyon in [24].

The necessity to remove a certain number of punctures from the surface in order to formulate the Temperleyan dimer model has a geometric flavour which is reminiscent of *Liouville conformal field theory*, in which correlation functions are only well defined with the appropriate number of singularities (known in that context as insertions). The required number of insertions is specified by the so-called Seiberg bounds; see, e.g., [16] on the sphere and [21] for Riemann surfaces (see also [7] for an introduction). As is the case here, the number of punctures is entirely specified by the topology of the surface, but in a way that is in some sense opposite to ours: indeed the required number of insertions is an increasing function of χ in Liouville CFT, whereas it is decreasing here.

This difference can be understood informally as follows. To first order, Liouville CFT may be viewed as yielding a random metric which is approximately hyperbolic. (This can, for instance, be made rigorous in the semiclassical limit where the parameter $\gamma \rightarrow 0$; see [29].) However, on the Riemann sphere, for which the theory is of particular interest, and on other surfaces of non-negative Euler characteristic, it is impossible to equip the surface with a metric of constant negative scalar curvature without first adding conical singularities at a certain number of points, as follows from the Gauss–Bonnet theorem. By contrast, the Temperleyan dimer model is trying to equip the surface with a *flat metric*, a fact which will be particularly apparent when we relate the height function to the winding of the Temperleyan forest in Section 4. (Note that the winding is computed on the universal cover, where we ignore the effects of curvature; see Section 2.2.) Yet, on a hyperbolic surface (i.e., when $\chi < 0$), no flat metric exists, again by the Gauss–Bonnet theorem. In summary, Liouville CFT may be viewed as trying to build a hyperbolic metric on a surface which may be parabolic or elliptic, whereas the Temperleyan dimer model tries to impose a flat metric on a surface which may be hyperbolic. In both cases, the difficulty is resolved by puncturing the surface at an appropriate number of singularities. (The above parallel between Liouville CFT and dimer model may perhaps also be compared with the two couplings between SLE and Gaussian free fields given, respectively, by the so-called forward and reverse couplings; the forward coupling describes Liouville conformal field theory, whereas the reverse coupling describes imaginary geometry [32, 33].)

1.3. Overview of the proof and organisation of the paper

As mentioned above, our general approach to study height fluctuations on Riemann surfaces is motivated and inspired by our earlier work [4] connecting the dimer model to *imaginary geometry* in the simply connected setting. Let us recall the main idea of that work briefly in the original simply connected setting before explaining how this relates with and differs from the case of Riemann surfaces.

- The starting point of [4] is Temperley’s bijection, which relates the dimer model on a Temperleyan discretisation of a simply connected domain to a pair of dual uniform spanning trees with, respectively, wired and free boundary conditions. Furthermore, the height function of the dimer model can be determined relatively simply from the associated pair of spanning trees: roughly speaking, the height difference between two points is simply the total winding (sum of turning angles) of the unique path connecting these points in either tree.
- The second observation is that, as follows from the landmark paper of Lawler, Schramm, and Werner [30], paths between points in the wired uniform spanning tree (UST) have a scaling limit which may be described in terms of Schramm–Loewner Evolutions with parameter $\kappa = 2$ or SLE_2 for short.
- Finally, in the continuum, there is a coupling known as imaginary geometry between SLE_κ curves and an appropriate multiple of the Gaussian free field. In this coupling, the values of the GFF may informally be identified as the winding of the associated SLE_κ curves. (Note that this requires careful interpretation, since the Gaussian free field is not defined pointwise and SLE_κ paths are not smooth.) Nevertheless, this coupling can be thought of as a continuous form of Temperley’s bijection in the case $\kappa = 2$.

The strategy of [4] is to exploit Temperley’s bijection on the one hand and imaginary geometry on the other hand to show that the winding of paths in a UST is indeed given by the appropriate multiple of the Gaussian free field asymptotically. (Essentially, one has to show that we can exchange the order of taking limits, or that a diagram commutes.) This strategy is very robust; in particular, it does not appeal to the more solvable aspects of the dimer model. Indeed, the convergence of loop-erased random walks (which describe the branches of a wired UST at the discrete level) to SLE_2 is known in great generality.

At a broad level, the approach we wish to pursue in the case of Riemann surfaces follows similar steps. However, it should be clear that each step will be substantially different in this new setup.

- (1) To begin with, we observe that Temperley’s bijection is local, and so, it may be applied on a Riemann surface to output a pair of random subgraphs that are dual to one another and locally tree-like, which we call *Temperleyan self-dual*

pair. While in the simply connected case, these are just dual spanning trees, the combinatorial structure (i.e., state space) of the Temperleyan self-dual pair on a Riemann surface is a priori not entirely clear, and its law even less so. Our first result will be a characterisation of the Temperleyan self-dual pair. We also describe its law explicitly, which has a simple and concise expression. It turns out that one of the two components, which we call the *Temperleyan cycle-rooted spanning forest* or Temperleyan CRSF for short, determines the other (and hence the entire dimer configuration) almost uniquely (up to some choice of orienting cycles), and it also has a tractable law. Essentially, a Temperleyan CRSF is a random spanning subgraph in which every cycle is non-contractible and satisfies a certain further topological condition. The introduction and study of Temperleyan CRSFs take up a substantial part in both of the articles in this series.

- (2) A second problem is that the connection between height function and winding is not as straightforward as in the simply connected case; in particular, in addition to the already mentioned fact that only the gradient of the height function is well defined (leading to a multivalued height function), there may not be a path connecting two given points in the Temperleyan CRSF, even when they are neighbours. We thus need to develop a systematic method for computing the height function given the Temperleyan CRSF.
- (3) The next difficulty is to obtain a scaling limit result for our Temperleyan CRSF. Let us immediately note here that there are specific difficulties with even defining variants of SLE on surfaces since the simply connected uniformisation theorem lies at the heart of the standard definition of SLE.
- (4) Finally, if we were to follow the blueprint of [4], one would also need to establish a version of imaginary geometry for Riemann surfaces, an endeavour which seems out of reach with current techniques. Fortunately, we are in fact able to entirely bypass this step because it turns out that the (discrete) “error” estimates needed anyway to exchange limits are strong enough to imply convergence regardless of any information in the continuum. Clearly, the downside of this approach that we cannot hope to describe the law of the limit in this manner.

The topological and geometric structures of the surface M raise substantial specific difficulties at each of the steps of this roadmap. In this first paper, we develop the relevant discrete/combinatorial arguments corresponding to the first two items above and present the conclusion according to the last item. The second paper in this sequence, [5], is devoted to the third bullet point, i.e., the proof of convergence of the Temperleyan CRSF in the scaling limit. This is the most involved step which requires several new ideas but can be read essentially independently.

At this point, it is also useful to discuss a bit more precisely the role that the different structures associated to a Riemann surface will play in our paper. In order to even discuss convergence on the surface M , one needs in principle to specify a Riemannian metric on M , and it is natural to require that this metric be consistent with the conformal structure of M . However, ultimately, we want to prove convergence to a conformally invariant object, so, in principle, only the conformal class of this metric matters. (Here, we mean that two metric tensors are equivalent if we obtain one from the other by multiplying by a smooth function e^ϕ .)

It turns out however that, after lifting to the universal cover (which in most of the interesting cases for this paper will be the unit disc; see below for more details), it will be more convenient for us to equip this cover with the standard Euclidean metric instead of the *a priori* more natural hyperbolic metric.

This impacts the way we formulate some of our assumptions in Section 2.4. Unfortunately, some of our technical estimates involve indirectly the Euclidean metric on the unit disc: for instance, through our assumptions that the graph G has a bounded density and each edge has a bounded winding (items (i) and (ii) in Section 2.4). Nevertheless, changing the metric to another one in the same class only changes the constants involved in these assumptions, and so, it does not impact the main results. We explicitly prove this fact in Lemma 2.5.

Organisation of the paper

- In Section 2, we recall some basic facts about dimers on surfaces, and in Section 2.3, we recall the notions of windings of curves which were developed in our previous article [4]. Sections 2.1 and 2.2 contain small reminders about Riemann surfaces. We introduce precisely our discrete setup in Section 2.4. As one of the main differences with the classical case is that the height function now becomes multivalued; we recall in Section 2.5 the language of height one-forms on graphs which is a classical way to handle such multivalued functions. In Section 2.5, we introduce another way to deal with multivalued functions by lifting them to single-valued functions on the universal cover. We also illustrate how an application of the Hodge decomposition theorem allows us to decompose the (multivalued) height function on the surface into a single-valued function (scalar component) and a canonical representative of the “multivalued (or topological) part”, the so-called instanton component.
- In Section 3, we introduce the generalisation of Temperley’s bijection to Riemann surfaces, essentially covering point (1) above. Section 3.1 contains the bijection itself. We also provide in Section 3.2 a simple criterion for a CRSF to be Temperleyan which is the basis of the second paper in the series [5].

- In Section 4, we carefully develop the relation between the height differences and the winding of Temperleyan forests, i.e., point 2 of the overview. This is in the spirit of the work of Kenyon–Propp–Wilson [28], who treated the (planar) simply connected case. In that case, of course, the bijection is with a uniform spanning tree. Compared to that work, there are two additional points that we need to handle. The first one is that the edges of the graph cannot be assumed to be straight lines as in [28]. The second and more significant one is that there are additional terms coming from the fact that the forest is not connected: given points x and y on the universal cover, the height difference $h(x) - h(y)$ (which is unambiguously defined) is essentially given by the intrinsic winding of any path connecting x to y , *plus* additional discontinuities every time the path jumps between components. The resulting key formula is stated in Theorem 4.10 (see also Lemma 4.6).
- In Section 5, we extend our local and global coupling results from [4] to the framework of Riemann surfaces. This coupling is a key ingredient of our approach, which allows us to show that given a finite number of points $(z_i)_{1 \leq i \leq k}$ on the surface (or, rather, their lifts to the universal cover), the respective geometries of the CRSF in neighbourhoods of these k points are almost independent. This is crucial because certain types of “error terms” in our approach vanish in expectation but are large in absolute value.
- In Section 6, we conclude by proving our main convergence result: Theorem 6.1. Sections 6.2 and 6.3 contain some technical a priori estimates on winding of the CRSF branches on the universal cover. Finally, in Section 6.4, we finish the proof of Theorem 6.1.
- Finally, appendix A contains some geometric facts about *spines* in the CRSF (that is, the lifts to the universal cover of a cycle of the CRSF). It is given as an appendix because the proofs rely on slightly more involved elements of the classical theory of Riemann surfaces than the rest of the paper and we wanted to avoid overloading the already lengthy setup section.

2. Background and setup

2.1. Riemann surfaces and embedding

In this article, we work with a Riemann surface M (a connected, one-dimensional complex manifold) satisfying the following properties.

- M is of finite topological type, meaning that the fundamental group $\pi_1(M)$ is finitely generated. In other words, we assume that the surface has finitely many “handles” and “holes”.

- M can be compactified by specifying a *boundary* ∂M . We denote by \bar{M} the compactified Riemann surface with the boundary. More precisely, every point is either in the interior and hence has a local chart homeomorphic to \mathbb{C} or is on the boundary and has a local chart homeomorphic to the closed upper half-plane $\bar{\mathbb{H}}$. Also, there are finitely many such charts which cover the boundary. Note that this condition implies that M has no punctures. (However, for future reference, we note here that we will later introduce punctures on M which will align with the removed vertices of Temperleyan graphs. The resulting punctured manifold will be denoted by M' .)

We say M is *nice* if M satisfies the above properties.

Given such a Riemann surface M , it is possible to equip it with a Riemannian metric g (with associated distance function denoted by d_M) which extends continuously to the boundary and turns M into a smooth Riemannian surface. This Riemannian metric can be constructed by considering the hyperbolic metric on the universal cover of the surface if it has no boundary; see [23, Section 2.4] in combination with the uniformisation theorem ([23, Theorem 4.4.1], also recalled and discussed below). If the surface has boundary components, we may apply the same result to the surface obtained by gluing it to itself along the boundary – also called the double.

Note that the choice of the metric above excludes surfaces, such as the hyperbolic plane. This will simplify certain topological issues later when we deal with the Schramm topology.

Classification of surfaces. Riemann surfaces can be classified into the following classes depending on their conformal type (see, e.g., [18] for an account of the classical theory):

- Elliptic: this class consists only of the Riemann sphere, i.e., $M \equiv \hat{\mathbb{C}}$.
- Parabolic: this class includes the torus, i.e.,

$$M \equiv \mathbb{T} = \mathbb{C}/(\mathbb{Z} + \tau\mathbb{Z}),$$

where $\Im(\tau > 0)$, the cylinder $M \equiv \mathbb{C} \setminus \{0\}$, and the complex plane itself (or the Riemann sphere minus a point), $M \equiv \mathbb{C}$.

- Hyperbolic: this class contains everything else. This includes examples such as the two-torus, the annulus, as well as proper simply connected domains in the complex plane, etc.

The proofs in this paper (and its companion [5]) are concerned with the hyperbolic case (subject to the above conditions) as well as the case of the torus. So, from now on, we always assume that M is such a manifold. We note that the case of simply connected proper domain in \mathbb{C} is covered in our previous work [4].

2.2. Universal cover

The universal cover \tilde{M} of the Riemann surface M will play an important role in our analysis. Recall that any Riemann surface admits a (regular) covering map by a covering space which is simply connected. This covering space may furthermore be endowed with a conformal structure, with respect to which the covering map is then analytic. Being simply connected, this covering space is (by the uniformisation theorem, see [23, Theorems 4.4.1 and 2.4.3]) conformally equivalent to a Riemann surface \tilde{M} , which is either the unit disc $\mathbb{D} \subset \mathbb{C}$ (hyperbolic case), the whole plane (parabolic case), or the Riemann sphere $\hat{\mathbb{C}}$ (elliptic case). Altogether, we obtain a map¹ $p : \tilde{M} \rightarrow M$ which is both analytic and a regular covering map of M . By definition of a covering map, p is a local homeomorphism: for every $z \in M$, there exists a neighbourhood N containing z so that p is injective in every component of $p^{-1}(N)$.

Recall further that in each of the three cases (elliptic, parabolic, and hyperbolic), \tilde{M} can be endowed with a metric (compatible with the complex structure) of constant Gaussian curvature equal to 1, 0, and -1 , respectively. Furthermore, there is a subgroup F of conformal automorphisms on \tilde{M} acting properly discontinuously and freely – i.e., without fixed points – (see, e.g., [23, Section 2.4] for definitions) such that p descends to a conformal equivalence between \tilde{M}/F and M . Put more simply, F is a discrete subgroup of the corresponding set of Möbius transformations describing the conformal automorphisms of \tilde{M} , and M is conformally equivalent to \tilde{M}/F (i.e., to either $\hat{\mathbb{C}}/F$, \mathbb{C}/F , or \mathbb{D}/F).

In case of the torus, this discrete subgroup is isomorphic to \mathbb{Z}^2 and the generators specify translations in the two directions of the torus. In the hyperbolic case, this class of subgroups is much bigger and are known as *Fuchsian groups* (see, e.g., [18] for a general account). This particular representation of a hyperbolic Riemann surface, sometimes called a Fuchsian model, will be particularly convenient because it allows us to describe the scaling limits of Temperleyan or cycle-rooted spanning forests *in the universal cover*, rather than on the surface itself. This will allow us to import directly a number of the ideas and results from [4] on the simply connected case.

Concretely, one of the main advantages of the Fuchsian model of a hyperbolic Riemann surface – and more generally of a lift to the universal cover which inherits the Euclidean metric – is that, *equipped with this Euclidean metric*, we have a flat embedding in which the sum of angles around a point is 2π , the total winding of a simple loop is 2π (see also Lemma 2.1 below). The Fuchsian structure itself plays a more technical role; see, for instance, appendix A. We rely on this for a canonical description of the height function in terms of the extended Temperley bijection.

¹For concreteness, one can fix a point x_0 in M and consider p such that $p(0) = x_0$. In the hyperbolic case, fix an arbitrary choice of rotation as well.

To conclude, moving to the universal cover allows us in a sense to trade a natural law on a graph possessing a nontrivial topology for a law with complicated periodicity-like constraints but on a planar graph.

2.3. Intrinsic and topological winding

The goal of this section is to recall several notions of windings of curves drawn in the plane, which we use in this paper. We refer to [4] for a more detailed exposition. A self-avoiding (or simple) curve in \mathbb{C} is an injective continuous map $\gamma : [0, T] \rightarrow \mathbb{C}$ for some $T \in [0, \infty]$.

The *topological winding* of such a curve γ around a point $p \notin \gamma[0, T]$ is defined as follows. We first write

$$\gamma(t) - p = r(t)e^{i\theta(t)},$$

where the function $\theta : [0, \infty) \rightarrow \mathbb{R}$ is taken to be continuous, which means that it is unique modulo a global additive constant multiple of 2π . We define the winding for an interval of time $[s, t]$, denoted by $W(\gamma[s, t], p)$, to be

$$W(\gamma[s, t], p) = \theta(t) - \theta(s).$$

(Note that this is uniquely defined.) Notice that if the curve has a derivative at an endpoint of γ , we can take p to be this endpoint by defining

$$W(\gamma[0, t], \gamma(0)) := \theta(t) - \lim_{s \rightarrow 0} \theta(s),$$

and similarly,

$$W(\gamma[s, T], \gamma(T)) := \lim_{t \rightarrow T} \theta(t) - \theta(s).$$

With this definition, winding is additive: for any $0 \leq s \leq t \leq T$,

$$W(\gamma[0, t], p) = W(\gamma[0, s], p) + W(\gamma[s, t], p).$$

The notion of *intrinsic winding* we describe now, also discussed in [4], is perhaps a more natural definition of windings of curves. This notion is the continuous analogue of the discrete winding of non-backtracking paths in \mathbb{Z}^2 which can be defined just by the number of anticlockwise turns minus the number of clockwise turns. Notice that we do not need to specify a reference point with respect to which we calculate the winding; hence, our name “intrinsic” for this notion.

We call a curve smooth if the map γ is smooth (continuously differentiable). Suppose that γ is smooth and, for all t , $\gamma'(t) \neq 0$. We write $\gamma'(t) = r_{\text{int}}(t)e^{i\theta_{\text{int}}(t)}$, where again $\theta_{\text{int}} : [0, \infty) \rightarrow \mathbb{R}$ is taken to be continuous. Then, define the intrinsic winding in the interval $[s, t]$ to be

$$W_{\text{int}}(\gamma[s, t]) := \theta_{\text{int}}(t) - \theta_{\text{int}}(s).$$




Figure 1. The graphs Γ , Γ^\dagger , \hat{G} , and G . The graph Γ is pictured in red and Γ^\dagger in black.

The total intrinsic winding is again defined to be $\lim_{t \rightarrow T} W_{\text{int}}(\gamma, [0, t])$ provided this limit exists. Note that this definition does not depend on the parametrisation of γ (except for the assumption of non-zero derivative). The following topological lemma from [4] connects the intrinsic and topological windings for smooth curves.

Lemma 2.1 ([4, Lemma 2.1]). *Let $\gamma[0, 1]$ be a smooth self-avoiding curve in \mathbb{C} ; then*

$$W_{\text{int}}(\gamma[0, 1]) = W(\gamma[0, 1], \gamma(1)) + W(\gamma[0, 1], \gamma(0)).$$

Note that the right-hand side in the above expression makes sense as soon as the endpoints of γ are smooth. From now on, we therefore extend the definition of intrinsic winding to all such curves using the equality of Lemma 2.1. Since the topological winding is clearly a continuous function on the curve away from the point p , this extension is consistent with any regularisation procedure. We also recall the following deformation lemma from [4, Remark 2.5].

Lemma 2.2. *Let D be a domain and $\gamma \subset \bar{D}$ a simple smooth curve (or piecewise smooth with smooth endpoints). Let ψ be a conformal map on D , and let $\arg_{\psi'(D)}$ be any realisation of argument on $\psi'(D)$. Then,*

$$W_{\text{int}}(\psi(\gamma)) - W_{\text{int}}(\gamma) = \arg_{\psi'(D)}(\psi'(\gamma(1))) - \arg_{\psi'(D)}(\psi'(\gamma(0))).$$

2.4. Discretisation setup

We will require the graphs we are working with to be embedded on M in a way that preserves its topology and geometry; we make this precise now.

Let M be a Riemann surface with b holes and g handles satisfying the assumptions of Section 2.1. We say that a graph Γ is *faithfully* embedded in M if it satisfies the following (see Figure 1).

- The embedding is proper; i.e., edges do not cross.

- Γ is connected.
- Every face of Γ not intersecting ∂M has the topology of a disc.
- Γ has a marked vertex (called a boundary vertex, the set of which is denoted by $\partial\Gamma$) for each component of ∂M .

Note that, with this convention, we can think of each boundary vertex as being “delocalised” along the whole boundary component, and indeed later, we will consider cycle-rooted spanning forests (CRSFs) with a wired condition on each boundary component. This “delocalised” boundary vertex is best illustrated to be sitting inside the hole; see Figure 1.

Let Γ^\dagger denote the dual graph of Γ , defined as follows. Every face of Γ contains a single vertex of Γ^\dagger and two vertices of Γ^\dagger are connected by an edge if their corresponding faces share an edge of Γ . We also assume the following for Γ^\dagger .

- The embedding is proper; i.e., the edges do not cross.
- Γ^\dagger is connected.
- Every face of Γ^\dagger corresponding to a non-boundary vertex of Γ is a topological disc.
- The vertices of Γ^\dagger corresponding to the faces of Γ which contain a given boundary vertex form a simple cycle homotopic to the corresponding boundary component of ∂M . We call this the *boundary cycle*, and we denote this by $\partial\Gamma^\dagger$.

We assign an *oriented* weight $w(e) = w_\Gamma(e) \geq 0$ to every oriented edge of Γ . This defines a continuous-time Markov chain on Γ , where, for every edge (x, y) in Γ , the chain jumps from x to y at rate $w(x, y)$, and with a slight abuse of terminology, we will from now on refer to this chain as the *random walk* on Γ . Note that we defined it in continuous time for simplicity, but we will in fact always be interested only in the geometry of its paths, so its time parametrisation is mostly irrelevant. With a slight abuse of notations, we will often identify the trajectory of the random walk (or other processes, such as loop-erased paths) with the continuous path obtained from the set of edges used by the walk, either seen as a path or as a subset of M . Similarly, with a slight abuse of notation, we will use direct image notations to denote sub-paths of the random walk; i.e., if $s < t$ are two times, we will denote by $X[s, t]$ the path followed by the random walk between times s and t . Finally, we will always stop the random walk whenever it reaches a boundary vertex, so we assume that all oriented edges starting at such a vertex have weight 0. We emphasise that there are no weights on the dual graph Γ^\dagger , and indeed, we will never consider the random walk on the dual graph.

Remark 2.3. To explain the asymmetry between Γ and Γ^\dagger , we point out that Γ can be thought of as a graph that is wired at the boundary of ∂M , whereas Γ^\dagger has free boundary conditions.

In order to apply locally the bijection between dimers and (a variant of) spanning trees, we need to define the superposition of Γ and Γ^\dagger . We introduce a new set of vertices W at each point where some edge e and dual edge e^\dagger intersect. Define the graph \widehat{G} whose vertices are $V = V(\Gamma) \cup V(\Gamma^\dagger) \cup W$ and whose edges are obtained by joining every vertex in W which is on the edge e with the endpoints of e and e^\dagger . (In other words, each pair of edges e and e^\dagger corresponds to four edges in \widehat{G} .) We call \widehat{G} the *superposition graph*. Then, we let G denote the graph where we remove from \widehat{G} the boundary vertices $\partial\Gamma$ and its corresponding half-edges. See Figure 1 for an illustration. Clearly, G is a bipartite graph where we take W to be the set of white vertices and the rest to be black. Also, every non-outer face of G is bounded by 4 edges (i.e., every non-outer face is a quadrangle). We can naturally define *unoriented weights* on G from the ones on Γ . If w is the white vertex in the middle of the edge (x, y) of Γ , we set $w_G(x, w) = w_\Gamma(x, y)$ and $w_G(y, w) = w_\Gamma(y, x)$ while $w_G(e) = 1$ if e is part of an edge of Γ^\dagger . It will be useful later for the definition of the height (see Section 4.1) to also choose in each face f of \widehat{G} a *diagonal* $d(f)$ which is just a curve connecting the primal and dual vertices adjacent to f without exiting f , together with a *midpoint* $m(f)$ which is just a point in the interior of the diagonal. If the primal vertex adjacent to f is a boundary point, choose the diagonal as a path connecting the boundary and the dual vertex. It will be clear from Section 4.1 that the choice of these diagonals and midpoints only affects some of the arbitrary conventions needed to define the height function and that any statement on the centred height $h - \mathbb{E}(h)$ is independent of this choice.

We will prove in Section 3 that if the manifold M is not a torus or an annulus, one must introduce *punctures* in the graph G for it to even admit dimer coverings. More precisely, we will remove finitely many white vertices from G and call a removed vertex and the edges associated with it a *puncture*. (Equivalently, the puncture may be thought of as forcing the vertex to be a monomer instead of in a dimer edge.) Note that each puncture creates an octagon G consisting of four white vertices, two black vertices from Γ , and two black vertices from Γ^\dagger (see Figure 2). We will call the resulting graphs, respectively, $G', \Gamma', (\Gamma^\dagger)'$ and we call the manifold obtained by removing the white vertices M' . We will see later in Section 3 that the number of punctures we remove must be $k = |\chi|$, where χ is the Euler characteristic of M . When this is the case, we will call G' a *Temperleyan graph* on M or *Temperleyan discretisation* of M .

When considering scaling limits, we will of course consider a sequence of graphs $\Gamma^{\#\delta}, (\Gamma^\dagger)^{\#\delta}, G^{\#\delta}$ from the setup above. We now introduce our assumptions on such a sequence.

Let $p : \widetilde{M} \rightarrow M$ be a map which is both analytic and a regular covering of M by its universal cover \widetilde{M} , which is either the unit disc $\mathbb{D} \subset \mathbb{C}$ or the whole plane \mathbb{C} . (Recall the discussion in Section 2.2 for the existence of this map.) In the end, the




Figure 2. Removing a white vertex (puncture) to create the dimer graph G' .

choice of this covering map does not affect the following assumptions; see Lemma 2.5 for a precise discussion about this covering map and also about the covering map for the punctured manifold.

Recall from Section 2.1 that we equipped M with a distance function d_M derived from a Riemannian metric, which extends to the closure \bar{M} ; we still denote this extended distance function by d_M .

- (i) (*Bounded density*) There exists a constant C independent of δ such that, for any $x \in M$, the number of vertices of $\Gamma^{\#\delta}$ in the ball $\{z \in M : d_M(x, z) < \delta\}$ is smaller than C .
- (ii) (*Good embedding*) The edges and diagonals of the graph are embedded as smooth curves and, for every compact set $K \subset \tilde{M}$, the intrinsic winding of every edge or diagonal in the lift $\tilde{\Gamma}^{\#\delta}$ intersecting K is bounded by a constant $C = C_K$ depending only on K . (Note that this allows edges to wind quite a bit near holes.)
- (iii) (*Invariance principle*) As $\delta \rightarrow 0$, the continuous-time random walk $\{\tilde{X}_t\}_{t \geq 0}$ on $\tilde{\Gamma}^{\#\delta}$ starting from the nearest vertex to 0 satisfies

$$(\tilde{X}_{t/\delta^2})_{t \geq 0} \xrightarrow[\delta \rightarrow 0]{(d)} (B_{\phi(t)})_{t \geq 0},$$

where $(B_t, t \geq 0)$ is a two-dimensional standard Brownian motion in \tilde{M} (killed when it leaves \tilde{M} , if $\tilde{M} = \mathbb{D}$) starting from 0, and ϕ is a nondecreasing, continuous, possibly random function satisfying

$$\phi(0) = 0 \quad \text{and} \quad \phi(\infty) = \infty.$$

The above convergence holds in law in Skorokhod topology.

We remark that the above condition is equivalent to asserting that simple random walk from some fixed vertex converges to Brownian motion on the Riemann surface itself up to time parametrisation (see, e.g., [22]).




Figure 3. An illustration of the crossing condition.

- (iv) (*Uniform crossing estimate*). Let R be the horizontal rectangle $[0, .3] \times [0, .1]$ and R' the vertical rectangle $[0, .1] \times [0, .3]$. Let

$$B_1 := B((.05, .05), 0.025)$$

be the *starting ball* and $B_2 := B((.25, .05), .025)$ the *target ball* (see Figure 3).

The uniform crossing condition is the following. There exist universal constants $\delta_0, \alpha_0 > 0$ such that, for every compact set $K \subset M$, there exists a δ_K such that for all $\delta \in (0, \delta_K)$ the following is true. Let $\tilde{K} = p^{-1}(K)$ be the lifts of K . Let R'' be a set of the form $cR + z$, where $c \geq \delta/\delta_0$ and $z \in \mathbb{R}^2$ (i.e., a scaling and translate of R), and is, such that $R'' \subset \tilde{K}$.

Let $B''_1 = cB_1 + z$ and $B''_2 = cB_2 + z$. For all $v \in \tilde{\Gamma}^{\#\delta} \cap B''_1$,

$$\mathbb{P}_v(\tilde{X} \text{ hits } B''_2 \text{ before exiting } R'') > \alpha_0.$$

We emphasise that this crossing condition is defined in the Euclidean metric in the disc, and not the (perhaps more standard) hyperbolic metric when dealing with the universal cover of a hyperbolic surface.

In what follows, sometimes for a compact set $S \subset \tilde{M}$, we will write δ_S to mean $\delta_{p(S)}$ as defined above.

- (v) (*Punctures*) We remove $k = |\chi|$ many points from (the interior of) M and call the resulting surface M' , where χ is the Euler characteristic of M . We assume that each of the (k many) monomers (white vertices removed from \hat{G} to obtain G') converge to a unique puncture as $\delta \rightarrow 0$ which is a given point in the interior of the Riemann surface. We also assume without loss of generality that, for all δ , the punctures are sufficiently far (say at graph distance at least some large fixed constant) from the boundary.

Finally, we also suppose that the following holds. Let u_i, v_i be the endpoints of the edge of Γ corresponding to the white monomer removed for $1 \leq i \leq k$. Then, there exist paths $\gamma_{u_i}, \gamma_{v_i}$ (viewed as an ordered collection of adjacent




Figure 4. An illustration of the assumption in item (v). The paths $\gamma_{u_1}, \gamma_{v_1}$ consist of p and q , which decompose the surface into a number of disconnected annuli: three in the first case and two in the second. The dotted edges are the ones removed from G to get G' .

edges $((x_1, x_2), (x_2, x_3), \dots, (x_{r-1}, x_r))$ or as a continuous path depending on context) satisfying the following. The paths start with $x_1 = u_i$ or $x_1 = v_i$, respectively, for $\gamma_{u_i}, \gamma_{v_i}$, and we assume that

$$M \setminus \bigcup_{1 \leq i \leq k} (\gamma_{u_i} \cup \gamma_{v_i} \cup (u_i, v_i)) \tag{2.1}$$

is topologically a disjoint union of annuli.

Remark 2.4. Typically, the existence of the paths γ_{u_i} and γ_{v_i} for δ small enough is guaranteed by the crossing condition and the *pants decomposition* of M . In other words, the paths decompose the surface into disjoint annuli; see Figure 4. However, it is possible to construct pathological examples (for large δ) where there are no paths such that (2.1) holds, so this needs to be made as a separate assumption. In fact, we will see that (v) is equivalent to the graph G admitting a dimer cover (see Figure 1), so we could equivalently assume the latter instead – this would be more elegant but less concrete.

Lemma 2.5 (Conformal invariance of assumptions). *The assumptions (ii), (iii), and (iv) on the sequence $(\Gamma^{\#\delta})_{\delta>0}$ are invariant with respect to the choice of the covering map p in the sense that if p' is any other analytic and regular covering of M , then these assumptions are satisfied for some (possibly different) choices of $\delta_0, \alpha_0, (\delta_K)_{K \subset M}, (C_K)_{K \subset \tilde{M}}$.*

Furthermore, if $(\Gamma^{\#\delta})_{\delta}$ satisfies the assumptions (i)–(v) for a choice of covering map p and constants $C, \delta_0, \alpha_0, (\delta_K)_{K \subset M}, (C_K)_{K \subset \tilde{M}}$, then these assumptions are invariant with respect to conformal transformations in the sense that if

$$\psi : (M, x_1, x_2, \dots, x_k) \rightarrow (N, x'_1, x'_2, \dots, x'_k)$$

is a conformal bijection mapping M to N and x_i to x'_i for $1 \leq i \leq k$, then $(\psi(\Gamma^{\#\delta}))_{\delta>0}$ also satisfies assumptions (i)–(v) for the choice of the covering map $p \circ \psi$ and constants $C, \delta_0, \alpha_0, (\delta_{\psi^{-1}(K)})_{K \subset M'},$ and $(C_K)_{K \subset \tilde{M}}$.

Proof. We start with the second of these two assertions, for which items (ii)–(v) are trivial. For (i), if M has no boundary, then $d_M(x, y) = d_N(\psi(x), \psi(y))$ (essentially because Möbius maps are isometries with respect to the hyperbolic metric on the unit disc). Thus, (i) holds. If M has a boundary, then note that the conformal bijection ψ between M and N extends to a conformal bijection $\bar{\psi}$ between the doubles of M, N ; and we are back to the previous case.

Let us now prove the first assertion. Note that, by the uniformisation theorem, there exists a conformal bijection (Möbius map) $\phi : \tilde{M} \rightarrow \tilde{M}$ such that $p \circ \phi = p'$. Note that ϕ is a Möbius map from the unit disc to itself in the hyperbolic case and from the complex plane to itself in the torus case. Assumption (ii) is trivially preserved in the torus case as Möbius maps are translations and rotations. For the hyperbolic case, we can employ the change in winding under conformal map formula (Lemma 2.2), and note that derivatives of Möbius maps are bounded on compact subsets of the open unit disc. In assumption (iii), we only require convergence of the random walk up to time change which is preserved under Möbius maps. Assumption (iv) is easily seen to be preserved in the torus case as a Möbius map maps a rectangle to another rotated and translated rectangle, which can be crossed by concatenating bounded number of vertical and horizontal rectangles of smaller scales. In the hyperbolic case, a rectangle is mapped to a domain bounded by four circular arcs, and the starting and target discs are mapped to discs inside this domain. Furthermore, compact sets are mapped to compact sets. Thus, concatenating domains of this type, it is easy to see that uniform crossing estimate holds for M , perhaps for different positive constants $\delta_0, \alpha_0, (\delta_K)_{K \subset M}$. ■

Remark 2.6. The invariance principle assumption (item (iii)) actually implies something stronger in combination with the other assumptions: for any point x in \tilde{M} , the random walk starting from a vertex $x^{\#\delta}$ nearest to x converges to a Brownian motion starting from x up to a time change as above. This is a consequence of the fact that random walk from 0 comes close to x with uniformly positive probability using the crossing estimate and the strong Markov property of Brownian motion.

In case $\partial M \neq \emptyset$, recall that the set of boundary cycles $(\partial\Gamma^\dagger)^{\#\delta}$ corresponds to the connected components of ∂M . One consequence of the invariance principle assumption (item (iii)) is that each boundary cycle converges in the Hausdorff metric (induced by d_M) to the associated component of ∂M .

Sometimes, we drop the superscript δ from $\Gamma^{\#\delta}, (\Gamma^\dagger)^{\#\delta}$ for clarity, when there is no possibility for confusion.

Remark 2.7. In the hyperbolic case, note also that while we have stated the assumptions (ii), (iii), and (iv) on the universal cover of M , these assumptions are also valid for the sequence $(\Gamma^{\#\delta})_{\delta>0}$ on the universal cover of the punctured surface M'

as well if we endow it with the metric it inherits from M via inclusion. This can be checked using the fact that there is a map from the universal cover of M' to $\tilde{M} \setminus p^{-1}(\{x_1, \dots, x_k\})$ (where x_1, \dots, x_k is the set of punctures) which is locally a conformal bijection. Also, note that (i) is also satisfied in M' trivially by the choice of the metric.

2.5. Height function and forms

A flow ω is a real-valued antisymmetric function on the oriented edges \vec{E} of G ; i.e., for every oriented edge (u, v) , $\omega(u, v) = -\omega(v, u)$. The total flow out of a vertex v is defined to be $\sum_{w \sim v} \omega(v, w)$. Similarly, the total flow into a vertex v is defined to be $\sum_{w \sim v} \omega(w, v)$. A flow f is a closed 1-form if the sum over any oriented contractible cycle is 0: i.e., for any oriented cycle $(v_0, v_1, \dots, v_n = v_0)$ in G so that the embedding of $\bigcup_{i=0}^{n-1} (v_i, v_{i+1})$ in M forms a contractible loop,

$$\sum_{i=0}^{n-1} \omega(v_i, v_{i+1}) = 0.$$

It is clear that if M is simply connected, then there exists a function f on the vertices of G (uniquely defined up to a global constant) such that $f(v) - f(u) = \omega(u, v)$ for all $u \sim v$.

We now associate to any dimer configuration \mathbf{m} on G a closed 1-form on \vec{E} . Let \mathbf{m} be a dimer configuration on G , and let $\vec{e} = (w, b)$ be an oriented edge, where w is a white vertex and b a black vertex. We define the flow $\omega_{\mathbf{m}}$ by setting $\omega_{\mathbf{m}}(\vec{e}) = 1_{\{e \in \mathbf{m}\}}$. Also, $\omega_{\mathbf{m}}$ is defined in an antisymmetric way: $\omega_{\mathbf{m}}((b, w)) = -\omega_{\mathbf{m}}((w, b))$. Note that the total flow out of a white vertex is 1 and that out of a black vertex is -1 .

To any flow ω on oriented edges one can associated a dual flow ω^\dagger defined on the oriented edges of the dual graph G^\dagger , where if e^\dagger crosses the edge $e = (w, b)$ with w on its right and b on its left, then we set $\omega^\dagger(e^\dagger) = \omega(e)$. Note also that if ω is divergence free (i.e., the flow out of every vertex is 0), then ω^\dagger is a closed 1-form on \vec{E}^\dagger .

Consider any reference flow ω_0 which has total flow out of white vertex equal to 1 and total flow out of a black vertex equal to -1 . Then, $\omega = \omega_{\mathbf{m}} - \omega_0$ defines a divergence free flow on \vec{E} . We call ω^\dagger the height 1-form corresponding to \mathbf{m} with reference flow ω_0 .

When G is embedded on a simply connected domain (so that, in particular, no cycle in G^\dagger is non-contractible), every closed 1-form ω on \vec{E}^\dagger becomes exact: i.e., there exists a function on the faces $F(G)$ of G , $h : F(G) \mapsto \mathbb{R}$ so that, for any two adjacent faces f, f' ,

$$h(f') - h(f) = \omega(f, f').$$

Observe further that this function is defined only up to a global constant. The function h is then called the *height function* of the dimer \mathbf{m} , admitting an abuse of terminology.

We recall the following simple but useful observation. A *path* in G (or G^\dagger) is a sequence of vertices (v_0, \dots, v_n) (or faces (f_0, \dots, f_n)) of G so that v_i is adjacent to v_{i+1} (or f_i is adjacent to f_{i+1} in G^\dagger) for all $0 \leq i \leq n - 1$.

Lemma 2.8 (Unique path-lifting property). *Let $\gamma = (f_0, f_1, \dots, f_n)$ be a path (not necessarily simple) in G^\dagger . Let \tilde{f}_0 be a lift (i.e., one pre-image) of f_0 to \tilde{M} . Then, there exists a unique path $\tilde{\gamma} = (\tilde{f}_0, \tilde{f}_1, \dots, \tilde{f}_n)$ in \tilde{M} so that \tilde{f}_i is the lift of f_i . Further,*

$$\sum_{i=0}^{n-1} \omega(\tilde{f}_i, \tilde{f}_{i+1}) = h(\tilde{f}_n) - h(\tilde{f}_0).$$

We now turn to the definition of height function in the more complicated case when M is no longer assumed to be simply connected. In that case, when we sum the values of the height 1-form (defined above) along any non-contractible cycle, we may get a nonzero value. One can use the *Hodge decomposition theorem* to isolate out the part of the height 1-form which is encoded by the topology of the underlying surface. The Hodge decomposition theorem works in great generality, but in the present context, it takes the following simple form. For any function f on the vertices of G , we define df to be the closed 1-form defined on \vec{E} as

$$df(u, v) = f(v) - f(u).$$

A *harmonic 1-form* \mathfrak{h} is a closed 1-form which is divergence free so that

$$\sum_{v \sim u} \mathfrak{h}(u, v) = 0.$$

Theorem 2.9 (Hodge decomposition [2, 8, 31]). *For any closed 1-form ω on G (or G^\dagger), there exist a function f on the vertices of G and a harmonic 1-form \mathfrak{h} defined on \vec{E} such that*

$$\omega = df + \mathfrak{h},$$

and f is unique up to an additive global constant, and \mathfrak{h} is unique. Furthermore, \mathfrak{h} is completely determined by summing ω over a finite set of oriented non-contractible cycles which forms the basis of the first homology group of M .

In this paper, we will analyse (f, \mathfrak{h}) corresponding to the divergence free flow $\omega_{\mathbf{m}} - \omega_0$, where \mathbf{m} is dimer configuration chosen from the law (1.1) subject to certain natural conditions and ω_0 is a carefully chosen reference flow. (In fact, we will consider the height function; see Section 6 for a precise statement.) We will call \mathfrak{h} the *instanton component*. We remark that changing the reference flow changes (f, \mathfrak{h})

by a deterministic additive factor and in particular does not affect the fluctuations of (f, \mathfrak{h}) around their mean. To be more precise, our main theorem (Theorem 6.1) will be stated in terms of the single-valued function associated to $\omega_{\mathbf{m}} - \omega_0$ on the universal cover of M' . See Remark 6.16 for a statement concerning both the scalar and instanton parts of the height function.

Throughout the paper, rather than working with the scalar and instanton components of the height 1-form, it will be more convenient to lift the height 1-form ω to the universal cover of M . Since the latter is always simply connected, this allows us to work with actual functions without having to worry about the Hodge decomposition Theorem 2.9. We will then check that the convergence of height function on the universal cover implies convergence of each of the components in the Hodge decomposition.

Our assumptions on the graph G where the dimer model lives are such that $\tilde{G} = p^{-1}(G)$ is a planar graph embedded on \tilde{M} . Moreover, the height 1-form ω on the dual edges of G lifts to a height 1-form $\tilde{\omega}$ on the dual edges of \tilde{G} . Since \tilde{M} is simply connected and since $\tilde{\omega}$ is a closed one-form on the dual edges of \tilde{G} (this is a local property, so remains true when we lift to \tilde{G}), we can define a height function $h = h(\mathbf{m}, G)$ (up to a global constant) on the dual graph \tilde{G}^\dagger . The instanton component \mathfrak{h} can be related to the height function h on the universal cover by summing up the value of $\tilde{\omega}$ along any path in the dual graph of \tilde{G} corresponding to a non-contractible loop in the dual edges of G . This is easier to explain on an example.

Example 2.10. If M is the flat torus $\mathbb{T} := \mathbb{C}/(\mathbb{Z} + \tau\mathbb{Z})$ for some complex number τ with $\Im(\tau) > 0$, then the universal cover is the complex plane \mathbb{C} . The universal cover can be thought of as many copies of the fundamental domain (a parallelogram determined by 1 and τ).

Fix v_0 in the fundamental domain. Then, by periodicity of dh , the height function h on \tilde{G} evaluated at a point $v = v_0 + m + \tau n$ (where $m, n \in \mathbb{Z}$) is given by

$$h(v) = h(v_0 + m + \tau n) = h(v_0) + am + bn,$$

for some $a, b \in \mathbb{R}$ which do not depend on either v_0, m, n . Let us describe what a, b are. Consider the two loops on the torus described by $L_1 := (t + 1/2\tau : t \in [0, 1])$ and $L_2 := (1/2 + t\tau : t \in [0, 1])$ in the fundamental domain (i.e., L_1 and L_2 are the two non-contractible loops in the torus which form the basis of the homology group). Then, a is the sum of the values of ω along any loop in \tilde{E}^\dagger which is homotopic to L_1 , whereas b is the sum of the values of ω along any loop which is homotopic to L_2 . Clearly, the choice of these curves in \tilde{G}^\dagger does not matter since the height 1-form is closed. Furthermore, in the Hodge decomposition of Theorem 2.9, the harmonic 1-form \mathfrak{h} is uniquely determined by the numbers a and b .

3. Temperley's bijection on Riemann surfaces

3.1. Notion of Temperleyan cycle-rooted spanning forest: Bijection

The goal of this section is to extend the classical Temperley bijection to graphs embedded in surfaces of higher genera. The bijection is between dimer covers of certain natural class of graphs introduced in Section 2.4 and certain objects which we call *Temperleyan forests*. The main result is formulated in Theorem 3.3.

Before defining the objects in question, let us briefly recall the classical Temperley bijection on simply connected surfaces [14]. The bijection map is local: given a dimer cover, one “extends” (in the direction from black to white) the dimer to an oriented edge as depicted in Figure 6. This yields two oriented trees which are dual to each other. In the opposite direction, given a pair of trees which are dual to each other, one can naturally orient them towards a pre-assigned “root vertex”; the dimer cover then consists of the first half of each oriented edge in the two dual trees. The primary issue that needs to be dealt with in higher genus is that the relevant objects are no longer pairs of trees that are dual to each other but pairs (t, t^\dagger) of *arborescences* (i.e., collection of oriented edges such that there is a single outgoing edge out of every vertex, except for boundary vertices) on Γ and Γ^\dagger , respectively, dual to each other, and where each component contains a non-contractible oriented cycle. Thinking only about the arborescence t on the primal graph, which we aim to ultimately understand via a variant of Wilson's algorithm, the requirement that its planar dual t^\dagger can be endowed with the orientation of an arborescence imposes nontrivial topological constraints on t which will be elucidated below.

Let Γ be a graph faithfully embedded on a surface with a certain specified set of boundary vertices. Assume that every edge e of Γ comes with a specified weight $w(e) > 0$. Before introducing the notion of Temperleyan forest, we start with the simpler notion of cycle-rooted spanning forest.

Definition 3.1. A *wired oriented cycle-rooted spanning forest* (which we abbreviate as wired oriented CRSF) of Γ with the specified boundary is an oriented subgraph t of Γ , where the following hold.

- Every non-boundary vertex of Γ has exactly one outgoing edge in t . Every boundary vertex has no outgoing edge. (As a result, any cycle of t must be consistently orientated.)
- Every cycle of t is non-contractible.

To each CRSF t we may associate the weight $\prod_{e \in t} w(e)$. We will consider the probability measure on CRSFs such that the probability of t is by definition proportional to its weight.

This is equivalent to the notion of essential CRSF on a graph with wired boundary introduced by Kassel and Kenyon [24]. Ignoring the orientation of t gives an unoriented graph its connected components will simply be called the *connected components* of t without any additional precision. Note that if t is a wired oriented CRSF, every connected component of t contains at most one cycle: more precisely, every boundary component must have zero cycles, while every non-boundary component contains exactly one cycle.

We will refer to the set of all non-contractible cycles of a wired oriented CRSF to mean the set of unique cycles corresponding to each (non-boundary) component of the wired oriented CRSF.

Let us come back to the setup of Section 2.4 and recall that we had the graph Γ , its dual Γ^\dagger , and the superposition graph G , all embedded nicely in a Riemann surface with g handles and b holes. Compared to the planar setting, there is an immediate topological difficulty, which is that this graph G in general does not admit a dimer cover. Indeed, by Euler’s formula, if G has v vertices, e edges, and f many contractible faces, then

$$v - e + f = \chi := 2 - 2g - b,$$

where χ is the Euler’s characteristic. On the other hand, if G admits a dimer cover, one must have $e = v + f$ since e is the number of white vertices which must match $v + f$ (the number of vertices and contractible faces of Γ^\dagger). Thus, we need to remove $k = |\chi|$ many edges from the superposition graph for it to have any chance of having a dimer cover (see Figure 2). These removed edges can be thought of as creating *punctures* in the surface. Call the graph with these punctures G' . (Recall that, as per item (v), we assume throughout that the punctures are macroscopically far apart in M and are away from the boundary as well and converge as $\delta \rightarrow 0$ to a set of pairwise distinct points in M .) Note that if M is a torus or an annulus, then $\chi = 0$, so no punctures need to be removed. See also Ciucu–Krattenthaler [13] and Dubédat [19] for other situations where punctured dimers arise.

To every wired oriented CRSF t of Γ' with boundary $\partial\Gamma$ one can associate a natural dual *free cycle-rooted spanning forest* t^\dagger (abbreviated free CRSF) of $(\Gamma^\dagger)'$ as follows. The vertices of t^\dagger are given by the vertices of $(\Gamma^\dagger)'$ (i.e., it spans $(\Gamma^\dagger)'$), and an edge e^\dagger is present in t^\dagger if and only if its dual e is absent in t . Note that *a priori* t^\dagger does not come with an orientation and that its cycles can overlap (or equivalently a component might contain several cycles); see Figure 5 for an example. In fact, it is not hard to see that, in cases similar to Figure 5, it is not possible to consistently orient the free CRSF, so there is no chance to apply any version of Temperley’s bijection. This motivates the introduction of the following definition.

Definition 3.2. We say that the wired oriented CRSF t is *Temperleyan* if each connected component of t^\dagger contains exactly one cycle.




Figure 5. A non-Temperleyan CRSF (in blue). The surface M is the “pair of pants”: a domain of the plane with two holes (in grey in the picture). In this example, t does not contain any cycle, and hence, any connected component flows to the boundary of M . Its dual t^\dagger must contain a component with two cycles which overlap. The cycles go around each of the two holes and must be connected, as otherwise there would have to be a path in t separating them; however, this is impossible as such a path would have to connect two distinct boundary points. So, t is *not* Temperleyan. Note that $\chi = 2 - 2g - b = -1$ here. See Lemma 3.5 for a more general argument.

An example of a wired oriented CRSF t that is *not* Temperleyan is provided in Figure 5.

Temperleyan forests also come with a natural law, which is simply the law of the CRSF conditioned on the event that each component of the dual contains exactly one cycle. We will soon formulate a more explicit criterion for a CRSF t to be Temperleyan (see Section 3.2 and in particular Theorem 3.7). For now, observe that if t is a Temperleyan wired oriented CRSF and t^\dagger is its dual, we can assign an orientation to each cycle in each component of t^\dagger arbitrarily from one of the two possible choices. Then, we orient all other edges of t^\dagger towards the cycle of that component. We call (t, t^\dagger) a Temperleyan pair or *Temperleyan self-dual pair* if t is a Temperleyan CRSF and t^\dagger is its dual with a choice of orientation as above. We warn the reader that the notation slightly obscures the choice of orientation needed to go from t to t^\dagger in order to highlight the duality. Hopefully, this should not introduce any confusion since we will always want to think of our CRSF as being oriented.

We now state the Temperley bijection for general surfaces. Recall that we assign *oriented* weights w_e to each edge e in Γ' , no weight (or unit weight) to edges of Γ^\dagger and that this turns G' into a weighted *unoriented* graph. Indeed, if $e = (x, y)$ is an oriented edge of Γ' , let w denote the white vertex in the middle of e . Then, we assign to the unoriented edge $\{x, w\}$ of G' the weight $w_{(x,y)}$ and to the edge $\{w, y\}$ of G' the weight $w_{(y,x)}$.




Figure 6. Illustration of the local transformation in Temperley’s bijection. Left: a dimer configuration on a superposition graph. Right: a collection of oriented edges forming dual oriented CRSF.

We define the measure \mathbb{P}_{Temp} as the measure on Temperleyan pairs with weights inherited from the weights of Γ , and we will call $(\mathcal{T}, \mathcal{T}^\dagger)$ a generic associated random variable, i.e.,

$$\mathbb{P}_{\text{Temp}}((\mathcal{T}, \mathcal{T}^\dagger) = (t, t^\dagger)) = \frac{1}{Z_{\text{Temp}}} 1_{\{(t, t^\dagger) \text{ Temperleyan pair}\}} \prod_{e \in t} w(e), \quad (3.1)$$

where Z_{Temp} is the partition function.

Theorem 3.3 (Temperley bijection on general surfaces). *Let $M, \Gamma', (\Gamma^\dagger)', G'$ be as in Section 2.4. Then, there exists a bijection ψ between the set of Temperleyan pairs and the set of dimer configurations on G' . Furthermore, if $(\mathcal{T}, \mathcal{T}^\dagger)$ has the law (3.1), then $m = \psi((\mathcal{T}, \mathcal{T}^\dagger))$ has law (1.1) with unoriented weights on G' described above.*

We emphasise that it is unclear at this point whether the measure $Z_{\text{Temp}} \mathbb{P}_{\text{Temp}}$ (and consequently the measure on dimers) is not trivially zero as we have only deduced that the construction of G' is necessary but have not deduced that it is sufficient. We prove this later in Proposition 3.4.

Proof of Theorem 3.3. We assume that the set of Temperleyan CRSF and the set of dimer covers are both nonempty. (This assumption is validated in Proposition 3.4.) Given a Temperleyan pair (t, t^\dagger) , we obtain a configuration of edges $m = \psi((t, t^\dagger))$ as follows: for every oriented edge $\vec{e} \in t$ (resp., $\vec{e} \in t^\dagger$), we can write $\vec{e} = e_1 \cup e_2$, where e_1, e_2 are the first and second halves of \vec{e} and set $e_1 \in m$ (see Figure 6). The resulting m is a matching on G' because every (non-boundary) vertex has a unique outgoing edge in either t or t^\dagger . Furthermore, since $t \cup t^\dagger$ spans the black vertices of G' , the matching is a perfect matching.

Also, ψ is injective: if (t_1, t_1^\dagger) and (t_2, t_2^\dagger) are distinct, then there must be a black vertex v on G' (i.e., a vertex of Γ' or $(\Gamma^\dagger)'$) such that the unique outgoing edge from v in t_1 or t_1^\dagger is different from the unique outgoing edge from v in t_2 or t_2^\dagger . Hence, v will be matched to two distinct white vertices in $\psi((t_1, t_1^\dagger))$ and $\psi((t_2, t_2^\dagger))$.

We now check ψ is onto. Given a matching m of G' , we can obtain a pair (t, t^\dagger) by extending the matched edges: for every non-boundary black vertex b of G' , let w be the white vertex matched with b in m . By construction, (bw) is a part of a unique edge $e = (bb')$ (either of Γ or Γ^\dagger), and we include the edge (bb') (oriented from b to b') in t or t^\dagger as appropriate. The fact that neither t nor t^\dagger contain contractible cycles follows from the same argument as in the standard, planar case: if, say, t has a contractible cycle C , then since $\Gamma \setminus C$ has a dimer cover, an elementary counting argument then shows that $v - e + f = 0$, where v, e, f are the number of vertices, edges, and faces of the contractible component that $\Gamma \setminus C$. On the other hand, Euler's formula implies $v - e + f = 1$ (again excluding the outer face). This shows that all cycles are non-contractible. Since every vertex v of t^\dagger has a unique outgoing edge, t^\dagger must have exactly one cycle per component, and so, t is Temperleyan. It also easily follows from the definition of the weights of the Temperleyan tree and the dimer covers that the bijection is weight preserving. This concludes the proof. ■

3.2. Criterion for a wired CRSF to be Temperleyan

In this section, we prove the following proposition.

Proposition 3.4. *The graph G' obtained in Section 2.4 has a dimer cover. In particular, the measure \mathbb{P}_{Temp} defined in (3.1) is a probability measure.*

As a by-product of the proof of Proposition 3.4, we derive a simple criterion for a CRSF to be Temperleyan (Theorem 3.7). We start with a lemma about the $\chi = 0$ case.

Lemma 3.5. *Let M be a nice Riemann surface with g handles and b boundary components and $\Gamma, \Gamma^\dagger, G$ be embedded as above (i.e., without punctures). A wired Temperleyan oriented CRSF of Γ exists if and only if M has the topology of either a torus or an annulus. Furthermore, in these cases, all oriented CRSF are Temperleyan.*

Proof. Note that if M has the topology of an annulus (with all the nice properties of Section 2.1), then finding a Temperleyan oriented CRSF is straightforward. Indeed, any wired spanning forest in the annulus (where both boundaries are wired) is Temperleyan: the dual is a graph containing a single cycle separating the two components touching each boundary. Also, notice that, for a torus, every oriented CRSF is Temperleyan [20]: essentially, an oriented CRSF must contain a cycle (since there is no boundary on a torus), and cutting along this cycle gives us a (bounded) cylinder or equivalently an annulus.

For the converse, observe that the extended Temperley’s bijection (Theorem 3.3) allows us to construct a dimer configuration from the Temperleyan CRSF and its dual (by endowing each cycle with arbitrary orientation and orienting every other edge towards the unique cycle of its component). Notice that although Theorem 3.3 is written for the punctured graph G' , since in the case of torus and annulus $G = G'$, the same proof goes through. Now, recall that a dimer configuration exists without removing any punctures only if $\chi = 0$ or equivalently $2g + b = 2$ (as discussed in Section 3.1). This equation has only two feasible solutions: $g = 1, b = 0$ (i.e., a torus) and $g = 0, b = 2$ (i.e., an annulus). This completes the proof of the “only if” part.

For the last assertion, note that every oriented CRSF in the annulus divides the annulus into several smaller disjoint annuli. The rest follows easily from the arguments in the first paragraph above. ■

Proof of Proposition 3.4. In light of Lemma 3.5, we assume that M is neither a torus nor an annulus. Using assumption item (v), recall that there exist a collection of paths $\gamma_{u_i}, \gamma_{v_i}$ on Γ' for $1 \leq i \leq k$ starting, respectively, from u_i and v_i such that

$$M \setminus \bigcup_{1 \leq i \leq k} (\gamma_{u_i} \cup \gamma_{v_i} \cup (u_i, v_i))$$

is topologically a disjoint union of annuli when we view the paths as continuous curves on M and (u_i, v_i) denotes the edge from u_i to v_i . We start building a dimer configuration \mathbf{m} on G' as follows. First, for each path γ in this collection, which by definition starts from either u_i or v_i for some $1 \leq i \leq k$, we orient it away from the puncture from which it emanates, i.e., away from u_i or v_i . We then consider the collection of dimer edges naturally associated with γ formed by the first halves of each of the oriented edges of γ , as in Figure 6 and the proof of Theorem 3.3. Then, include all these edges in \mathbf{m} .

The remaining disjoint annuli are nice in the sense of Section 2.1, so we can apply Lemma 3.5 to obtain a Temperleyan forest on each of these annuli, and hence by Theorem 3.3 a dimer cover on the part of G' corresponding to these annuli. Patching these together gives us a dimer configuration \mathbf{m} on all of G' , as desired. ■

Remark 3.6. Note that this implies that the assumption in item (v) is, as already mentioned earlier, necessary in order for G' to admit a dimer cover. Indeed, if there is a dimer cover, then by applying the local mapping in Figure 6, starting from the vertices u_i and v_i adjacent to the punctures, we can obtain paths as described in item (v). Each component of the complement of these paths also has a dimer cover but no puncture. According to Lemma 3.5, this is only possible if (2.1) is satisfied.

We now deduce from the above an extremely convenient criterion for a wired oriented CRSF to be Temperleyan. Let us suppose that M is not a torus or an annulus,

whence $k = 2g + b - 2 > 0$. (We already know that every CRSF is Temperleyan otherwise.) Define the branch starting from a primal vertex v of Γ' to be the path obtained by going along the unique outgoing edge from each vertex (which necessarily ends when a loop is formed or a boundary is hit). Recall that, at a puncture (i.e., a white monomer), there are exactly two vertices u_i, v_i of Γ' on either side of the puncture. Let $\mathfrak{B}_{i1}, \mathfrak{B}_{i2}$ be the branches in Γ' of u_i, v_i for $1 \leq i \leq k = 2g + b - 2$. We call the union of these branches, together with the removed edges e_i , the *skeleton* \mathfrak{s} of t . That is,

$$\mathfrak{s} = \bigcup_{i=1}^k (\mathfrak{B}_{i1} \cup e_i \cup \mathfrak{B}_{i2}).$$

Simply put, the skeleton of t is the union of the branches emanating out of the punctures.

Theorem 3.7. *Suppose that M is not a torus. A wired oriented CRSF is Temperleyan if and only if every component of $M \setminus \mathfrak{s}$ has the topology of an annulus.*

(Of course, in the case of an annulus, as already noted, there are no punctures, and so, $\mathfrak{s} = \emptyset$, so the condition is automatically satisfied, as we already know.) The criterion for a CRSF to be Temperleyan is thus just to say that the skeleton cuts the surface into disjoint annuli. Figure 4 gives two examples on a surface M (the “pair of pants”) where $M \setminus \mathfrak{s}$ consists of topological annuli, where \mathfrak{s} is the paths p, q in that figure. (We get three annuli in the first example and two in the second one.)

Proof. Let t be a Temperleyan CRSF, and let t^\dagger be its dual with a choice of orientation. Note that the vertices (in G') of \mathfrak{s} are all matched with each other in the dimer configuration associated to (t, t^\dagger) . Therefore, in each component of $M \setminus \mathfrak{s}$, all vertices are also matched with each other. By Lemma 3.5, this implies that these components are annuli.

Conversely, note that by definition t^\dagger cannot cross \mathfrak{s} , and so, t can be restricted to each component of $M \setminus \mathfrak{s}$ to form a wired CRSF. By the other implication in Lemma 3.5, it follows that t is Temperleyan in each such component, and so, it is globally Temperleyan. ■

3.3. Wilson’s algorithm to generate wired oriented CRSF

Recall the measure \mathbb{P}_{Temp} from (3.1), i.e., the law of $(\mathcal{T}, \mathcal{T}^\dagger)$ on Temperleyan pairs which is naturally associated with the dimer model. We will not study directly \mathbb{P}_{Temp} but rather a slightly altered version which can be sampled through Wilson’s algorithm and is defined as follows: we first sample a wired oriented CRSF of Γ with law

$$\mathbb{P}_{\text{Wils}}(\mathcal{T} = t) = \frac{1}{Z_{\text{Wils}}} 1_{\{t \text{ Temperleyan}\}} \prod_{e \in t} w(e);$$

and given $\mathcal{T} = t$, we pick \mathcal{T}^\dagger , an oriented dual uniformly among all possibilities of orientation of the dual.

Clearly, the only difference between the two is due to the fact that any cycle of the dual t^\dagger of a Temperleyan oriented CRSF t can be oriented in two possible ways to determine a dual pair (t, t^\dagger) . We deduce the following relationship.

Lemma 3.8. *Let (t, t^\dagger) be a Temperleyan pair such that t^\dagger contains exactly k^\dagger non-contractible cycles. Then, the Radon–Nikodym derivative satisfies*

$$\frac{d \mathbb{P}_{\text{Temp}}}{d \mathbb{P}_{\text{Wils}}}(t, t^\dagger) = \frac{Z_{\text{Wils}}}{Z_{\text{Temp}}} 2^{k^\dagger}.$$

In particular, conditioned on having k^\dagger non-contractible cycles for the dual forest, the laws \mathbb{P}_{Temp} and \mathbb{P}_{Wils} coincide.

Let Γ, Γ^\dagger be faithfully embedded on a nice Riemann surface M . We now describe Wilson’s algorithm to generate a wired (but not necessarily Temperleyan) oriented CRSF on Γ . We prescribe an ordering of the vertices (v_0, v_1, \dots) of Γ .

- We start from v_0 and perform a loop-erased random walk until a non-contractible cycle is created or a boundary vertex (i.e., a vertex in $\partial\Gamma$) is hit.
- We start from the next vertex in the ordering which is not included in what we sampled so far and start a loop-erased random walk from it. We stop if we create a non-contractible cycle or hit the part of vertices we have sampled before.

There is a natural orientation of the subgraph created since from every non-boundary vertex there is exactly one outgoing edge through which the loop-erased walk exits a vertex after visiting it. Let $\tilde{\mathbb{P}}_{\text{Wils}}$ be the law of the resulting wired oriented CRSF.

Proposition 3.9. *We have*

$$\tilde{\mathbb{P}}_{\text{Wils}}(t) = \frac{1}{\tilde{Z}_{\text{Wils}}} \prod_{e \in t} w(e). \tag{3.2}$$

In particular, $\tilde{\mathbb{P}}_{\text{Wils}}$ generates a wired oriented CRSF of Γ as described by Definition 3.1 with law given by (3.2). Furthermore, conditionally on being Temperleyan, $\tilde{\mathbb{P}}_{\text{Wils}}$ coincides with the first marginal of \mathbb{P}_{Wils} .

Proof. This follows from Theorem 1 and Remark 2 of Kassel–Kenyon [24]. ■

As a consequence of the result of this section, we obtain a relatively usable description of the law \mathbb{P}_{Temp} of interest: start with $\tilde{\mathbb{P}}_{\text{Wils}}$, which can be sampled via Wilson’s algorithm, and condition on the topological event described in Theorem 3.7. (In the second paper [5], we denote this conditioning event by $\mathcal{A}_{\mathcal{M}}$.) Finally, bias the law by the Radon–Nikodym derivative coming from Lemma 3.8. One of the main

difficulties in [5] stems from the fact that $\mathcal{A}_{\mathcal{M}}$ is asymptotically degenerate (i.e., the probability of $\mathcal{A}_{\mathcal{M}}$ for a wired CRSF tends to zero as the “mesh size” δ tends to zero). The Radon–Nikodym derivative in Lemma 3.8 on the other hand is relatively benign and does not introduce much further complication.

4. Winding and height function

In this section, we explain the connecting between winding of the Temperleyan forest and height function in our setup. In the process, we develop a systematic way to relate these two notions in a very general setting which we believe is of interest even in simply connected domains.

A first difficulty is that measuring angles and computing the total angle along say, a simple loop is well known to be more complicated on a surface than on a plane. In the following, we will sidestep completely this issue by only establishing the connection for the lifts of the Temperleyan forest and the height function where we will be able to use the usual planar theory of angles and winding. This will also have the advantage that the dimer height become an actual function instead of a one-form.

Using this approach, we are left with a situation analogous to the original generalised Temperley’s bijection in [28]: there is a fixed graph with a given planar embedding and we have to relate the dimer height to the winding of the associated “tree” at a deterministic level. (We also point out that a version on a torus already appeared in [38].) In particular, the fact that in our current setting the graphs (and trees) arise as lifts and therefore have many symmetries should not play a significant role. (Recall that [28] is valid in great generality.)

There are however still significant technicalities to work out. First, we cannot assume that our embedding uses only straight lines (since in the hyperbolic case, the lift of an edge cannot be a straight line simultaneously in different copies of the fundamental domain). Second, our “tree” objects are actually forests, so the winding between vertices belonging to different components must be properly defined. Third, we need to check that the objects defined on the lift can be reasonably mapped back to the surface and state any simplification of the general theory that would arise specifically in our setting.

Finally, the way the connection is phrased in [28] is relatively unwieldy, so it would be nice to make it more “user-friendly”, for example, to make it symmetric with respect to the primal and dual forests. Recall that the winding field is naturally defined on either primal or dual vertices, while the dimer height function is naturally defined on faces of the superposition graph, so we need some natural local rule to move between faces and vertices. Our solution for this technicality will be to extend

both the primal and dual trees along the diagonal of each face, up to the midpoint of that face, a construction which results in what we call augmented trees. We will see that, with this construction, the primal and dual trees play a completely symmetric role, as desired.

Our strategy in this section will be to proceed by successive generalisations from the definition of a winding field on a single deterministic embedded tree to dual pairs of spanning tree and the relation to dimer height function in Section 4.1. Section 4.2 covers the situation specific to our Riemann surface setting, with the main result of this section stated as Theorem 4.10. Finally, Section 4.3 treats the relation between the instanton component and the set of non-contractible loops.

4.1. Winding field of embedded trees and choice of reference flow

Embedded tree. We temporarily forget about dimers and Temperleyan forests and focus on how to compute the *winding field* of a *deterministic* tree.

Consider a graph (finite or infinite) which is a tree oriented towards either a leaf or one of its ends. Embed it on \mathbb{C} with smooth edges (although having points of non-differentiability at the vertices, which we call *corners*, is allowed so that the branches are only piecewise smooth). Call the embedding \mathcal{T} ; we want to think of it as a union of smooth curves in \mathbb{C} . Recall that, using Lemma 2.1, the intrinsic winding of any finite self-avoiding path in this tree is well defined. From every point x on \mathcal{T} , we let γ_x be the unique path from x following the orientation (so it stops either at a leaf or at the end towards which the tree is oriented) and call it the branch of x . Suppose for now that

$$|W_{\text{int}}(\gamma_x)| < \infty \quad \text{for all } x \in \mathcal{T}. \tag{4.1}$$

We can then define a winding field $\{h_{\mathcal{T}}(x) : x \in \mathcal{T}\}$ of the tree simply as $h_{\mathcal{T}}(x) := W_{\text{int}}(\gamma_x)$.

The first generalisation step is to extend the definition of $h_{\mathcal{T}}$ even if (4.1) is not satisfied which follows from the following elementary lemma. If $x \notin \gamma_y$ and $y \notin \gamma_x$, notice that γ_x and γ_y eventually merge and γ_y merges either to the right or to the left of γ_x (since the tree is embedded in \mathbb{C} , this makes sense). Let γ_{xy} be the unique path connecting x and y in \mathcal{T} .

Lemma 4.1. *In the above setup, if γ_y merges with γ_x to its right, then*

$$h_{\mathcal{T}}(x) - h_{\mathcal{T}}(y) = W_{\text{int}}(\gamma_{xy}) + \pi.$$

If γ_y merges with γ_x to its left, then

$$h_{\mathcal{T}}(x) - h_{\mathcal{T}}(y) = W_{\text{int}}(\gamma_{xy}) - \pi.$$

If $y \in \gamma_x$ and y is not a corner (or $x \in \gamma(y)$ and x is not a corner), then

$$h_{\mathcal{T}}(x) - h_{\mathcal{T}}(y) = W_{\text{int}}(\gamma_{xy}).$$

Proof. Notice that the last assertion follows simply from additivity of intrinsic winding. Indeed, for example, if $y \in \gamma_x$, there is no discontinuity of intrinsic winding at y since y is not a corner.

For the rest, take $m \in \gamma_x \cap \gamma_y$. Notice that, by additivity of winding,

$$\begin{aligned} h_{\mathcal{T}}(x) - h_{\mathcal{T}}(y) &= W_{\text{int}}(\gamma_{xm}) - W_{\text{int}}(\gamma_{ym}) \\ &= W_{\text{int}}(\gamma_{xm}) + W_{\text{int}}(\gamma_{my}) \\ &= W_{\text{int}}(\gamma_{xy}) + \varepsilon\pi, \end{aligned}$$

where $\varepsilon = +1$ (or $\varepsilon = -1$) if γ_y merges with γ_x to its right (or left). This is clear since we need to do a half-turn to move from γ_{xm} to γ_{my} at m and the turn is clockwise if γ_y merges to the right of γ_x and anticlockwise otherwise. ■

We remark that the last assertion of the above lemma could be made to work even at corners by adding the appropriate angle; however, in what follows, we avoid using winding field at corners, so we do not need to introduce these additional difficulties. Finally, the formulas for $h(x) - h(y)$ described in Lemma 4.1 can be taken to be the definition of the winding field (or rather its gradient) for any tree embedded with piecewise smooth edges, even if (4.1) is not satisfied. This will be the typical situation for our setup.

Dual pairs and dimers. Take an infinite planar graph Γ , properly embedded in \mathbb{C} with smooth edges, and suppose that the embedding is locally finite in the sense that every vertex and face have finite degree, every face is a topological disc, and any bounded set intersects only finitely many edges. Let \mathcal{T} be a one ended spanning tree of Γ . (We emphasise that we are still considering a spanning tree for the moment and not yet a forest.) Since Γ is locally finite, \mathcal{T} admits a dual \mathcal{T}^\dagger which is also a one-ended spanning tree of the dual graph Γ^\dagger . Let G be the superposition graph (in \mathbb{C}), as introduced in Section 2.4 (with the obvious modifications to take into account that our graph is infinite and without creating any puncture). It will be useful to *augment the trees* \mathcal{T} and \mathcal{T}^\dagger as follows. Recall that, in Section 2.4, for each face f of G , we fixed a diagonal $d(f)$ which is just a smooth simple curve connecting the primal and dual vertex staying inside the face and a *midpoint* $m(f)$ which was just a point in $d(f)$. For each face, add to \mathcal{T} (resp., \mathcal{T}^\dagger) the portion of the diagonal $d(f)$ connecting the point $m(f)$ to the unique primal (dual) vertex touching $d(f)$. This way the primal and dual trees meet in each face f at the (smooth) point $m(f)$ on the diagonal $d(f)$ of that face. With a small abuse of notation, we will still denote by \mathcal{T} and \mathcal{T}^\dagger these augmented trees.




Figure 7. The definition of the height function in terms of winding.

Note that we can orient $\mathcal{T}, \mathcal{T}^\dagger$ towards their respective ends. This allows us to define two winding fields as above (one for \mathcal{T} and one for \mathcal{T}^\dagger). Having oriented \mathcal{T} and \mathcal{T}^\dagger , we can also apply the local operation of Figure 6 as described in Section 3.1 to obtain a dimer covering of G . We will first need to define a suitable reference flow on G , which will then allow us to speak of the height function associated to the dimer configuration and then show the relation between this dimer height function and the two above winding fields.

Definition 4.2 (Reference flow). Let w, b be two adjacent white and black vertices, and let f_l, f_r be the faces to the left and right of the oriented edge (bw) . Define

$$\omega_{\text{ref}}(wb) = \frac{1}{2\pi} (W_{\text{int}}((m(f_r), b) \cup (b, m(f_l))) + \pi), \tag{4.2}$$

where $(m(f_l), b)$ is the portion of $d(f_l)$ joining $m(f_l)$ and b , and similarly $(m(f_r), b)$. Define $\omega_{\text{ref}}(bw) = -\omega_{\text{ref}}(wb)$ (see Figure 7).

Lemma 4.3. ω_{ref} defined above is a valid reference flow. That is, the total mass sent out of any white vertex w , $\sum_{b \sim w} \omega_{\text{ref}}(wb)$, is equal to 1; and the total mass received by any black vertex b , $\sum_{w \sim b} \omega_{\text{ref}}(wb)$, is also 1.

Proof. Let w be a white vertex. Notice that, in $\sum_{b \sim w} \omega_{\text{ref}}(wb)$, the oriented diagonals form a clockwise loop whose total winding is -2π . Adding $+\pi$ for each of the surrounding black vertices and dividing by 2π , we see that the total flow out of w is indeed 1 as desired.

For a black vertex b , the argument is better explained by considering a picture (see Figure 8). Fatten the “star” formed by the half-diagonals incident to b into a star shaped domain. Then, notice that the total flow out of b is simply the limit of the total winding, divided by 2π , of the boundary of this domain (again in the clockwise orientation this time), as the domain thins into the “star”. Indeed, the $-\pi$ term in the




Figure 8. Proof of Lemma 4.3.

definition of ω_{ref} in (4.2) counts the half-turn as we move from the left side to the right side of a half-diagonal. Since the total winding of such a curve is -2π , it follows that the total flow out of b is -1 , as desired. ■

We are now ready to relate the three notions of height function defined by the pair of dual spanning trees $\mathcal{T}, \mathcal{T}^\dagger$. To do so, note that we can extend the definition of the winding field of both \mathcal{T} and \mathcal{T}^\dagger from Lemma 4.1 to the augmented trees.

Proposition 4.4. *In the above setup, let $h_{\mathcal{T}}$ and $h_{\mathcal{T}^\dagger}$ be the winding field of \mathcal{T} and \mathcal{T}^\dagger , respectively. Let h_{dim} be the height function corresponding to the dimer configuration obtained from $(\mathcal{T}, \mathcal{T}^\dagger)$ with reference flow ω_{ref} . For any two faces f, f' ,*

$$h_{\mathcal{T}}(m(f')) - h_{\mathcal{T}}(m(f)) = h_{\mathcal{T}^\dagger}(m(f')) - h_{\mathcal{T}^\dagger}(m(f)) = 2\pi(h_{\text{dim}}(f') - h_{\text{dim}}(f)).$$

Remark 4.5. The winding fields in \mathcal{T} and \mathcal{T}^\dagger now play a completely symmetric role.

Proof. For a face f , we define γ_f (respectively, γ_f^\dagger) to be the branch of the augmented tree \mathcal{T} (resp., \mathcal{T}^\dagger) starting from $m(f)$. Fix faces f and f' and assume without loss of generality that $\gamma_{f'}$ is to the right of γ_f when they merge. Note that this means that $\gamma_{f'}^\dagger$ is to the left of γ_f^\dagger . Let $\gamma_{ff'}$ and $\gamma_{f'f}^\dagger$ denote the paths in \mathcal{T} and \mathcal{T}^\dagger connecting, respectively, $m(f)$ to $m(f')$ and $m(f')$ to $m(f)$. With the above convention on the relative position of γ_f and $\gamma_{f'}$, the concatenation of $\gamma_{ff'}$ with $\gamma_{f'f}^\dagger$ has to be a simple clockwise loop and there is no jump of the winding at $m(f)$ and $m(f')$ because by assumption the midpoint $m(f)$ is a smooth point of $d(f)$, so

$$W_{\text{int}}(\gamma_{ff'}) + W_{\text{int}}(\gamma_{f'f}^\dagger) = -2\pi$$

The first equality easily follows by applying the definition of winding field from Lemma 4.1.

For the last equality, let f_l, f_r be adjacent faces. Let the common (oriented) edge be (bw) with b being the black and w the white vertex, and let f_r lie to its right. We




Figure 9. An example for the computation of the height along an arbitrary trajectory. The path P in Γ is the union of solid blue and dashed blue edges, and here, it is composed of $k = 3$ partitions. The solid blue denotes the segments in the tree \mathcal{T} with arrows giving the orientation of each edge. The union of solid red and solid blue edges is the modified path. The orange dotted paths denote the direction in which the primal tree goes off to ∞ and the green dotted path denotes the same for the dual tree. This determines the ε'_i 's and δ'_i 's. Here, $\varepsilon_1 = -1$, $\varepsilon_2 = 1$, $\varepsilon_3 = -1$, and $\delta_1 = 1$, $\delta_2 = -1$.

assume without loss of generality that b is a primal vertex. From the definition of ω_{ref} and recalling the sign convention of the flow defining the height function (Section 2.5),

$$\begin{aligned} 2\pi(h_{\text{dim}}(f_r) - h_{\text{dim}}(f_l)) &= \omega_{\text{ref}}(wb) - 2\pi 1_{(bw) \text{ occupied by dimer}} \\ &= W_{\text{int}}((f_r, b) \cup (b, f_l)) + \pi - 2\pi 1_{(bw) \text{ occupied by dimer}}. \end{aligned}$$

Note that γ_{f_r} and γ_{f_l} merge at b . Also, note from Temperley bijection that if (bw) is occupied by a dimer, then γ_b starts by using the (half) edge bw which implies that γ_{f_l} lies to the left of γ_{f_r} . Otherwise, γ_{f_l} lies to the left of γ_{f_r} . We conclude using the definition of winding field from Lemma 4.1 and the above equation. ■

From tree to forest. One can think of Proposition 4.4 as giving a way to compute the (gradient of the) winding field by following the path between two vertices in the tree. Before dealing with forests, we generalise the above by showing how to compute a height difference along an arbitrary path in Γ , which might use edges not in the tree.

We advise the reader to look at Figure 9 while reading the following definitions. Fix f, f' two faces of G , and let P be a self avoiding path starting at $m(f)$, going along $d(f)$ up to the vertex of Γ adjacent to f , then moving along edges of Γ up to f' , and finally following $d(f')$ up to $m(f')$. Partition this path into connected components belonging to \mathcal{T} (called segments from now on) separated by edges not belonging to \mathcal{T} , allowing for the segments to be reduced to a single vertex (so any two segments are joined by a single edge). We call these segments and edges, respectively,

(P_1, \dots, P_k) and (e_1, \dots, e_{k-1}) , and we let (x_i, y_i) be the starting and ending points of P_i and z_i be the white intersection vertex in e_i .

If $x_i \neq y_i$, let $\varepsilon_i = +1$ (resp., -1) if γ_{y_i} lies to the right (resp., left) of γ_{x_i} for $1 \leq i \leq k$. If $x_i = y_i$ (i.e., when P_i is a single vertex), then let $\varepsilon_i = +1$ (resp., -1) if γ_{x_i} starts to the left (resp., right) of P . Note that by definition the edges e_i are not in \mathcal{T} , so z_i are all in \mathcal{T}^\dagger . We let $\delta_i = +1$ (resp., -1) if $\gamma_{z_i}^\dagger$ starts to the left (resp., right) of P .

Lemma 4.6. *Under the above setup,*

$$\begin{aligned} h_{\mathcal{T}}(m(f)) - h_{\mathcal{T}}(m(f')) &= W_{\text{int}}(P) + \pi \left(\sum_{i=1}^k \varepsilon_i + \sum_{i=1}^{k-1} \delta_i \right) \\ &= 2\pi(h_{\text{dim}}(f) - h_{\text{dim}}(f')). \end{aligned}$$

Proof. Modify P as follows. Observe that the oriented edges $e_i = (y_i, x_{i+1})$ has two faces of G to their left. Call the one incident to y_i, f_i^- and the one incident to x_{i+1}, f_i^+ and m_i^-, m_i^+ the respective midpoints of the diagonals of these faces. By convention, also call $x_1 = m_0^+$ and $y_k = m_k^-$. We also join y_i to x_{i+1} using the diagonal segments of f_i^- and f_i^+ . Finally, we delete the edges e_i . This completes the modification of the path P , and note that this modification does not change the intrinsic winding, so we are still going to denote the modified path P (see Figure 9).

Note that, after this modification, connected components connecting diagonal midpoints staying in \mathcal{T} alternate with components connecting diagonal midpoints staying in \mathcal{T}^\dagger , so we can apply Lemma 4.1 and Proposition 4.4 for all such connected components. The terms ε_i and δ_i match the $\pm\pi$ terms in Lemma 4.1, so we obtain the result. ■

We now define the winding field given by a dual pair of spanning forests. We are in the same setup as above for the graphs Γ, Γ^\dagger , except that $\mathcal{T}, \mathcal{T}^\dagger$ are not necessarily spanning trees but spanning forests, and they are dual to each other.

Definition 4.7. Suppose that we are in the above setup and $(\mathcal{T}, \mathcal{T}^\dagger)$ is a dual pair of spanning forests where each component is oriented towards one of its ends. The winding field of $(\mathcal{T}, \mathcal{T}^\dagger)$ is a function defined on all diagonal midpoints which satisfies the first equality from Lemma 4.6. It is well defined and unique up to a global shift in \mathbb{R} .

Proof of Definition 4.7 being consistent. Once we show consistency, uniqueness up to a global shift by \mathbb{R} is clear. To that end, note that one can apply Temperley’s bijection to $(\mathcal{T}, \mathcal{T}^\dagger)$ to obtain a dimer configuration whose height function is always well defined up to a global shift in \mathbb{R} and which can be computed along any path. The second equality of Lemma 4.6 shows that the definition is consistent. ■

Remark 4.8. Note that there is no assumption about the number of ends of $(\mathcal{T}, \mathcal{T}^\dagger)$ in the above definition: it makes sense as long as an orientation is fixed towards a unique end in each component. Further, note that, in the above proof, we only need the pair $(\mathcal{T}, \mathcal{T}^\dagger)$ to be associated to a dimer configuration by Temperley’s bijection, so Definition 4.7 also extends to the case where Γ and $(\mathcal{T}, \mathcal{T}^\dagger)$ are obtained by lifting a graph and a Temperleyan pair (t, t^\dagger) (satisfying the assumptions in Section 2.4) from the manifold M' to its cover. Note however that we need to be a bit careful with the boundary vertices of Γ as they are somewhat special. In order not to deal with boundary components separately, we replace the boundary vertex by a cycle whose length is the same as the degree of the vertex, and with every vertex in the cycle having one neighbour in Γ (see Figure 10). It is not too hard to see that the neighbours can be chosen so that the embedding of Γ still satisfies the first three items of the definition of faithful embedding (see Section 2.4). Let us call these the boundary cycles of Γ . (The boundary cycles defined in Section 2.4 belonged to Γ^\dagger , so this should not cause confusion.) We now replace the boundary vertex in t by this cycle in a natural way. In this setup, every component of t, t^\dagger has a unique non-contractible cycle, with the boundary component of t corresponding to the boundary cycles we just defined. It is clear from the unique path lifting property that an oriented loop in t, t^\dagger corresponds to a collection of bi-infinite simple paths in the lift, each one obtained by going along the loop infinitely many times in clockwise and anti-clockwise direction. Hence, after the above surgery of the boundary vertex, t, t^\dagger lift to a forest $\mathcal{T}, \mathcal{T}^\dagger$ with every component having exactly two ends. Fixing an arbitrary orientation of the boundary cycles, we can also orient the lifts towards one of its ends. We can then readily apply Definition 4.7.

4.2. Winding of CRSF and height function

In this section, we give further details on the connections between winding and height functions which are specific to our setting, i.e., graphs and spanning forests obtained as lifts. More precisely, our main goal will be to provide for every pair of faces f, f' a canonical path $\gamma_{ff'}$ that makes the formula from Lemma 4.6 as simple as possible; see Theorem 4.10 for the final statement.

At this point, we need to spare a few words related to the removal of white vertices from the graph G to obtain a graph with a dimer cover. This operation can be interpreted as inserting certain discrete version of magnetic operators on the free field (e.g., in the sense of [19]). If we want to interpret the height function as winding, the height function would be additively multivalued where it picks up an additional $\pm 2\pi$ winding when it goes around a removed white vertex. Recall that we introduced a puncture corresponding to each face obtained by the removal of a white vertex (cf. Figure 2) and that the new manifold is called M' . According to the above discussion,




Figure 10. Splitting of the boundary vertex of Γ (red graph) into a cycle. The black graph is Γ^\dagger .




Figure 11. The conformal maps between an annulus and its universal cover. The dotted lines separate different copies of the fundamental domain. The blue and red curves show two loops in the surface and the associated spines in the cover.

it is natural to lift G' to the universal cover \tilde{M}' of M' and not the cover of M (for the rest of this section, when we speak of lift or cover, we will always refer to \tilde{M}') and call it \tilde{G}' . Note that punctures are mapped to the boundary of the disc and every non-outer face of G' is mapped to a quadrangle.

Fix a Temperleyan CRSF t of Γ' and its oriented dual t^\dagger , both augmented by joining the midpoints of diagonals as explained in Section 4.1. Each component of t or t^\dagger contains either a single cycle (which is oriented) or a boundary component. Using the surgery in Remark 4.8, each boundary component of t also has a unique non-contractible cycle. We call each component of the lift of a cycle a *spine*. Note that, by definition, each spine is a bi-infinite path oriented towards one of its ends.

We now state a useful lemma about the geometry of spines.

Lemma 4.9. *If M is hyperbolic, then any spine S is either a simple path in \mathbb{D} connecting two points in $\partial\mathbb{D}$ or a simple loop containing a unique point in $\partial\mathbb{D}$.*

If M is the torus, any spine S is a bi-infinite periodic path and all spines and dual spines have the same asymptotic direction.

The proof of the first part of Lemma 4.9 uses tools from the theory of Riemann surfaces, so we postpone it to appendix A in order to not disrupt the rest of the argument too much. For now, the reader may refer to Figure 11 for an illustration of the annulus case.

By Lemma 4.9, we see that $(\mathcal{T}, \mathcal{T}^\dagger)$ form a dual pair of spanning forest of $\tilde{\Gamma}'$ and its dual so Section 4.1 applies to them. We now turn to the definition of the terms involved in our main statement, Theorem 4.10.

Given two spines S and S' , we have a well-defined region $\Omega_{S,S'}$ between them which is bounded by S , S' and (in the hyperbolic case) two portions of $\partial\mathbb{D}$. (If S and S' are loops, then these portions of $\partial\mathbb{D}$ are understood as only prime ends associated with the same point.) We say that a spine or dual spine S'' separates S from S' if it connects these two portions of $\partial\mathbb{D}$. Since the graph \tilde{G}' only has accumulation points on $\partial\mathbb{D}$, it is easy to see that two spines can only be separated by finitely many others. Furthermore, two adjacent dual spines must be separated by a primal spine.

Now, pick two faces f, f' of \tilde{G}' , and let S, S' be the spines of the components of \mathcal{T} containing $m(f), m(f')$. Let $\Omega_{S,S'}$ be the component between them as above. Draw a simple curve connecting S and S' in $\Omega_{S,S'}$ and oriented from S to S' . For each primal spine $\sigma \subset \Omega_{S,S'}$ (hence not including S and S'), let ε_σ be the algebraic number of times σ is crossed by this curve (with the convention of counting $+1$ if σ is crossed from its left to its right by the curve and -1 otherwise). If σ is not crossed, define ε_σ to be 0. Define δ_σ similarly for primal spines. Notice that

$$\sum_{\sigma \subset \Omega_{S,S'}} \varepsilon_\sigma + \delta_\sigma \tag{4.3}$$

is a topological term which does not depend on the choice of the curve (where the sum is over all primal and dual spines contained in $\Omega_{S,S'}$). Also, its only dependence on f and f' is through the spines S and S' .

For any face f , recall that γ_f is the infinite oriented path in \mathcal{T} starting from $m(f)$ and going off to infinity along the unique outgoing oriented edges. Let ζ be the limit point of γ_f (which exists due to Lemma 4.9). Let ζ' be the limiting endpoint of S' which lies in the same connected component of $\partial\tilde{M} \cap \Omega_{S,S'}$ as ζ . Note that we are not defining ζ' to be the limit point of $\gamma_{f'}$ as it could be the case that ζ and ζ' might lie in different boundary components of $\partial\tilde{M} \cap \Omega_{S,S'}$ (depending on the orientation of S'). In case S, S' are loops through the same boundary point, we want ζ and ζ' to be in the same prime end of $\Omega_{S,S'}$.

- If the limit point of $\gamma_{f'}$ does not lie in the same component of $\partial\tilde{M} \cap \Omega_{S,S'}$ as ζ and f' is in $\Omega_{S,S'}$, then we set $\varepsilon_{S'} = -1$.
- In all other cases, define $\varepsilon_{S'} = +1$.

Finally, define $\varepsilon_S = +1$.

Let (ζ, ζ') be the arc joining ζ and ζ' (which could be a single point if $\zeta = \zeta'$) in $\partial\tilde{M}$ in the hyperbolic case. Let h_{dim} be the dimer configuration corresponding to the pair $(\mathcal{T}, \mathcal{T}^\dagger)$ with reference flow given by Definition 4.2. We can now state the following final result relating the dimer height function to the winding of trees.

Theorem 4.10. *In the hyperbolic case, let $\gamma := \gamma_{ff'}$ be the curve formed by concatenating γ_f , (ζ, ζ') and the path in \mathcal{T} joining ζ' to $m(f')$. Orient γ from $m(f)$ to $m(f')$. We have the following deterministic relation:*

$$\begin{aligned} h_{\mathcal{T}}(m(f')) - h_{\mathcal{T}}(m(f)) &= 2\pi(h_{\text{dim}}(f') - h_{\text{dim}}(f)) \\ &= W(\gamma, m(f)) + W(\gamma, m(f')) \\ &\quad + \pi \sum_{\sigma \in \bar{\Omega}_{S,S'}} (\varepsilon_\sigma + \delta_\sigma). \end{aligned} \tag{4.4}$$

Here, the sum $\sum_{\sigma \in \bar{\Omega}_{S,S'}} \varepsilon_\sigma + \delta_\sigma$ is as in (4.3) but also includes S and S' .

In the case of the torus, we have

$$\begin{aligned} h_{\mathcal{T}}(m(f')) - h_{\mathcal{T}}(m(f)) &= -2\pi(h_{\text{dim}}(f') - h_{\text{dim}}(f)) \\ &= W(\gamma_f, m(f)) - W(\gamma_{f'}, m(f)) + W(\gamma_f, m(f')) \\ &\quad - W(\gamma_{f'}, m(f')) + \pi \sum_{\sigma \in \bar{\Omega}_{S,S'}} (\varepsilon_\sigma + \delta_\sigma). \end{aligned} \tag{4.5}$$

Proof. Let us consider the hyperbolic case first. Take two vertices x, x' in S, S' , respectively. Observe that we can find a path joining x, x' which goes along the primal components and moves from one component to the other by “jumping” over dual spines (i.e., going along edges whose dual belongs to a dual spine). Also, one can make sure that the path is minimal, in the sense that every component between those of x, x' (i.e., components which separate S, S') is visited at most once. Call this path (x, x') . Now, letting $x \rightarrow \zeta$ and $x' \rightarrow \zeta'$, we can also ensure that the portion of $\Omega_{S,S'}$ bounded by (x, x') , (ζ, ζ') and the spines S, S' contains none of f or f' . Note that, for this path, the choice of the $\varepsilon_S, \varepsilon_{S'}$ exactly matches with the choice defined for Lemma 4.6. We refer the readers to the second figure of Figure 12 for the case when the limit point of $\gamma_{f'}$ does not lie in the same component of $\partial\tilde{M} \cap \Omega_{S,S'}$ and f' lies in $\Omega_{S,S'}$, in which case

$$\varepsilon_{S'} = -1.$$




Figure 12. An illustration of spines in the hyperbolic case. The spines S and S' are separated by a dual spine drawn in dashed blue. This spine also separates S'' and S' (left). The second and third figures from the left illustrate the choice of ε_S in Theorem 4.10: in the second, $\varepsilon_S = 1 = -\varepsilon_{S'}$. In the third figure (note that the orientation of S' is reversed), $\varepsilon_S = \varepsilon_{S'} = 1$.

Let $\tilde{\gamma}$ be the path obtained by concatenating $(m(f), x), (x, x'), (x', m(f'))$. From Lemma 4.6, it is clear that

$$2\pi(h_{\dim}(f') - h_{\dim}(f)) = W_{\text{int}}(\tilde{\gamma}) + \pi \sum_S (\varepsilon_S + \delta_S).$$

Indeed, notice that since the path is minimal, the ε_S, δ_S terms are defined so as to match with the definition of ε_i, δ_i in Lemma 4.6. Furthermore, $W_{\text{int}}(\tilde{\gamma}) = W_{\text{int}}(\gamma)$ because of the above choice of x, x' . We finish the proof using Lemma 2.1 and the fact that the topological winding of $\tilde{\gamma}$ and γ are equal.

For the torus case, the same proof applies by noticing that $W((x, x'), m(f)) + W((x, x'), m(f'))$ converges to 0 as x and x' go to infinity along S, S' in the same asymptotic direction. (The signs in (4.5) come from the fact that one goes along $\gamma_{f'}$ in the opposite direction.)

Finally, the fact that the term $\pi \sum_S (\varepsilon_S + \delta_S)$ converges follows simply from the fact that the number of non-contractible components is a.s. finite in the limiting CRSF (and hence so is the number of spines which separates S, S' into different components). ■

4.3. Mapping back to the manifold

In this subsection, we briefly explain how the notion of height function constructed above on the universal cover corresponds to a height one-form on the surface coming from the dimer cover. We also describe how the instanton component is a function of only the non-contractible loops of the Temperleyan CRSF.

Suppose that γ is a path in M' connecting two faces f, f' , and let $h(f, f')$ denote the height difference in M' calculated along γ . Now, imagine two copies of the lift

of γ to the universal cover \tilde{M}' and observe that, using our formula ((4.4) or (4.5)) for computing the height difference along the lifts in the universal cover, it is not clear how these quantities are related for different lifts and how these formulas map back to the surface M' . In our next lemma, we show that this difference among different copies can be explicitly calculated and is deterministic, so it does not contribute to the fluctuations of the height function, which is the quantity of interest in this article.

Recall from Sections 2.2 and 2.4 that the covering map p maps \tilde{M}'/F to \tilde{M}' , where F is a subgroup of \mathbb{Z}^2 in the torus case, and is a Fuchsian group (a discrete subgroup of the Möbius group) in the hyperbolic case. Thus, one can map one lift of a path in G' to another copy simply by using an element from F , which is either a translation in the torus case or a Möbius map in the hyperbolic case. In the following lemma, we show how this mapping affects the winding. We encourage the reader to recall the formulas in Section 2.3 before reading the lemma.

Lemma 4.11. *Suppose that we are in the hyperbolic case, so $\tilde{M}' = \mathbb{D}$. For any Möbius map ϕ from \mathbb{D} to itself, mapping \tilde{G}' to itself, for any two diagonal midpoints m, m' ,*

$$h_{\mathcal{T}}(\phi(m')) - h_{\mathcal{T}}(\phi(m)) = h_{\mathcal{T}}(m') - h_{\mathcal{T}}(m) + \arg_{\phi'(\mathbb{D})}(\phi'(m(f))) - \arg_{\phi'(\mathbb{D})}(\phi'(m(f'))).$$

In particular, $h_{\mathcal{T}} - \arg(\phi')$ can be mapped back as a one-form on the dual of G' , away from vertices corresponding to punctures. In the parabolic case, the same identity holds with no argument terms.

Proof. We focus on the hyperbolic case since the case of the torus is trivial. If ϕ is a conformal map mapping \tilde{G}' to itself, then it must also map \mathcal{T} to itself and in particular

$$\gamma_{\phi(f)\phi(f')} = \phi(\gamma_{ff'})$$

in the notations of Theorem 4.10. Clearly, the topological terms $\epsilon_{\sigma}, \delta_{\sigma}$ from Theorem 4.10 are invariant under ϕ :

$$\sum_{\sigma \in \Omega_{S,S'}} (\epsilon_{\sigma} + \delta_{\sigma}) = \sum_{\sigma \in \Omega_{\phi(S),\phi(S')}} (\epsilon_{\sigma} + \delta_{\sigma}).$$

For the term $W(\gamma, m(f)) + W(\gamma, m(f'))$, it transforms as an intrinsic winding under conformal maps, so by Lemma 2.2

$$\begin{aligned} &W(\phi(\gamma), \phi(m(f))) + W(\phi(\gamma), \phi(m(f'))) \\ &= W(\gamma, m(f)) + W(\gamma, m(f')) + \arg_{\phi'(\mathbb{D})}(\phi'(m(f))) - \arg_{\phi'(\mathbb{D})}(\phi'(m(f'))), \end{aligned}$$

which concludes the proof. ■

Remark 4.12. A consequence of Lemma 4.11 is the following. Fix any two faces f, g of G' and any path γ in the dual of G' starting at f and ending at g . Let $h(f, g)$ be the sum of the height one-form calculated along γ . Suppose that $\tilde{\gamma}$ and $\tilde{\gamma}'$ are two copies of the lifts of γ with endpoints \tilde{f}, \tilde{g} and \tilde{f}', \tilde{g}' , respectively. Recalling the definition of height function h_{\dim} from Sections 4.1 and 4.2, observe that

$$\begin{aligned} h_{\dim}(\tilde{f}) - h_{\dim}(\tilde{g}) - \mathbb{E}(h_{\dim}(\tilde{f}) - h_{\dim}(\tilde{g})) \\ = h_{\dim}(\tilde{f}') - h_{\dim}(\tilde{g}') - \mathbb{E}(h_{\dim}(\tilde{f}') - h_{\dim}(\tilde{g}')) \end{aligned}$$

since the arg terms cancel. In other words, the dimer height fluctuation does not depend on the choice of the copy of γ used in the cover and hence can be mapped back into M' without any ambiguity.

Let us now consider the instanton component. Intuitively, it should be completely determined by the homology class of the non-contractible loops of the Temperleyan forests (and in particular it should be determined by the spines). To that end, take a rooted loop λ , and consider the height difference along a particular choice of the lift of the loop, starting and ending at copies of the root; call this $h(\lambda)$. Lemma 4.11 tells us that we can compute $h(\lambda)$ unambiguously on any choice of the lift, but we should in fact be able to say more. In particular, computing $h(\lambda)$ using (4.4) and (4.5) a priori requires more precise geometric information involving windings. The following lemma shows that in fact the computation of $h(\lambda)$ does not involve any winding term and that the instanton component is at least a “reasonable” function of the non-contractible loops.

Take an ordered finite set of continuous simple loops which forms the basis of the first homology group of M' , all endowed with a fixed orientation. Let H be the finite set of numbers which denote the height change along these loops. It is well known (see Theorem 2.9) that \mathfrak{h} is completely determined by H (at the discrete level).

Lemma 4.13. *Let $(\mathbb{C}, \mathbb{C}^\dagger)$ be the set of oriented non-contractible loops of the primal and dual Temperleyan CRSF. Let \mathfrak{h} be the instanton component of the height 1-form of the dimer configuration corresponding to this CRSF pair given by the extended Temperley bijection. Then, \mathfrak{h} is a function of $(\mathbb{C}, \mathbb{C}^\dagger)$ only.*

Proof. We focus first on the hyperbolic case. Let λ be a loop from the basis of the homology group considered above; without loss of generality, we can assume that λ starts in a diagonal midpoint. Choose an arbitrary lift of λ and call its starting and ending points m and m' and the corresponding faces f and f' .

Let ϕ be the Mobius map sending m to m' and mapping \tilde{G}' to itself. By Theorem 4.10 and Lemma 2.1, $h_{\mathcal{T}}(m') - h_{\mathcal{T}}(m) = W_{\text{int}}(\gamma) + \pi \sum_{\sigma \in \bar{\Omega}_{S, S'}} (\epsilon_\sigma + \delta_\sigma)$ with γ obtained in this case by concatenating γ_f , the arc (ζ, ζ') and the time reverse of $\gamma_{f'}$.

(We assume for a brief moment that γ is smooth close to ζ and ζ' , and therefore piece-wise smooth everywhere.) By additivity of the intrinsic winding, we have

$$W_{\text{int}}(\gamma) = W_{\text{int}}(\gamma_f) + W_{\text{int}}((\zeta, \zeta')) - W_{\text{int}}(\gamma_{f'}),$$

but since $\gamma_{f'} = \phi(\gamma_f)$, by Lemma 2.2,

$$W_{\text{int}}(\gamma_f) - W_{\text{int}}(\gamma_{f'}) = -\arg(\phi'(\zeta)) + \arg(\phi'(m)).$$

Clearly, the above argument extends by taking an arbitrary regularisation of γ if it is not smooth near ζ and ζ' (since the right-hand side is independent of the regularisation procedure). Thus, we have proved that

$$\begin{aligned} h_{\mathcal{T}}(m') - h_{\mathcal{T}}(m) &= \arg(\phi'(m)) - \arg(\phi'(\zeta)) + W_{\text{int}}((\zeta, \zeta')) \\ &\quad + \pi \sum_{\sigma \in \bar{\Omega}} (\epsilon_{\sigma} + \delta_{\sigma}). \end{aligned} \tag{4.6}$$

The function ϕ only depends on λ and the choice of lift m for its starting point. Also, clearly, $\zeta, \zeta', \bar{\Omega}_{S, S'}$, and the $\epsilon_{\sigma}, \delta_{\sigma}$ only depend on the set of spines which concludes the proof of the hyperbolic case.

The proof for the parabolic case is identical and in fact even simpler since γ_f and $\gamma_{f'}$ are translate of each other and there is no term for the winding of (ζ, ζ') in Theorem 4.10. ■

We record now the following corollary regarding the weak limit of H as the mesh size goes to 0. For this corollary, we need to rely on results from the second article in the series [5] (and is not a deterministic result like the rest of this section) which will be reproduced later in the paper in Theorem 6.3. We write a superscript δ to account for the dependence in δ as in Section 2.4.

Corollary 4.14. $H^{\#\delta}$ described above converges in law as $\delta \rightarrow 0$ to some H which is measurable with respect to the limit of $(\mathbb{C}^{\#\delta}, (\mathbb{C}^{\dagger})^{\#\delta})$.

Proof. The term $\arg(\phi'(m)) - \arg(\phi'(\zeta)) + W_{\text{int}}((\zeta, \zeta'))$ is clearly a continuous function of $(\mathbb{C}, \mathbb{C}^{\dagger})$. Since $(\mathbb{C}^{\#\delta}, (\mathbb{C}^{\dagger})^{\#\delta})$ converges to an almost surely disjoint set of loops using Theorem 6.3, $\sum_{\sigma \in \bar{\Omega}} (\epsilon_{\sigma} + \delta_{\sigma})$ in the right-hand side of (4.6) also converges in law, thereby completing the proof. ■

5. Local coupling

In this section, we prove a local discrete coupling result which extends the ideas of [4] to the setup of Riemann surfaces. Roughly speaking, the goal of such a result is to

show that the local geometry in a small neighbourhood of a Temperleyan CRSF is given by that of a uniform spanning tree in the surface or, alternatively (and more usefully), in some reference planar domain. Moreover, locally, around a finite number of given points, the configurations can be coupled to independent such USTs.

Recall that the actual Temperleyan CRSFs cannot be completely sampled using the standard Wilson’s algorithm. However, due to Theorem 3.7, given the skeleton \varkappa emanating from either side of the punctures, the rest of the Temperleyan CRSF can be sampled from Wilson’s algorithm. The main result in this section is Proposition 5.10.

The argument follows the same line of arguments as in [4], so let us first recall the strategy there. This consists of two main steps. Consider k points v_1, \dots, v_k in $(\Gamma')^{\#\delta}$. In the first step, we choose cutsets at a small but macroscopic distance around each of the k points such that the cutsets separate the points from each other and from the rest of the graph. We reveal *all* the branches emanating from these cutsets. This leaves k unexplored subgraphs $\Gamma_i^{\#\delta}$, one around each point. In this step, the key point is to make sure that the $\Gamma_i^{\#\delta}$ are macroscopic (e.g., contain a ball of a radius roughly of the same order of magnitude as the distance to the cutsets). Clearly, the conditional law of the tree in each $\Gamma_i^{\#\delta}$ is that of a wired UST. Moreover, these wired USTs are also independent conditionally given the cutset exploration. (Of course, unconditionally there is still some dependence.) The second step is then to say that, in each $\Gamma_i^{\#\delta}$, one can couple a wired UST with a full plane one, which shows, among other things, that the unconditional distribution is close to being independent.

To adapt this strategy to Riemann surfaces, if the $\Gamma_i^{\#\delta}$ are sufficiently small so that it is simply connected, then the conditional law of the CRSF in each $\Gamma_i^{\#\delta}$ will also be that of a wired UST so that the second step can be used directly. For the first step however (cutset exploration), we will need to redo parts of the proof to take into account the possible loops in the CRSF.

We first state few useful lemmas from the second article [5, Lemmas 2.5 and 2.6].

Lemma 5.1 (Beurling-type estimate). *For all $r, \varepsilon > 0$, there exists $\eta > 0$ such that for any $\delta < \delta(\eta)$ and for any vertex $v \in \Gamma^{\#\delta}$ such that $\eta/2 < d_M(v, \partial M) < \eta$, the probability that a simple random walk exits*

$$B_M(v, r) := \{z \in M : d_M(z, v) < r\}$$

before hitting $\partial\Gamma^{\#\delta}$ is at most ε .

Lemma 5.2. *Let $K_0 \subset K'_0 \subset M$ be open connected sets, and let $K \subset K'_0, K'$ be compact sets which are the closures of K_0, K'_0 . Also, assume that K contains a loop which is non-contractible in M . Then, there exist $\delta_{K, K'} > 0$ and $\alpha_{K, K'} > 0$ depending only on K, K' such that for all $\delta < \delta_{K, K'}$ and all $v \in \Gamma^{\#\delta}$ such that $v \in K$, simple random walk starting from v has probability at least $\alpha_{K, K'}$ of forming a non-contractible loop before exiting K' .*

Next, we recall a lemma from [5] about the regularity of the skeleton branches away from the punctures. As mentioned above, the skeleton \mathfrak{s} cannot be directly sampled using Wilson’s algorithm; however, the following (technical) result from [5] guarantees that it will not be qualitatively different from a regular loop-erased random walk away from the punctures. Let $\bar{\lambda}_{\mathcal{M}}$ denote the law of \mathfrak{s} . Let

$$\eta_r := \{ \eta_r^{(i)} = (\eta_{r,1}^{(i)}, \eta_{r,2}^{(i)}) \}_{1 \leq i \leq k}$$

denote the portion of the branches of the skeleton until they exit a ball of radius r around the puncture, where r is some fixed small quantity. For each such path $\eta_{r,j}^{(i)}$ (with $1 \leq i \leq k, j = 1, 2$), run independent simple random walks from their tips, where each walk is conditioned to exit the ball of radius $r' > r$ before hitting $\eta_{r,j}^{(i)}$ for another small choice of r' . After hitting r' , the walks are allowed to run unconditionally until they either hit η_r or $(\partial\Gamma')^{\#\delta}$ or their loop erasure creates a non-contractible loop. Let $\tilde{\mu}_{r,r'}$ denote the joint law of the loop erasures of the resulting pair of walks.

Proposition 5.3 ([5, Proposition 6.5]). *For any event E which is measurable with respect to $\mathfrak{s} \setminus \eta_r$,*

$$\underline{\lambda}_{\mathcal{M}}(E) \leq C \left(\sqrt{\sum_{j=1}^{2k} \sup_{\eta_{r/2^j}} \tilde{\mu}_{\mathcal{M}_{r/2^j}, r/2^j-1}(E)} \right),$$

where the supremum is over all possible $\eta_{r/2^j}$.

Thanks to that proposition, we will be able to use Wilson’s algorithm to prove estimates about all the branches in Temperleyan CRSF, including the skeleton.

5.1. Cutset exploration

We now describe the construction more precisely. Fix k points v_1, \dots, v_k in $(\Gamma')^{\#\delta}$, and let N_{v_i} be small enough neighbourhoods around each v_i such that $\{N_{v_i}\}_{1 \leq i \leq k}$ do not intersect each other or the pre-specified boundary $\partial\Gamma^{\#\delta}$ of $\Gamma^{\#\delta}$ (if the boundary is non-empty) and for each $1 \leq i \leq k, p^{-1}(N_{v_i})$ is a disjoint union of sets in \tilde{M} such that p restricted to each component is a homeomorphism to its image. Also, for each i , fix one pre-image \tilde{v}_i of v_i , and let $\tilde{N}_{\tilde{v}_i}$ denote the component of \tilde{v}_i in $p^{-1}(N_{v_i})$. Let $R(\tilde{v}_i, \tilde{N}_{\tilde{v}_i})$ be the inradius seen from \tilde{v}_i in \tilde{M} with respect to the Euclidean metric, and call

$$r = \min_{1 \leq i \leq k} R(\tilde{v}_i; \tilde{N}_{\tilde{v}_i}). \tag{5.1}$$

(We point out that in our setup it is natural to define these quantities with respect to Euclidean geometry on \tilde{M} and project to M because our assumption and in particular

the uniform crossing condition are stated with respect to this geometry, whereas the intrinsic geometry of M does not really appear in our setup.)

We denote by $B_{\text{euc}}(\tilde{v}, r)$ the Euclidean ball of radius r around \tilde{v} , and for $r' > r$, let $A_{\text{euc}}(\tilde{v}, r, r') = B_{\text{euc}}(\tilde{v}, r') \setminus B_{\text{euc}}(\tilde{v}, r)$ be the Euclidean annulus of inradius r and outer radius r' . Let H_i be a set of vertices in $p(A_{\text{euc}}(\tilde{v}_i, r/2, r)) \cap (\Gamma')^{\#\delta}$ which disconnects v_i from ∂N_{v_i} . The *cutset exploration* is done in two steps.

- (a) First sample the skeleton according to the correct law (i.e., according to CRSF branches with the global conditioning that the branches divide the manifold into annuli).
- (b) Then, reveal all the branches from H_i , $1 \leq i \leq k$, using (unconditioned) Wilson’s algorithm, resulting in a subgraph $\mathcal{T}_{H_i}^{\#\delta}$, $1 \leq i \leq k$.

We say a vertex v_j has *cutset isolation radius* $6^{-k}r$ at scale r if

$$p(B_{\text{euc}}(\tilde{v}_j, 6^{-k}r)) \text{ does not contain any vertex from } \mathcal{T}_{H_i}^{\#\delta}.$$

Let us define J_{v_j} as the minimum value such that v_j has isolation radius $6^{-J_{v_j}}r$, and let $J = \max_j J_{v_j}$.

We want to show that J has exponential tail (which results in polynomial tail of the isolation radius). To this end, we rely on the following bound on the distance between a loop-erased walk and a point, which is a version of [4, Proposition 4.11].

Lemma 5.4. *Let γ be a loop-erased random walk starting from a vertex v until it hits $(\partial\Gamma')^{\#\delta}$ or a non-contractible loop is created. Let $u \neq v$ be another vertex, and let \tilde{u} be one of the images under p^{-1} of u . Let r be small enough such that $U := \bar{B}_{\text{euc}}(\tilde{u}, r)$ is contained in the pre-image of N_u which contains \tilde{u} . Then, there exist constants $\alpha, c > 0$ (depending only on the initial assumptions of the graph) such that, for all $\delta < \delta_U$ and all $n \in (0, \log_2(\frac{Cr\delta_0}{\delta}) - 1)$, where δ_U, δ_0 are as in the crossing condition (assumption (iv)),*

$$\mathbb{P}(\gamma \cap p(B_{\text{euc}}(\tilde{u}, 2^{-n}r)) \neq \emptyset) < c\alpha^n.$$

Proof. This is almost identical to [4, Lemma 4.11] except we now need to take into account the topology as well. To emphasise the differences with Lemma 4.11 in [4], we recall the strategy there. The idea is that if the loop-erased walk on the manifold comes inside $p(B_{\text{euc}}(\tilde{u}, 2^{-n}r))$, then some lift of the random walk must necessarily come within $B_{\text{euc}}(\tilde{u}, 2^{-n}r)$ – call this region (exponential) scale n . Furthermore, after the last such visit, this (lift of the) random walk must cross n many annuli without making a loop around \tilde{u} (which we called a “full turn”). This has a probability bounded by e^{-cn} .

In the current situation, however, the random walk might form a non-contractible loop before exiting U , and therefore its relevant lift described above does not necessarily have to cross n many annuli (see, e.g., Figure 13) after coming within Euclidean




Figure 13. A schematic representation (in solid line) of the lift of the loop erasure of the random walk on the torus until a non-contractible loop is formed. Call this path $\tilde{\gamma}$, and let $\gamma = p(\tilde{\gamma})$. The dashed square denotes the fundamental domain and the dashed paths denote some other lifts of γ . In this case, $\tilde{\gamma}$ stops inside the fundamental domain.

distance $2^{-n}r$ of \tilde{u} before we stop it. Thus, we cannot simply apply the argument of the previous paper. We overcome this using the following idea which we first explain informally. The random walk on the manifold may visit $p(B_{\text{euc}}(\tilde{u}, r))$ multiple times. Consider one such time when it enters $p(B_{\text{euc}}(\tilde{u}, r))$ (not necessarily the first time) and suppose its relevant lift \tilde{X} which also enters $B_{\text{euc}}(\tilde{u}, r)$ at that time. Let X' be the portion of the loop erasure of X on M at that time such that if the random walk hits X' later on, then a non-contractible loop in M is formed. Assume that X' comes inside $p(B_{\text{euc}}(\tilde{u}, 2^{-m}r))$. In order for the whole loop erasure γ to come inside $p(B_{\text{euc}}(\tilde{u}, 2^{-n}r))$, the following events must take place on one such visit. First, the lift \tilde{X} has to cross the first m scales without performing a full turn (otherwise, it would close a non-contractible loop on the manifold and stop); this has a probability bounded by e^{-cm} . Then, \tilde{X} needs to come within Euclidean distance $2^{-n}r$ of \tilde{u} , but we bound crudely the probability of this event by 1. Finally, \tilde{X} needs to come back out after the last visit to $B_{\text{euc}}(\tilde{u}, 2^{-n}r)$ in such a way that no full turn occurs between distances $2^{-n}r$ and $2^{-m}r$ (for the same reason as in the simply connected case). The probability of this event can be bounded by $e^{-c(n-m)}$. The intersection of these three events gives the right overall upper bound:

$$e^{-cm} e^{-c(n-m)} = e^{-cn}.$$

Let us fill in the details. Let C_i be the Euclidean circle of radius $r_i := 2^{-i}r$ around \tilde{u} for $i \geq 0$. We start a random walk X from v (to emphasise again: X is on the manifold M). Let $\{\tau_k\}$ be a sequence of stopping times defined inductively as follows. Set $\tau_0 = 0$. Having defined τ_k to be a time when the random walk X crosses or hits some circle $p(C_{i(k)})$, define τ_{k+1} to be the smallest time after τ_k when the random walk crosses or hits either $p(C_{i(k)-1})$ or $p(C_{i(k)+1})$; this defines τ_k by induction for every $k \geq 1$. If $i(k) = 0$, define τ_{k+1} to be the smallest time after τ_k when the random walk crosses or hits $p(C_{i(k)+1})$. Let k_{exit} be the smallest integer such that in

the interval $[\tau_{k_{\text{exit}}}, \tau_{k_{\text{exit}}+1}]$ either a non-contractible loop is created by the loop erasure of the random walk or the random walk hits the boundary $\partial\Gamma^{\#\delta}$. Let $\kappa_0, \kappa_1, \dots, \kappa_N$ be the set of indices $k < k_{\text{exit}} + 1$ when X_{τ_k} has crossed or hit $p(C_0)$. Also, note that the portions $X[\tau_{\kappa_0+1}, \tau_{\kappa_1}]$, $X[\tau_{\kappa_1+1}, \tau_{\kappa_2}]$, \dots are contained inside $p(C_0)$, while the portions $X[\tau_{\kappa_0}, \tau_{\kappa_0+1}]$, $X[\tau_{\kappa_1}, \tau_{\kappa_1+1}]$, \dots are contained outside $p(C_1)$. Our first claim is that N has exponential tail, i.e., $\exists \alpha_1 \in (0, 1)$, such that, for all $n \geq 1$,

$$\mathbb{P}(N > n) < \alpha_1^n. \tag{5.2}$$

Indeed, using Lemma 5.2, every time the walk hits $p(C_0)$, there is a positive probability independent of the past that the walk creates a non-contractible loop before returning to $p(C_1)$. Therefore, N has geometric tail. This proves (5.2).

Let \mathcal{S} be the set of points $\{X(\tau_k)\}_{k \geq 1}$. Note that, given \mathcal{S} , the pieces of random walk $\{X[\tau_k, \tau_{k+1}]\}$ are independent of each other. We call such pieces *elementary pieces* of the walk. If $i(k) \neq 0$, $X[\tau_k, \tau_{k+1}]$ is (given \mathcal{S}) a random walk starting from X_{τ_k} conditioned to exit the annulus $p(A_{\text{euc}}(\tilde{u}, r_{i(k)-1}, r_{i(k)+1}))$ at $X_{\tau_{k+1}}$. If $i(k) = 0$, $X[\tau_k, \tau_{k+1}]$ is a random walk which is conditioned to next hit $p(B_{\text{euc}}(\tilde{u}, r/2))$ at $X_{\tau_{k+1}}$. (Note that this property is lost if we solely condition on $\{X(\tau_k)\}_{k \leq \kappa_N}$ since the conditioning on N is complicated.) By [4, Lemma 4.7], conditionally on \mathcal{S} , each random walk piece $X[\tau_k, \tau_{k+1}]$ has a uniformly positive probability to do a *full turn* in the annulus $p(A_{\text{euc}}(\tilde{u}, r_{i(k)-1}, r_{i(k)}))$ for $\delta \leq \delta_U$ and given range of n (where δ_U comes from the uniform crossing assumption). Here, we define a full turn to be the event that the random walk crosses every curve joining the inner and outer boundary in the specified annulus. Indeed, although [4, Lemma 4.7] gives the uniform positive probability estimate for crossing in the Euclidean plane, the estimate is valid here by considering the relevant lift of $X[\tau_k, \tau_{k+1}]$ which is inside $\bar{B}_{\text{euc}}(\tilde{u}, r) = U$ and applying the uniform crossing estimate in the universal cover. Let us also point out that δ is chosen small enough so that the uniform crossing estimate is valid inside U .

We now define certain low probability events E_1, \dots, E_n such that one of them must take place if the loop erasure of X is to enter $B_n := p(B_{\text{euc}}(\tilde{u}, 2^{-n}r))$. Let us condition on \mathcal{S} . We say the event E_j occurs if the following hold.

- (i) The portion $X[\tau_{\kappa_{j-1}+1}, \tau_{\kappa_j}]$ intersects B_n . We then let λ_j be the smallest k in $[\tau_{\kappa_{j-1}+1}, \tau_{\kappa_j}]$ such that X_{τ_k} crosses (or hits) $p(C_n)$.
- (ii) Let X'_j be the portion of the loop erasure of $X[0, \tau_{\kappa_{j-1}+1}]$ such that if the walk $X[\tau_{\kappa_{j-1}+1}, \tau_{\kappa_j}]$ intersects X'_j , a non-contractible loop would be created. Let m be the maximum index such that X' intersects the circle $p(C_m)$. Then, for the event E_j to hold, we also require, in addition to the previous point, that, for any $\kappa_{j-1} + 1 \leq k \leq \lambda_j$, if $i(k) < m$, then the walk $X[\tau_k, \tau_{k+1}]$ does not perform a full turn in $p(A_{\text{euc}}(\tilde{u}, r_{i(k)-1}, r_{i(k)}))$. If no such m exists (e.g., if X'_j is empty), we do not require anything further.

- (iii) Let ℓ_j be the last k before κ_j such that X_{τ_k} crosses (or hits) $p(C_n)$. Let ℓ'_j be the first time after ℓ_j that the walk intersects $p(C_{m+1})$. Then, in addition to the previous two points, for E_j to hold we require that the walk $X[\tau_k, \tau_{k+1}]$ does not perform a full turn for any $\ell_j \leq k \leq \ell'_j$.

From the discussion above, it is clear that we have the following lemma.

Lemma 5.5. *We have*

$$\{\gamma \cap B_n \neq \emptyset\} \subseteq \left\{ N \leq n, \bigcup_{j=1}^n E_j \right\} \cup \{N > n\}.$$

Thus, all we need to show is that the event on the right-hand side of Lemma 5.5 has exponential tail bound. Now, we claim that

$$\mathbb{P}(E_j | \mathcal{S}) \leq e^{-c'n}, \quad \text{and so,} \quad \mathbb{P}(E_j) \leq e^{-c'n}.$$

This is justified using the uniform positive probability of the walk X performing a full turn, even conditionally given \mathcal{S} . Indeed, conditioned on \mathcal{S} , we have that the event (ii) in the definition of E_j has probability at most $e^{-c'm}$, and conditioned on event in (ii), the event in (iii) has probability at most $e^{-c''(n-m)}$ since, conditionally given \mathcal{S} , the random walk portions $X[\tau_k, \tau_{k+1}]$ are independent. Thus, the overall probability given \mathcal{S} is at most e^{-cn} with $c = c' \wedge c''$. All in all, using (5.2), Lemma 5.5, and a union bound, we obtain

$$\mathbb{P}(\{\gamma \cap B_n \neq \emptyset\}) \leq ne^{-c'n} + e^{-c'n} \leq ce^{-c'n}.$$

thereby concluding the proof. ■

Lemma 5.6. *Let \mathfrak{s} denote the skeleton. Let u be a vertex which is not one of the punctures, and let \tilde{u} be one of the images under p^{-1} of u . Let s be small enough such that $U := \bar{B}_{euc}(\tilde{u}, s)$ is contained in the pre-image of N_u which contains \tilde{u} . Then, there exist constants $\alpha, c > 0$ (depending only on the initial assumptions of the graph) such that for all $\delta < \delta_U$ and all $n \in (0, \log_2(\frac{Cs\delta_0}{\delta}) - 1)$, where δ_U, δ_0 are as in the crossing condition (assumption (iv)),*

$$\mathbb{P}(\mathfrak{s} \cap p(B_{euc}(\tilde{u}, 2^{-n}s)) \neq \emptyset) < c\alpha^n.$$

Proof. This is a corollary of Proposition 5.3 and Lemma 5.4. Consider the case $k = 1$ and recall the notations from Proposition 5.3. It is easy to see, using the Markov property of the independent random walks used to sample $\tilde{\mu}_{r/2,r}$ that after hitting r , the result of Lemma 5.4 kicks in, yielding the required exponential bound for $\tilde{\mu}_{r/2,r}(E)$. The same argument obviously applies to $\tilde{\mu}_{r/4,r/2}(E)$, yielding the desired statement.

The case $k > 1$ is analogous. We skip the details here. ■

We now state the result showing an exponential tail of J , which is a combination of Lemmas 5.4 and 5.6 and Schramm’s lemma, and is identical to the proof of [4, Lemma 4.20] (see also the proof of [4, Theorem 4.21]); as it is identical, we skip the proof here.

Lemma 5.7. *There exist constants $c, c' > 0$ such that the following holds. Let \tilde{D} be a compact set containing $B(\tilde{v}_i, r)$ for $1 \leq i \leq k$, where r is as in (5.1). Then, for all $\delta \in (0, \delta_{\tilde{D}})$ and for all $m \in (0, \log_6(\delta_0 r / \delta) - 1)$,*

$$\mathbb{P}(J > m) \leq ce^{-c'm}.$$

Finally, we state a lemma which says that with exponentially high probability, a branch of the CRSF, after entering an exponential scale t , does not backtrack to a smaller scale. Such a lemma for SLE curves can be found in [36] (see also [4, Lemma 3.4] in the simply connected case).

Lemma 5.8. *Fix u, U, δ_U as in Lemma 5.4. Suppose that γ is the loop erasure of a simple random walk in $\Gamma^{\# \delta}$ starting from vertex u until it hits the boundary or creates a non-contractible loop. Suppose that $\tilde{\gamma}$ is the lift of γ started from \tilde{u} (parametrised from \tilde{u} to its endpoint). There exist constants $c, c' > 0$ (depending only on the initial assumptions of the graph) such that for all $\delta < \delta_U$ and all $n \in (0, \log_2(\frac{Cr\delta_0}{\delta}) - 1)$*

$$\mathbb{P}(\tilde{\gamma} \text{ enters } B_{\text{euc}}(\tilde{u}, r2^{-n}) \text{ after exiting } B_{\text{euc}}(\tilde{u}, r2^{-n/2})) \leq ce^{-c'n}.$$

Proof. Let E be the event that $\tilde{\gamma}$ enters $B_{\text{euc}}(\tilde{u}, r2^{-n})$ after exiting $B_{\text{euc}}(\tilde{u}, r2^{-n/2})$. Let r be as in Lemma 5.4, and assume that n is even without loss of generality. The argument for this is very similar to Lemma 5.4 and in fact simpler, so we content ourselves with a sketch. Let r, C_i, B_i be as in the proof of Lemma 5.4. We look at the lift \tilde{X} of the simple random walk X starting at \tilde{u} , and we stop when either X hits the boundary or a non-contractible loop is created. Let τ_k be the set of stopping times defined as in the proof of Lemma 5.4 but for the lift \tilde{X} instead of X , and let \mathcal{S} be the set $\{\tilde{X}(\tau_k)\}_{k \geq 1}$. Observe that lift of the loop erasure of X is the loop erasure of \tilde{X} since erasing contractible loops commutes with lifting to the universal cover. If the loop erasure of \tilde{X} has to backtrack to scale n , the random walk has to necessarily enter scale n after leaving scale $n/2$. Let J be the largest j such that \tilde{X}_{τ_j} crosses or hits C_n after leaving $C_{n/2}$, and let I be the largest index i smaller than J when \tilde{X}_{τ_i} crosses or hits $C_{n/2}$. Conditioned on \mathcal{S} , if E occurs, then \tilde{X} enters B_n at least once after leaving $B_{n/2}$, and none of the elementary pieces of the walk between τ_I and τ_J can perform a full turn. But again, conditioned on \mathcal{S} , there is a uniformly positive probability to do a full turn for each elementary piece. Since there are at least $n/4$ such elementary pieces contained in $[\tau_I, \tau_J]$, we conclude the proof of the lemma applying the upper bound on the full-turn estimate on each elementary piece. ■

5.2. Full coupling

The results of the previous section covered the first step in the proof of the coupling. As we mentioned above, the second step is identical to the simply connected case, so we only recall the main statements.

We first recall the result which we will need from [4]. We use the notations and assumptions already in force. Let $\tilde{D} \subset \tilde{D}' \subset \tilde{M}'$ be two simply connected compact domains, and fix $\tilde{v} \in \tilde{D}^{\#\delta}$. Let $\mathcal{T}^{\tilde{D}'}, \mathcal{T}^{\tilde{D}}$ denote the wired UST, respectively, in $(\tilde{D}')^{\#\delta}$ and $\tilde{D}^{\#\delta}$. Let $r_{\tilde{v}}$ denote $R(\tilde{v}, \tilde{N}_{\tilde{v}})$ as in (5.1) (the largest Euclidean radius so that p is injective).

Lemma 5.9 ([4, Theorem 4.21]). *There exists $c, c' > 0$ such that the following holds. Fix $\tilde{v}, \tilde{D}, \tilde{D}'$ as above. There exists a coupling between $\mathcal{T}^{\tilde{D}}$ and $\mathcal{T}^{\tilde{D}'}$ and a random variable $R' > 0$ such that*

$$\mathcal{T}^{\tilde{D}'} \cap B_{\text{euc}}(\tilde{v}, R')^{\#\delta} = \mathcal{T}^{\tilde{D}} \cap B_{\text{euc}}(\tilde{v}, R')^{\#\delta}.$$

Furthermore, for all $\delta < \delta_{\tilde{D}'}$, and for all $m \in (0, \log_6(\delta_0 r_{\tilde{v}}/\delta) - 1)$, if we write $R' = 6^{-K_v} r_{\tilde{v}, \tilde{D}}$, where $r_{\tilde{v}, \tilde{D}}$ is the minimum of $r_{\tilde{v}}$ and the Euclidean distance between \tilde{v} and $\partial\tilde{D}$, then

$$\mathbb{P}(K_v \geq m) \leq ce^{-c'm}.$$

We now put together the cutset exploration with the above one-point coupling from the simply connected case as follows. Recall the notations from Section 5.1. Let $\mathcal{T}_{H_i}^{\#\delta}$ be the branches revealed in the cutset exploration around each v_i . Let

$$\mathcal{T}_H^{\#\delta} = \bigcup_{i=1}^k \mathcal{T}_{H_i}^{\#\delta}.$$

Let $\Gamma_i^{\#\delta}$ be the connected component of $(\Gamma')^{\#\delta} \setminus \mathcal{T}_H^{\#\delta}$ containing v_i . Observe that, given $\mathcal{T}_H^{\#\delta}$, the law of $\mathcal{T}^{\#\delta} \cap \Gamma_i^{\#\delta}$ is independent wired UST in $\Gamma_i^{\#\delta}$ (with the natural boundary) by the generalised Wilson algorithm. Applying the coupling of Lemma 5.9 in the lift of $\Gamma_i^{\#\delta}$ to \tilde{M}' and some fixed compact \tilde{D} in \tilde{M}' containing some lifts $\tilde{v}_1, \dots, \tilde{v}_k$ of points v_1, \dots, v_k , we obtain a coupling of the oriented CRSF in $\Gamma^{\#\delta}$ and independent wired USTs $(\mathcal{T}_i^{\tilde{D}})_{1 \leq i \leq k}$ in \tilde{D} . Furthermore, the r in (5.1) and $r_{\tilde{v}}$ in Lemma 5.9 differ by a constant multiplicative factor which depends only on the choice of lifts of the vertices. Furthermore, we may (and will) choose the $(\mathcal{T}_i^{\tilde{D}})_{1 \leq i \leq k}$ to be independent of $\mathcal{T}_H^{\#\delta}$. In fact, note that, for each $1 \leq i \leq k$, $\mathcal{T}_i^{\tilde{D}}$ may be chosen independent of the restriction of $\mathcal{T}^{\#\delta}$ to $\bigcup_{j \neq i} p(B(\tilde{v}_j, r/2))$.

Proposition 5.10. *The above coupling has the following properties: there exist random variables R_1, \dots, R_k such that*

$$\mathcal{T}^{\#\delta} \cap p(B(\tilde{v}_i, R_i)) = p(\mathcal{T}_i^{\tilde{D}} \cap B(\tilde{v}_i, R_i)). \tag{5.3}$$

Furthermore, if we write $R_i = 6^{-I_{v_i}} r_i$, where r_i is the minimum of r as in (5.1) and the distance between \tilde{v}_i and $\partial\tilde{D}$, then for all $\delta \leq \delta_{\tilde{D}}$, and for all $1 \leq i \leq k$, for all $n \in (0, \log_6(\delta_0 r_i / \delta) - 1)$,

$$\mathbb{P}(I_{v_i} \geq n) \leq ce^{-c'n}$$

for some constants $c, c' > 0$ (depending only on the initial assumptions on the graph). In particular, $\mathbb{P}(I_{v_i} \geq n) \leq ce^{-c'n} \vee \delta^{c'}$.

The set of non-contractible loops of $\mathcal{T}^{\#\delta}$ is measurable with respect to $\mathcal{T}_H^{\#\delta}$. In particular, $(\mathcal{T}_i^{\tilde{D}})_{1 \leq i \leq k}$ are also independent of the set of oriented non-contractible loops in the oriented CRSF.

Observe that when I_{v_i} is very big or r_i is very small, it is possible that the ball $B(v_i, R_i)$ is reduced to a point so the statement (5.3) is trivial (that is, (5.3) holds for a single vertex). This happens with a probability which is at most $\delta^{c'}$ for some c' .

Proof. This follows immediately from Lemma 5.7 and Lemma 5.9 once we observe that I_{v_i} is within $O(1)$ of the sum of J_{v_i} in Lemma 5.7 and K_{v_i} in Lemma 5.9, both of which have exponential tails.

The proof of the final assertion is a topological fact. Indeed, conditioned on $\mathcal{T}_H^{\#\delta}$, it is a deterministic fact that none of the oriented non-contractible loops pass through the portion of the CRSF in $\Gamma_i^{\#\delta}$ (which is a wired UST). Indeed, otherwise, we would have a path joining two points of the wired boundary of a wired UST which is impossible. Thus, the wired UST in each $\Gamma_i^{\#\delta}$ is conditionally independent of the non-contractible loops and hence so are $\mathcal{T}_i^{\tilde{D}}$ s. ■

6. Convergence of height function and forms

In this section, we precisely state our main result (that is Theorem 6.1) and then prove it. Recalling the sketch from Section 1.3, we see that what remains to be done is to go from the convergence of the Temperleyan CRSF to the convergence of its winding field using the coupling of Section 5. The global idea is the same as in [4], but some of the estimates on winding of LERW have to be redone. The primary difficulty is that spines are made of multiple copies of LERW, and hence, Wilson’s algorithm cannot be used to estimate the winding of all the copies. These are dealt with in Sections 6.2 and 6.3. The proof of the main result is finally concluded in Section 6.4.

6.1. Precise statement of the result

From now on, we work with the manifold M' . Recall that it is obtained by removing $2g + b - 2$ points from the interior of M and that the white vertices removed from

$G^{\#\delta}$ to obtain a Temperleyan graph in Theorem 3.7 will converge to these points. Denote by \tilde{M}' the universal cover of M' ; in the sequel, “lift” refers to lift on \tilde{M}' . Let $h^{\#\delta}$ be the height function defined as in Section 4 sampled using (1.1). Recall that $h^{\#\delta}$ is a function from the dual of $(\tilde{G}')^{\#\delta}$ to \mathbb{R} which is defined up to a global additive constant and $\bar{h}^{\#\delta} := h^{\#\delta} - \mathbb{E}(h^{\#\delta})$ is independent of the choice of the reference flow (see Remark 4.12). Recall the probability measures \mathbb{P}_{WilS} and \mathbb{P}_{Temp} from Section 3.3.

First, we extend $h^{\#\delta}$ to a function $h_{\text{ext}}^{\#\delta} : \tilde{M}' \rightarrow \mathbb{R}$ by defining $h_{\text{ext}}^{\#\delta}(x)$ to be the value of $h^{\#\delta}$ on the face containing x . Recall the graphs $(\Gamma^{\#\delta})'$, $(\Gamma^{\dagger, \#\delta})'$ from Section 2.4. Let $\mathcal{T}^{\#\delta}$ be the Temperleyan forest sampled using \mathbb{P}_{Temp} . Recall that we endowed M with a Riemannian metric d_M as in Section 2, and we equip M' with the same metric via the inclusion map. This induces a volume form $d\mu$ in \tilde{M}' .

Theorem 6.1. *Let $\tilde{f} : \tilde{M}' \rightarrow \mathbb{R}$ be a smooth compactly supported function with $\int_{\tilde{M}'} \tilde{f} d\mu = 0$. Suppose the assumptions in Section 2.4 hold. Then,*

$$\left(\int \tilde{f}(x) \bar{h}_{\text{ext}}^{\#\delta}(x) d\mu(x), \mathcal{T}^{\#\delta} \right)$$

converges jointly in law as $\delta \rightarrow 0$. The first coordinate also converges in the sense of all moments. Furthermore, the limit of the first coordinate is measurable with respect to the limit \mathcal{T} of $\mathcal{T}^{\#\delta}$, is universal (in the sense that it does not depend on the graph sequence $(G')^{\#\delta}$), and is conformally invariant.

To clarify, convergence in the sense of all moments means that, for all i , denoting $X^{\#\delta} := \int \tilde{f}(x) \bar{h}_{\text{ext}}^{\#\delta}(x) d\mu(x)$, $\mathbb{E}(|X^{\#\delta}|^i)$ converges as $\delta \rightarrow 0$. Notice also that since $\int \tilde{f} d\mu = 0$, the fact that $\bar{h}_{\text{ext}}^{\#\delta}(x)$ is defined only up to a global additive constant is irrelevant.

Let us explain briefly what is meant by conformal invariance in our setting. Suppose $|\chi| = k$ and $(M, x_1, x_2, \dots, x_k)$ and (N, y_1, \dots, y_k) are conformally equivalent in the sense that there is a map $\phi : M \rightarrow N$ so that $\phi(x_i) = y_i$ for all $1 \leq i \leq k$ and ϕ is a conformal bijection between the two Riemann surfaces. Note that if p_M and p_N are their respective covering maps (satisfying the properties in Section 2.4), then there exists a Möbius map ψ such that $p_N = p_M \circ \psi$. Let h_M (resp., h_N) denote the limit field obtained in the setup of Theorem 6.1 with the punctures being x_i and y_i for $1 \leq i \leq k$. Then, $h_N = h_M \circ \psi$ in the sense that, for any compactly supported test function f on the universal cover as in Theorem 6.1, $\int h_M f d\mu$ has the same law as $\int h_N (f \circ \psi^{-1}) |(\psi^{-1})'|^2 d\mu$.

Remark 6.2. Before we start giving the details, we remark that the upcoming proof of the theorem goes through if we can establish the convergence of the Temperleyan CRSF (in the Schramm sense) to a conformally invariant limit in which all curves are a.s. simple, together with a few additional ingredients. In addition to the assumptions

on the discretisation of the surface described in Section 2.4 and in particular the crossing assumption, these are the following: for each fixed vertex $v \in (\Gamma')^{\#\delta}$, let γ_v be the branch of the Temperleyan forest starting from v . Then, we require that γ_v satisfies Lemma 5.4 (the branch does not come close to a given point) as well as the moments on winding of Lemma 6.9 and Lemma 6.10. Note that the convergence to Brownian motion in item (iii) is not really needed here; instead, convergence of random walk to a process which is absolutely continuous with respect to Brownian motion would be sufficient. Stated this way, we note that Theorem 6.1 is novel even in the simply connected case. For instance, this result is applied in [3] and in the forthcoming [6].

For ease of reference, let us assemble in a single statement the convergence results from the second paper in the series that are needed in the following (see [5, Theorem 3.1, Corollary 3.14, and Proposition 6.1]).

Theorem 6.3. *Let M be a Riemann surface satisfying the assumptions of Section 2.1, which we equip with the distance d_M described in Section 2.1. Let $(\Gamma^{\#\delta})_{\delta>0}$ be a sequence of graphs with boundary $\partial\Gamma^{\#\delta}$ faithfully embedded in M satisfying the assumptions of Section 2.4. Then, the limit in law as $\delta \rightarrow 0$ of the Temperleyan CRSF sampled using \mathbb{P}_{Wils} or \mathbb{P}_{Temp} exists in the Schramm topology relative to d_M . Both limits are independent of the sequence $\Gamma^{\#\delta}$ subject to the assumptions in Section 2.4. These limits are also conformally invariant.*

Furthermore, let $K^{\#\delta}$ (resp., $(K^\dagger)^{\#\delta}$) denote the number of non-contractible loops in the CRSF (resp., its dual). Then, for any $q > 1$, there exists a constant $C_q > 0$ independent of δ such that

$$\mathbb{E}_{\text{Wils}}(q^{K^{\#\delta}}) \leq C_q,$$

where \mathbb{E}_{Wils} denotes the expectation under \mathbb{P}_{Wils} . The same bound holds also with $K^{\#\delta}$ replaced by $(K^\dagger)^{\#\delta}$.

Furthermore, let $(C_1^{\#\delta}, \dots, C_{K^{\#\delta}}^{\#\delta})$ be the set of non-contractible cycles of the Temperleyan CRSF,

$$(C_1^{\#\delta}, \dots, C_{K^{\#\delta}}^{\#\delta}) \xrightarrow[\delta \rightarrow 0]{(d)} (C_1, C_2, \dots, C_K),$$

where C_1, \dots, C_K are almost surely disjoint non-contractible simple loops in M .

6.2. Some a priori tail estimates on winding

The goal of this section is to obtain tail estimates on the winding of a branch of a Temperleyan forest. This will be achieved in Lemma 6.7 which is the main result of this section. Conceptually, the arguments are similar to [4, Section 4]. However, as in Section 5, there are additional difficulties linked with the fact that Wilson’s algorithm can stop because of the formation of a non-contractible loop. Furthermore,

we ultimately want to consider the winding of entire spines; yet only one copy of the lift of a loop is directly connected to a loop-erased random walk path. We first treat the part directly obtained by loop erasure in this section and defer estimates on the copies for the next section.

Throughout this section, we deal with a CRSF sampled from \mathbb{P}_{Wils} (cf. Section 3.3). First, we sample the skeleton (cf. Theorem 3.7) which we denote by \mathfrak{s} . Let $\tilde{\mathfrak{s}}$ be the lift of \mathfrak{s} to \tilde{M}' . Recall that \mathfrak{s} decomposes M' into a finite number of disjoint annuli and that, conditionally on \mathfrak{s} , in each of them we can sample the rest of the Temperleyan CRSF using Wilson’s algorithm.

We now work conditionally on \mathfrak{s} . Let γ_v be the branch of the CRSF starting from v (which can thus be sampled using Wilson’s algorithm). For any vertex $\tilde{v} \in p^{-1}(v)$, let $\tilde{\gamma}_{\tilde{v}}$ be the lift of γ_v starting from \tilde{v} and up until the time when γ_v closes a non-contractible loop or hits the boundary $\partial\Gamma^{\#\delta} \cup \mathfrak{s}$. Let τ_{nc} be the stopping time when the simple random walk generating $\tilde{\gamma}_{\tilde{v}}$ creates a non-contractible loop or hits the boundary. Recalling the definition of spines from Section 4, note that $\tilde{\gamma}_{\tilde{v}}$ will include a part of a spine whenever the random walk does not hit $\partial\Gamma^{\#\delta} \cup \mathfrak{s}$ when it stops (see Figure 13).

Let us orient $\tilde{\gamma}_{\tilde{v}}$ starting from \tilde{v} and going away from it in some continuous manner in $[0, 1]$. For $t > 0$, if $\tilde{\gamma}_{\tilde{v}}$ exits $B(\tilde{v}, e^{-t-1})$, let t_1 be the first time $\tilde{\gamma}_{\tilde{v}}$ exits $B(\tilde{v}, e^{-t-1})$, and let t_2 be the last time it exits $B(\tilde{v}, e^{-t})$. In this case, if $\tilde{\gamma}_{\tilde{v}}$ ends in $B(\tilde{v}, e^{-t})$ we set $t_2 = 1$.

We first state a simple deterministic lemma connecting the winding of curves avoiding each other.

Lemma 6.4. *Consider the annulus $A = \{z \in \mathbb{C} : r < |z - x| < R\}$, where $r > 0$ and $R \in (r, \infty]$, and let γ_0 be a simple curve in A connecting the outer and inner boundaries of A , assumed to be parametrised in $[0, 1]$. Then, for any simple curve γ in $A \setminus \gamma_0$, we have*

$$|W(\gamma, x)| \leq \sup_{0 \leq s_1 < s_2 \leq 1} |W(\gamma_0[s_1, s_2], x)| + 4\pi.$$

For $R = \infty$, the statement is interpreted as γ_0 connects the boundary of the $B(x, r)$ to ∞ .

Note that in the above lemma we do not need any assumption on the regularity of the curves beyond continuity.

Proof. Observe that any two points in an annulus can be joined by a smooth curve lying completely in a half-annulus containing those two points. Join each endpoint of γ to its nearest points (with its intrinsic metric inherited from the Euclidean plane) in γ_0 using smooth curves η_1, η_2 satisfying the above-mentioned property. Assume that

η_1 and η_2 intersect γ_0 at $\gamma_0[s_1]$ and $\gamma_0[s_2]$, respectively. Consider the loop L obtained by γ , η_1 , η_2 , and $\gamma_0[s_1, s_2]$. Note that L might be non-simple, but any loop which is created by the union of γ and the two geodesics lie in the annulus A slitted by γ_0 which is simply connected and does not contain x . Hence, all these loops contribute winding 0 around x . Furthermore, it is easy to see that erasing all these loops (say in a chronological manner starting from one of the tips of γ) creates a simple loop. Since the winding of L after erasing these loops is 0 or $\pm 2\pi$, the total winding of L around x is also either 0 or $\pm 2\pi$. Since by definition η_1 and η_2 lie in the smaller half-annulus containing the two points, its winding at most π . Using this observation, the proof of the lemma is complete. ■

Assume without loss of generality and to simplify notations that t is an integer. Let $r_{\tilde{v}}$ be such that $B(\tilde{v}, 10r_{\tilde{v}})$ does not intersect $\tilde{\mathfrak{s}}$ and p is injective in $B(\tilde{v}, r_{\tilde{v}})$. (This is a slight modification of the previous definition of $r_{\tilde{v}}$.) Let

$$e^{-t_0} = r_{\tilde{v}}.$$

We consider concentric circles C_j of radius $\{e^{-t_0-j}\}_{0 \leq j \leq t+2}$ (with B_j the ball inside it) around \tilde{v} . Take a random walk X in $(\Gamma')^{\#\delta}$ starting from \tilde{v} stopped when it hits $(\partial\Gamma')^{\#\delta} \cup \mathfrak{s}$. In case $(\partial\Gamma')^{\#\delta} \cup \mathfrak{s} = \emptyset$ (i.e., in the case of the torus), the random walk continues forever. Let \tilde{X} be the lift of this walk starting from \tilde{v} . Now, let $\{\tau_k\}_{k \geq 0}$ be the set of stopping times as described in the proof of Lemma 5.4 for the random walk \tilde{X} . That is, if we hit or cross the circle $C_{i(k)}$ at time τ_k , we wait until we hit or cross $C_{i(k) \pm 1}$ at time τ_{k+1} . If $i(k) = 0$ (or $t + 2$), we wait until we hit or cross C_1 (or C_{t+1}). Let B_i denote the disc inside C_i .

Let $\tau_{i_1}, \tau_{i_2}, \dots$ be the successive times in the sequence $(\tau_k)_{k \geq 1}$ when the walk hits C_0 . Note that the interval $[\tau_{i_j}, \tau_{i_{j+1}})$ is spent completely outside B_1 . We now claim that the random walk cannot wind too much outside B_1 .

Lemma 6.5. *There exist $c, c' > 0$ such that for all $\delta < \delta_{p(\bar{B}_0)}$, $n \geq 1$, $j \geq 1$, and $u \in B_1$ such that $\mathbb{P}(X_{\tau_{i_{j+1}}} = u) > 0$,*

$$\mathbb{P}\left(\sup_{\mathcal{Y} \subset \tilde{X}[\tau_{i_j}, \tau_{i_{j+1}}]} |W(\mathcal{Y}, \tilde{v})| > n \mid X_{\tau_{i_{j+1}}} = u\right) \leq ce^{-c'n},$$

$$\mathbb{P}\left(\sup_{\mathcal{Y} \subset \tilde{X}[\tau_{i_j}, \tau_{i_{j+1}}]} |W(\mathcal{Y}, \tilde{v})| > n \mid X_{\tau_{i_{j+1}}} \in \partial\Gamma^{\#\delta} \cup \mathfrak{s}\right) \leq ce^{-c'n}.$$

Here, the supremum is over all continuous paths obtained by erasing portions of $\tilde{X}[\tau_{i_j}, \tau_{i_{j+1}}]$.

Proof. The proof of Lemma 6.5 is very similar to [4, Lemma 4.8]. But it needs an input from Riemannian geometry to control the winding of the spines near the boundary of the universal cover. We postpone this proof to appendix A. ■

Lemma 6.5 takes care of the winding of the excursions outside B_1 . For excursions inside, most of the technical work was done in [4]. Let \tilde{Y}^j be the loop erasure of $\tilde{X}[0, j]$. For any j , we parametrise \tilde{Y}^j in some continuous way away from \tilde{v} . Fix $m \geq 1$, and let t_1 be the first time $\tilde{Y}^{\tau_{im}}$ exits $B(\tilde{v}, e^{-t-1})$ and t_2 the last time $\tilde{Y}^{\tau_{im}}$ exits $B(\tilde{v}, e^{-t})$.

Lemma 6.6. *There exist $C, c > 0$ such that, for all $\delta < \delta_p(\bar{B}_0)$, $t \in (t_0, \log(C\delta_0/\delta))$, and $n \geq 1, m \geq 1$,*

$$\mathbb{P}\left(\sup_{\mathcal{Y} \subseteq \tilde{Y}^{\tau_{im}}[t_1, t_2] \cap B(\tilde{v}, e^{-t-1})^c} |W(\mathcal{Y}, \tilde{v})| > n \mid \tau_{im} < \infty\right) < Ce^{-cn}.$$

We emphasise that, in the above lemma, m is non-random as it is going to be important in what follows.

Proof. The proof of this lemma is identical to that of [4, Lemma 4.13], and the only difference is in the setup, which we point out. In [4], we were working with a simply connected domain and we waited until the random walk exited it. The argument proceeds by conditioning on the positions $\mathcal{S} := \{\tilde{X}_{\tau_k}\}_{k \geq 1}$ and arguing that, in each interval of walk between successive points in \mathcal{S} , the winding of any continuous subpath has exponential tail. Then, we proved that we only need to look at a random number of intervals which itself has exponential tail. In the present setup, we can condition on \mathcal{S} so that $\tau_{im} < \infty$ is satisfied. The exponential tail of winding inside any inner annulus follows from [4, Lemma 4.7] (this is exactly the same as the simply connected case), and for the outer annulus, we use Lemma 6.5. The number of relevant intervals to consider also has exponential tail following verbatim the proof of [4, Lemma 4.15]. ■

With these lemmas, we can now state and prove the main result of this section, which controls the (topological) winding in an annulus.

Lemma 6.7. *There exist constants $C, C', c > 0$ so that for all $n \geq 1$, for all \tilde{v} , for all $\delta < \delta_{\bar{B}}(\tilde{v}, r_{\tilde{v}})$, and for all $0 < t < \log(C'r_{\tilde{v}}\delta_0/\delta)$,*

$$\mathbb{P}\left(\max_{\mathcal{Y} \subseteq \tilde{\gamma}_{\tilde{v}}[t_1, t_2] \cap B(\tilde{v}, r_{\tilde{v}}e^{-t-1})^c} |W(\mathcal{Y}, \tilde{v})| > n\right) < Ce^{-cn^{1/3}}, \tag{6.1}$$

where the maximum is taken over all continuous segments from $\tilde{\gamma}_{\tilde{v}}[t_1, t_2]$. Also,

$$\mathbb{P}\left(\max_{\mathcal{Y} \subseteq \tilde{\gamma}_{\tilde{v}} \cap B(\tilde{v}, r_{\tilde{v}}e^{-1})^c} |W(\mathcal{Y}, \tilde{v})| > n\right) < Ce^{-cn}. \tag{6.2}$$

A few remarks are in order. Firstly, the stretched exponential tail is an artefact of the proof, and we believe that an exponential tail bound could be proved with more care if necessary. Secondly, note that the intersection in the argument of the max above

is used so that we do not need to look into what happens very close to \tilde{v} . This is a technicality which simplifies the proof, but later it is not going to matter. In the end, we can decompose the whole path $\tilde{\gamma}_{\tilde{v}}$ into a disjoint union over t of $[t'_1, t'_2]$, where t'_1 is the last exit of $B(\tilde{v}, e^{-t-1})$ and t'_2 is the last exit of $B(\tilde{v}, e^{-t})$ which will accomplish the desired moment bound of truncated winding by exploiting the moment bound for each of these segments. We also emphasise that, although the non-contractible loop in the component containing v could be a branch from the skeleton, $\tilde{\gamma}_{\tilde{v}}$ itself does not contain any portion of the skeleton. This is simply because we have sampled the skeleton first before defining γ_v .

Proof of Lemma 6.7. The main idea is to use a union bound for the winding of each excursion between annuli. One difficulty with working with $\tilde{\gamma}_{\tilde{v}}$ is that if we condition on where or when a non-contractible loop is formed then we break the independence between the pieces of random walk. Indeed, a walk conditioned on not forming a non-contractible loop will avoid certain portions of its previous trajectory, which *a priori* might bias it to have a lot of winding. Recall that, at the point when a non-contractible loop is formed, the walk could be very close to the starting point. Thus, the idea is to run the random walk until the boundary $(\partial\Gamma')^{\#\delta}$ is hit (which could be potentially longer than dictated by Wilson’s algorithm) and control the winding of its loop erasure uniformly over all scales using Lemmas 6.5 and 6.6. After this, we need to separately compare the winding made by the last excursion into the ball to the previous estimate.

Recall that t_0 satisfies $e^{-t_0} = \tilde{r}_{\tilde{v}}$. Note that (6.2) follows easily from Lemma 6.5. Indeed, let us call τ_{nc} the first time where a non-contractible loop is created and, let τ_{∂} be the first hitting of the boundary. Fix $N = \max\{j : \tau_{i_j} < \tau_{nc} \wedge \tau_{\partial}\}$. Recall that N has exponential tail since every time we hit C_0 , conditioned on what happened before, we have a positive probability to create a non-contractible loop or hit the boundary before hitting C_1 , by Lemma 5.2. Thus, we can work on the event $N \leq n$ at a cost which is exponentially small in n . Now, Lemma 6.5 entails that on each of the n pieces, the winding has uniform exponential tail. This completes the proof of (6.2). We emphasise here that, to bound portions of the loop erasure outside radius e^{-1} , we can run the random walk until it hits the boundary $(\partial\Gamma')^{\#\delta}$ and not worry about the case when the walk stops deep inside the ball, as the loop erasure outside radius e^{-1} is already taken care of.

Now, we turn to (6.1). As before, let $\mathcal{S} := \{\tilde{X}_{\tau_k}\}_{k \geq 1}$. The idea is to compare the winding of $\tilde{\gamma}_{\tilde{v}}$ to the winding of $\tilde{Y}^{\tau_{i_m}}$ because there the conditioning on \mathcal{S} preserves the independence and Lemma 6.6 controls the winding. The main work will be to do a case-by-case analysis of what can be erased and created between τ_{i_m} and the creation of a non-contractible loop.

Let N be as above, and note that using the same idea of exponential tail of N and exponential tail of winding up to a fixed number of hits of the outermost circle

(Lemma 6.6), we get

$$\mathbb{P}\left(\sup_{\mathcal{Y} \subseteq \tilde{Y}^{\tau_{i_N}} [t_1, t_2] \cap B(\tilde{v}, e^{-t-1})^c} |W(\mathcal{Y}, \tilde{v})| > n\right) < Ce^{-c'n}. \tag{6.3}$$

Let λ_1 and λ_2 be, respectively, the first time $\tilde{Y}^{\tau_{i_N}}$ exits $B(\tilde{v}, e^{-t-1})$ and the last time $\tilde{Y}^{\tau_{i_N}}$ exits $B(\tilde{v}, e^{-t})$. To ease notations, from now on, all maximums in this proof come with the additional condition that the paths stay outside $B(\tilde{v}, e^{-t-1})$ without writing it explicitly. We will also assume without loss of generality that the parametrisations of $\tilde{\gamma}_{\tilde{v}}$ and $\tilde{Y}^{\tau_{i_N}}$ are identical up to the first point where their traces differ.

We now consider two possible cases. If a non-contractible cycle is created or the boundary is hit in the interval $[\tau_{i_N}, \tau_{(i_N+1)}]$, then, since $\tilde{X}[\tau_{i_N}, \tau_{(i_N+1)}]$ does not intersect $B(\tilde{v}, e^{-t_0-1})$, only pieces of $\tilde{Y}^{\tau_{i_N}}[0, \lambda_2]$ can be erased and nothing can be added in the time interval $[\tau_{i_N}, \tau_{(i_N+1)}]$ to get $\tilde{\gamma}_{\tilde{v}}[0, t_2]$. In particular, we see that in this case

$$\max_{t_1 < s_1 < s_2 < t_2} |W(\tilde{\gamma}_{\tilde{v}}(s_1, s_2), \tilde{v})| \leq \max_{\lambda_1 < s_1 < s_2 < \lambda_2} |W(\tilde{Y}^{\tau_{i_N}}(s_1, s_2), \tilde{v})|,$$

and we are done using (6.3).

The only other possibility is that a non-contractible cycle is created between $\tau_{(i_N+1)}$ and $\tau_{i_{(N+1)}}$ (otherwise that would contradict the maximality of N). Since p is injective in $B(\tilde{v}, e^{-t_0})$, this occurs if and only if the walk hits a copy inside $B(\tilde{v}, e^{-t_0})$ of a portion of the walk that is further away from \tilde{v} , as illustrated schematically in Figure 14. From now on, we assume that we are on this event.

First, we make a topological observation. Let β be the first exit time of $B(\tilde{v}, e^{-t_0})$ by $\tilde{Y}^{\tau_{i_N}}$. We claim that $\tilde{X}[\tau_{i_{N+1}}, \tau_{nc}]$ cannot hit $\tilde{Y}^{\tau_{i_N}}[0, \beta]$. Indeed, suppose by contradiction that it does so at some time $T \in [\tau_{i_{N+1}}, \tau_{nc}]$. Then, after erasure, we are left with a path completely contained in $B(\tilde{v}, e^{-t_0})$, where p is injective. No non-contractible loop can then be created before $\tau_{i_{N+1}}$, which contradicts the maximality of N as explained above.

Let k_{nc} be the index during which the non-contractible loop happens, i.e., the unique index k such that $\tau_{nc} \in [\tau_k, \tau_{k+1}]$. By the topological claim in the previous paragraph, no full turn can occur in any of the intervals $[\tau_k, \tau_{k+1}]$ for $i_N \leq k < k_{nc}$. However, by [4, Corollary 4.5], a full turn can occur in any interval $[\tau_k, \tau_{k+1}]$ independently with uniformly positive probability given \mathcal{S} . Hence, there exist constants c, C depending only on the crossing estimate such that

$$\mathbb{P}(k_{nc} - i_N > n) \leq \mathbb{P}(k_{nc} - i_N > n, N \leq n) + \mathbb{P}(N > n) \leq Ce^{-cn}.$$

Combining the above with [4, Lemma 4.7] (which bounds the winding of the random walk during an interval of the type $[\tau_i, \tau_{i+1}]$), we obtain the following stretched exponential tail bound:

$$\mathbb{P}\left(\max_{\mathcal{Y} \subseteq \tilde{X}[\tau_{i_N}, \tau_{nc}]} |W(\mathcal{Y}, \tilde{v})| > n\right) \leq Ce^{-c\sqrt{n}} \tag{6.4}$$




Figure 14. The concatenation of the red and blue curves is $\tilde{Y}^{\tau_{i_N}}$ and the green part is $\tilde{X}[\tau_{i_N}, \tau_{nc}]$. The dotted path denotes a copy of $\tilde{Y}^{\tau_{i_N}}$ so overall the event represented is one where a non-contractible loop is formed inside $B(\tilde{v}, e^{-t_0})$.

for some $C, c > 0$ depending only on the crossing estimate and where, as before, \mathcal{Y} is any continuous portion obtained from $X[\tau_{i_N}, \tau_{nc}]$ which preserves the order of the random walk path. In particular, this gives a good control of the winding of the piece of $\tilde{\gamma}_{\tilde{v}}$ added to $Y^{\tau_{i_N}}$.

However, we are still not done as illustrated in Figure 14: the times t_1, t_2 (which were first entry, last exit times for $\tilde{\gamma}_{\tilde{v}}$) may be different from λ_1, λ_2 (which were the first entry, last exit of $\tilde{Y}^{\tau_{i_N}}$), so we need additional arguments. Let

$$\lambda_E := \inf\{\lambda : \exists t \in [\tau_{i_N}, \tau_{nc}], \tilde{Y}^{\tau_{i_N}}(\lambda) = \tilde{X}(t)\}$$

be the last time in $\tilde{Y}^{\tau_{i_N}}$ which is not erased. Note that with this notation we showed above that $\lambda_E \geq \beta$. This implies immediately that

$$t_1 = \lambda_1.$$

Now, we need to consider several cases depending on where λ_E is in relation to λ_2 .

- If $\lambda_E \in [\lambda_1, \lambda_2]$, then $\tilde{\gamma}_{\tilde{v}}[t_1, t_2]$ can be decomposed as the union of a piece of $\tilde{Y}^{\tau_{i_N}}[\lambda_1, \lambda_2]$ and some erasure of $X[\tau_{i_N}, \tau_{nc}]$, so on that event

$$\max_{t_1 < s_1 < s_2 < t_2} |W(\tilde{\gamma}_{\tilde{v}}(s_1, s_2), \tilde{v})| \leq \max_{\lambda_1 < s_1 < s_2 < \lambda_2} |W(\tilde{Y}^{\tau_{i_N}}(s_1, s_2), \tilde{v})| + \Delta,$$

where Δ is a variable with stretched exponential tail using (6.4).

- If $\lambda_E \geq \lambda_2$ and $\tilde{\gamma}_{\tilde{v}}(\lambda_E, t_2)$ (the loop erasure after λ_E) does not enter $B(\tilde{v}, e^{-t})$, then $\tilde{\gamma}_{\tilde{v}}[t_1, t_2] = \tilde{Y}^{\tau_{i_N}}[\lambda_1, \lambda_2]$ and there is nothing more to prove.

- If $\lambda_E \geq \lambda_2$ and $\tilde{\gamma}_{\tilde{v}}(\lambda_E, t_2]$ enters $B(\tilde{v}, e^{-t})$ (this is the case pictured in Figure 14), then we decompose $\tilde{\gamma}_{\tilde{v}}$ as the union of $\tilde{Y}^{\tau_{i_N}}[\lambda_1, \lambda_2]$, $\tilde{Y}^{\tau_{i_N}}[\lambda_2, \lambda_E]$, and $\tilde{\gamma}_{\tilde{v}}[\lambda_E, t_2]$. The winding of any continuous portion of the first part has exponential tail by (6.3). The winding of any continuous portion of the last part has stretched exponential tail since it is bounded by Δ as in the first case. So, we only need to take care of the middle part.

To that end, we decompose $\tilde{Y}^{\tau_{i_N}}[\lambda_2, \lambda_E]$ into excursions inside and outside B_0 as follows:

$$\begin{aligned} &\tilde{Y}^{\tau_{i_N}}[\lambda_2, \lambda_E] \\ &= \tilde{Y}^{\tau_{i_N}}[\lambda_2, \beta_0] \cup \tilde{Y}^{\tau_{i_N}}[\beta_0, \beta_1] \cup \tilde{Y}^{\tau_{i_N}}[\beta_1, \beta_2] \cup \dots \cup \tilde{Y}^{\tau_{i_N}}[\beta_{k_0}, \lambda_E], \end{aligned} \tag{6.5}$$

where $\beta_0 = \beta$, and for $k \geq 1$, β_{2k} is the first exit of $B(\tilde{v}, e^{-t_0})$ after β_{2k-1} and β_{2k-1} is the first entry into $B(\tilde{v}, e^{-t_0-1})$ after β_{2k-2} . Note that any portion of $\tilde{Y}^{\tau_{i_N}}[\beta_{2k}, \beta_{2k+1}]$ is outside $B(\tilde{v}, e^{-t_0-1})$, so its winding has exponential tail using Lemma 6.5. Note also that $\tilde{Y}^{\tau_{i_N}}[\beta_{2k+1}, \beta_{2k+2}]$ lies inside the annulus bounded between C_t and C_0 and never intersects the curve $\tilde{X}[\tau_{i_N}, \tau_{nc}]$ by maximality of λ_E . Let Y' be the loop erasure of the portion of \tilde{X} from τ_{i_N} until its first hit of C_t . Hence, using Lemma 6.4,

$$\max_{\beta_{2k} < s_1 < s_2 < \beta_{2k+1}} |W(\tilde{Y}^{\tau_{i_N}}(s_1, s_2), \tilde{v})| \leq \max_{s_1 < s_2} |W(Y'[s_1, s_2], \tilde{v})| + 4\pi.$$

Note that the right-hand side has stretched exponential tail by (6.4). Thus, each term in the decomposition (6.5) has stretched exponential tail, and clearly, the number of terms is bounded by N which itself has exponential tail. Combining these two, we can conclude that the absolute value of winding of any continuous portion of $\tilde{Y}^{\tau_{i_N}}[\lambda_2, \lambda_E]$ has stretched exponential tail (with exponent $1/3$), thereby completing the proof of this case.

This concludes the proof of Lemma 6.7. ■

6.3. From partial path to the full spine

As shown in Theorem 4.10, the path $\tilde{\gamma}_{\tilde{v}}$ defined above is only a part of what is needed to compute the height function. Recall the notion of spine from Section 4.2 and that when we follow the outgoing edges in the Temperleyan forest from any given \tilde{v} , we always end up in a unique spine (since boundary loops have been added around every hole).

Recall from Theorem 4.10 that we are interested in the winding of the path starting from \tilde{v} and then moving along the spine to infinity or the boundary of the disc. Observe that the initial portion of this path is $\tilde{\gamma}_{\tilde{v}}$ (Figure 13), and then, it moves along

copies of the non-contractible loop in the component of \tilde{v} in the Temperleyan CRSF. In this section, we will call this path $p_{\tilde{v}}$ (i.e., $\tilde{\gamma}_{\tilde{v}}$ followed by a semi-infinite piece of the spine). Observe that this notation is in contrast to the notation used in Theorem 4.10, where the path $p_{\tilde{v}}$ was called γ_f , whereas here we emphasise that $\tilde{\gamma}_{\tilde{v}}$ denotes what is sampled from Wilson’s algorithm given the skeleton. When we apply Wilson’s algorithm to sample $\tilde{\gamma}_{\tilde{v}}$, given the skeleton, we may form a new non-contractible loop, in which case we have discovered a portion of a spine (the unique spine attached to \tilde{v}). The winding of this portion is controlled by Lemma 6.7. However, we also need to control the winding of the rest of $p_{\tilde{v}}$. The purpose of this section is precisely to achieve this (done in Lemma 6.8). We also need to control the winding of the other semi-infinite piece of the spine attached to \tilde{v} , which is done in Lemma 6.9. Finally, if we want to control the height gap between \tilde{u} and \tilde{v} , we need to control the winding of $p_{\tilde{u}}$ around \tilde{v} as well, which is done in Lemma 6.10.

Let us parametrise $p_{\tilde{v}}$ in $[0, \infty)$ in any continuous manner away from \tilde{v} . Let t_0 be as in Section 6.2, i.e., $e^{-t_0} = \tilde{r}_{\tilde{v}}$. For $t \geq t_0$, let t_1 be the first exit time of $p_{\tilde{v}}$ from $B(\tilde{v}, e^{-t-1})$, and let t_2 be the last exit time of the same from $B(\tilde{v}, e^{-t})$. Note again that here it is possible that the unique spine attached to \tilde{v} could be the same as the spine corresponding to a skeleton branch.

Lemma 6.8. *For all $k \geq 1$, there exists a constant $m > 0$ so that, for all $\delta < \delta_{\tilde{B}(\tilde{v}, e^{-t_0})}$ and for all $t_0 \leq t < \log(C'\delta_0/\delta)$,*

$$\mathbb{E}\left(\max_{\mathcal{Y} \subseteq p_{\tilde{v}}[t_1, t_2] \cap B(\tilde{v}, e^{-t-1})^c} |W(\mathcal{Y}, \tilde{v})|^k\right) \leq m,$$

where the supremum is taken over all continuous segments. Also,

$$\mathbb{E}\left(\max_{\mathcal{Y} \subseteq p_{\tilde{v}} \cap B(\tilde{v}, e^{-t_0})^c} |W(\mathcal{Y}, \tilde{v})|^k\right) \leq m.$$

Proof. In this proof, we write $p_{\tilde{v}} = p$ to lighten notation. Let us first consider the case when p does not contain a portion of the spine corresponding to the skeleton.

We parametrise $\tilde{\gamma}_{\tilde{v}}$ in $[0, 1]$ as before and assume that the parametrisation of $\tilde{\gamma}_{\tilde{v}}$ and p is the same until they start to differ. We drop \tilde{v} and write $\tilde{\gamma}_{\tilde{v}} = \gamma$ throughout the rest of this proof for notational clarity. To differentiate the first and last exit times for these two parametrisations, we write $t_1(p)$ (resp., $t_2(p)$) for the first exit of $B(\tilde{v}, e^{-t-1})$ (resp., last exit of $B(\tilde{v}, e^{-t})$) of p . We also write $t_1 = t_1(\gamma)$ and $t_2 = t_2(\gamma)$ for clarity.

Note that it is always the case that $t_1(\gamma) = t_1(p)$ and that $t_2(\gamma) \leq t_2(p)$. On the event $t_2(p) = t_2(\gamma)$, the maximum in this lemma is over the same set as the one in Lemma 6.7, so there is nothing to prove. We focus now on the case where $t_2(p) > t_2(\gamma)$, meaning that a part of the spine comes back close to x on the cover \tilde{M}' after the time that a non-contractible loop was created.

We start with the case of the torus. Let us denote by S the spine attached to \tilde{v} (meaning the full bi-infinite path). Since S is a periodic path, it has a well-defined direction d (say represented by a unit vector in \mathbb{R}^2) and separates the plane into two sets right and left of S . Let us assume without loss of generality that the direction d is horizontal and that \tilde{v} is below S . Let τ be a time at which the vertical coordinate of γ reaches its maximum (for topological reasons, this has to occur on the spine S), and let us define γ' by appending to $\gamma[0, \tau]$ a vertical segment going up to infinity.

Now, it is easy to see that the winding of a straight segment around a point is bounded by π , so

$$\max_{\mathcal{Y} \subseteq \gamma' \cap B(\tilde{v}, e^{-t-1})^c} |W(\mathcal{Y}, \tilde{v})| \leq \max_{\mathcal{Y} \subseteq \gamma \cap B(\tilde{v}, e^{-t-1})^c} |W(\mathcal{Y}, \tilde{v})| + \pi. \tag{6.6}$$

Using Lemma 6.4, this implies that we can always bound

$$\begin{aligned} \max_{\mathcal{Y} \subseteq \rho_{\tilde{v}} \cap B(\tilde{v}, e^{-t-1})^c} |W(\mathcal{Y}, \tilde{v})| &\leq \max_{\mathcal{Y} \subseteq \gamma' \cap B(\tilde{v}, e^{-t-1})^c} |W(\mathcal{Y}, \tilde{v})| + 4\pi \\ &\leq \max_{\mathcal{Y} \subseteq \gamma \cap B(\tilde{v}, e^{-t-1})^c} |W(\mathcal{Y}, \tilde{v})| + 5\pi. \end{aligned} \tag{6.7}$$

By Lemma 6.7 (the $\tilde{\gamma}_{\tilde{v}}$ there is γ here) applied to all scales up to scale t , the right-hand side is a sum of $O(t)$ variables with uniform stretched exponential tails.

Let us now bound $\mathbb{P}(t_2(\mathfrak{p}) > t_2(\gamma); \mathfrak{p} \cap \mathfrak{s} = \emptyset)$. Note that, on the event $t_2(\mathfrak{p}) > t_2(\gamma)$, \tilde{v} has to be at a distance less than e^{-t} of its spine. Using the fact that we can sample points in any order in Wilson’s algorithm, we can bound that probability by first doing a cutset exploration around \tilde{v} at scale $r_{\tilde{v}}$. After this cutset exploration, the probability that any of the branches sampled intersects $B(\tilde{v}, e^{-t})$ is at most $C e^{-c(t-t_0)}$ for some constants C, c (see Lemma 5.7). However, after the cutset exploration around \tilde{v} , the spine attached to \tilde{v} is necessarily fully sampled (see, e.g., Proposition 5.10, final assertion). Hence,

$$\mathbb{P}(t_2(\mathfrak{p}) > t_2(\gamma); \mathfrak{p} \cap \mathfrak{s} = \emptyset) \leq C e^{-c(t-t_0)}. \tag{6.8}$$

Thus, overall, either $\max_{\mathcal{Y} \subseteq \rho_{[\tau_1, \tau_2]} \cap B(\tilde{v}, e^{-t-1})^c} |W(\mathcal{Y}, \tilde{v})|$ is the same as the variable in Lemma 6.7, or with probability at most $C e^{-c(t-t_0)}$, it is at most a sum of $O(t)$ variables with uniform stretched exponential tail. The moment bound now easily follows.

For the case where the universal cover is the unit disc, a similar argument can be done provided that we can construct a path γ' which is nice and avoids the spine. Let us also introduce τ_S as the first time γ hits the spine and τ_{nc} as the ending time of γ . As recalled in more detail in appendix A, it is a well-known fact from Riemann surfaces that there exists a Möbius transform $\phi = \phi_{\gamma, \tilde{v}}$ such that

$$\mathfrak{p}(v) = \gamma[0, \tau_S] \cup \bigcup_{n \geq 0} \phi^{(n)}((\tau_S, \tau_{nc})),$$

where $\phi^{(n)} = \phi \circ \dots \circ \phi$ and the union is disjoint. In other words, we keep applying the same Möbius transform to obtain all the copies. Furthermore, ϕ can be written as $\phi = \Phi^{-1} \circ \mu \circ \Phi$, where Φ is a Möbius map from the unit disc to the upper half-plane and μ is either the translation by ± 1 or a multiplication by $\lambda > 1$.

If it is a translation, we use the same argument as before. Otherwise, it is a scaling by $\lambda > 1$ in which case $\phi(S)$ necessarily converges to infinity. Furthermore, let τ be the time at which $\mathfrak{S}(\Phi(\gamma))$ reaches its minimum; we define γ' by appending a straight vertical segment ℓ from $\Phi(\gamma(\tau))$ to \mathbb{R} (note that this segment may not intersect $\Phi(\gamma)$ nor the subsequent scalings of the image of the portion of the spine $\Phi(\gamma([\tau_S, \tau_{nc}]))$ and then mapping $\Phi(\gamma[0, \tau] \cup \ell)$ back to the disc by Φ^{-1} . By construction, it is trivial to check that $\Phi(\gamma')$ does not intersect $\Phi(p(\tau, \infty))$, so γ' does not intersect $p(\tau, \infty)$. On the other hand, Φ^{-1} is the image of a line segment by a Möbius transform, so it is a circular arc, and therefore, its winding around any point is bounded by 2π . We can then conclude using exactly the same reasoning as in the torus case, with only the constant π replaced by 2π on the right-hand side of (6.6), and hence obtain the analogue of (6.7).

Finally, let us consider the case when p contains a portion of the spine corresponding to the skeleton. Note that the whole argument was geometric and the only place where we used the probability was in (6.8) which is provided by Lemma 5.6. ■

Lemma 6.9. *For all $k \geq 1$, there exists a constant $C > 0$ such that the following holds. Fix a compact set $K \subset \tilde{M}'$. Suppose that we are in the setup of Lemma 6.8. Let $p = p_{\tilde{v}}$ be as above and \bar{p} the curve which starts at \tilde{v} , hits the spine, and then goes to infinity in the direction opposite to the orientation of the spine. Orient and parametrise both p, \bar{p} from \tilde{v} to ∞ so that t_2 is the last time they exit $B(\tilde{v}, e^{-t})$. Then, for all $\delta < \delta_K, t \in [t_0, \infty], \tilde{v} \in K$,*

$$\mathbb{E}\left(\sup_{\mathcal{Y} \subset p[t_2, \infty)} |W(\mathcal{Y}, \tilde{v})|^k\right) \leq C((1+t) \wedge \log(1/\delta))^k,$$

$$\mathbb{E}\left(\sup_{\mathcal{Y} \subset \bar{p}[t_2, \infty)} |W(\mathcal{Y}, \tilde{v})|^k\right) \leq C((1+t) \wedge \log(1/\delta))^k.$$

We emphasise here that the spine might be identical to a spine of a skeleton branch.

Proof. First, we tackle the case of p . If $t < \log(C'\delta_0/\delta)$, we need to add $O(t)$ many variables given by Lemma 6.8, and hence, we are done by Minkowski's inequality. If $t > \log(C'\delta_0/\delta)$, first, we apply the previous bound up to $t = \log(C'\delta_0/\delta)$. For the remainder, we can simply bound the winding by the volume of the ball of radius $C'\delta/\delta_0$ around \tilde{v} which is $O(1)$ by our assumptions on the graph (see assumption (i) in Section 2.4). Indeed, this quantity is bounded by the number of times p crosses a straight line joining \tilde{v} to the boundary of the ball, which is simply bounded by the

volume of the ball. For \bar{p} , we can use Lemma 6.4 and the bound for p and the proof is complete. ■

Lemma 6.10. *For all k , there exist $C, c > 0$ such that for all points \tilde{u}, \tilde{v} , for all compact sets \tilde{K} containing $\tilde{u}, \tilde{v} \in \tilde{M}'$, for all $\delta < \delta_{\tilde{K}}$,*

$$\mathbb{E}(|W(p_{\tilde{u}}, \tilde{v})|^k) \leq C(1 + \log |\tilde{u} - \tilde{v}| \wedge \log(1/\delta))^c.$$

Proof. Fix $T = \log(\frac{C'\delta_0}{\delta})$ with C' as in Lemma 6.8. We parametrise $p_{\tilde{v}}, p_{\tilde{u}}$ as in Lemma 6.9. If $|\tilde{u} - \tilde{v}| < e^{-T}$, we use Lemmas 6.4 and 6.9 to conclude.

If $|\tilde{u} - \tilde{v}| \geq e^{-T}$, we take $t = -\log |\tilde{u} - \tilde{v}|$. Then, using Lemma 6.4, we can write

$$|W(p_{\tilde{u}}, \tilde{v})| \leq 2\pi + \sum_{k=t}^{\infty} \sup_{\mathcal{Y} \subset p_{\tilde{v}}[k_2, \infty)} |W(\mathcal{Y}, \tilde{v})| 1_{|p_{\tilde{u}} - \tilde{v}| \in [e^{-k}, e^{-k-1}]},$$

where k_2 is the last exit from the $B(\tilde{v}, e^{-k})$ by $p_{\tilde{u}}$. Using Lemma 5.4 or Lemma 5.6 whichever is appropriate, we get an exponential bound on the expectation of the indicator event above. (Notice that, for scales $k > \log(C'\delta_0/\delta)$, the probability actually becomes 0, so this is a finite sum.) Therefore, we can again use Lemma 6.9 and Cauchy–Schwarz to conclude. ■

6.4. Convergence of height function

In this section, we prove Theorem 6.1.

Let us describe informally the general structure of the proof. Most of the work will be in the universal cover to obtain the convergence of expressions of the form $\mathbb{E} \prod_i (h(\tilde{z}_i) - h(\tilde{w}_i))$ for distinct points \tilde{z}_i and \tilde{w}_i (Lemma 6.12). By integrating this expression, we will then obtain the convergence of the height function, seen as a function on the universal cover.

For the first part, the idea is to introduce a regularised height function $h_t^{\#\delta}$ which is a continuous function of the CRSF so that the convergence of $h_t^{\#\delta}$ as $\delta \rightarrow 0$ is immediate. This leaves us with two issues: the comparison between $h^{\#\delta}$ and $h_t^{\#\delta}$ for fixed δ and the convergence of h_t as $t \rightarrow \infty$ in the limit. Both questions are actually solved simultaneously by the estimate of Lemma 6.12 which compares $h^{\#\delta}$ and $h_t^{\#\delta}$ with an error term that becomes small both in δ and t independently.

Setup and notations. Recall that our goal is to prove that

$$\int_{\tilde{M}'} \bar{h}_{\text{ext}}^{\#\delta}(z) \tilde{f}(z) d\mu(z) \tag{6.9}$$

converges in law and also in the sense of moments as $\delta \rightarrow 0$. (Recall that \bar{X} denotes the centred random variable $\bar{X} = X - \mathbb{E}(X)$ whenever this expectation is well defined.)

Implicitly, $h^{\#\delta}$ is sampled according to the dimer law (1.1). In other words, by Theorem 3.3, we need to first sample a Temperleyan pair $(\mathcal{T}, \mathcal{T}^\dagger)$ and apply the bijection ψ in that theorem.

However, it turns out to be more convenient in the proof to work with a pair $(\mathcal{T}, \mathcal{T}^\dagger)$ sampled from \mathbb{P}_{Wils} . To this pair $(\mathcal{T}, \mathcal{T}^\dagger)$ we can apply the bijection ψ from Theorem 3.3 and obtain a dimer configuration \mathbf{m} whose centred height function can be studied. We will first prove (6.9) for \mathbb{P}_{Wils} and then later explain how this implies the same for \mathbb{P}_{Temp} .

Recall that since $\int_{\tilde{M}'} \tilde{f}(z) d\mu(z) = 0$, the expression in (6.9) is in fact well defined. We wish to compute the moments of this integral but only in terms of height differences since the field is a priori defined only up to constant. To do this, we use the following trick. Note that since

$$\int \tilde{f} d\mu = 0, \quad \int \tilde{f}^+ d\mu = \int \tilde{f}^- d\mu =: Z(\tilde{f}),$$

where $f^+ = \max\{f, 0\}$ and $f^- = -\min\{f, 0\}$. Now, we can write

$$\int_{\tilde{K}} \bar{h}_{\text{ext}}^{\#\delta}(z) \tilde{f}(z) d\mu(z) = \int_{\tilde{K} \times \tilde{K}} (\bar{h}_{\text{ext}}^{\#\delta}(z) - \bar{h}_{\text{ext}}^{\#\delta}(w)) \frac{\tilde{f}^+(z) \tilde{f}^-(w)}{Z(\tilde{f})} d\mu(z) d\mu(w), \tag{6.10}$$

which implies

$$\begin{aligned} & \left(\int_{\tilde{K}} \bar{h}_{\text{ext}}^{\#\delta}(z) \tilde{f}(z) d\mu(z) \right)^k \\ &= \int_{\tilde{K}^{2k}} \prod_{i=1}^k (\bar{h}_{\text{ext}}^{\#\delta}(z_i) - \bar{h}_{\text{ext}}^{\#\delta}(w_i)) \frac{\tilde{f}^+(z_i) \tilde{f}^-(w_i)}{Z(\tilde{f})} d\mu(z_i) d\mu(w_i). \end{aligned} \tag{6.11}$$

Therefore, we are interested in the k -point function, $\mathbb{E}[\prod_{i=1}^k (\bar{h}_{\text{ext}}^{\#\delta}(z_i) - \bar{h}_{\text{ext}}^{\#\delta}(w_i))]$. Pick k distinct pairs of points

$$(z_1, w_1), \dots, (z_k, w_k) \in \tilde{K},$$

and let $(f(z_1), f(w_1)), \dots, (f(z_k), f(w_k))$ be the faces containing them. Let $z_i^{\#\delta}, w_i^{\#\delta}$ be the midpoint of the diagonals of $f(z_i), f(w_i)$.

Now, recall Theorem 4.10 which relates the dimer height difference between two faces f and f' to the winding of a specific path $\gamma_{f,f'}$ connecting $m(f)$ and $m(f')$ and additional terms of the form $\pm\pi$ associated with jumping over components of the CRSF. Let $\gamma_i^{\#\delta}$ be the path $\gamma_{f,f'}$ when $f = f(z_i)$ and $f' = f(w_i)$, and orient it from $z_i^{\#\delta}$ to $w_i^{\#\delta}$. Then, with these notations, Theorem 4.10 says that

$$h(z_i^{\#\delta}) - h(w_i^{\#\delta}) = W(\gamma_i^{\#\delta}, z_i^{\#\delta}) + W(\gamma_i^{\#\delta}, w_i^{\#\delta}) + \Psi_i^{\#\delta}, \tag{6.12}$$

where $\Psi_i^{\#\delta}$ is the $\pi \sum_S (\varepsilon_S + \delta_S)$ term in Theorem 4.10. We drop the superscript δ from now on for clarity. Thus, from now on, we focus on proving convergence of

$$\mathbb{E}_{\text{Wils}} \left(\prod_{i=1}^k (\bar{W}(\gamma_i, z_i) + \bar{W}(\gamma_i, w_i) + \bar{\Psi}_i) \right). \tag{6.13}$$

Let

$$\mathcal{X} = \{z_i, w_i : 1 \leq i \leq k\}, \text{ and assume } p(u) \neq p(v) \text{ for all } u \neq v \in \mathcal{X}.$$

Clearly, (6.13) can be expanded as a sum of 2^k many terms of the form

$$\mathbb{E}_{\text{Wils}} \left(\prod_{x \in S} (\bar{W}(\gamma_x, x) + \bar{\Psi}_x/2) \right) =: \mathbb{E}_{\text{Wils}} \left(\prod_{x \in S} \bar{F}_x(\gamma_x) \right), \tag{6.14}$$

where S is a subset of vertices of \mathcal{X} with distinct indices and γ_x is γ_i for some i such that $x = z_i$ or w_i . (Of course, the products in the expansion of (6.14) have further restrictions, but we ignore that for clarity.)

We are interested in estimating and proving convergence of (6.14). To that end, we employ an idea similar to that in [4]: we truncate the CRSF branch at a macroscopic scale and deal with the truncated macroscopic winding and the remaining microscopic winding part separately.

We now fill in the details. Parametrise $\gamma = \gamma_x$. Define $(\gamma_x(t))_{t \geq 1}$ to be γ_x from the first entry time into the ball $B(x, e^{-t})$. (Note that at the moment, the balls $B(x, e^{-t})$ might overlap for different $x \in \mathcal{X}$.) We emphasise here that we parametrise γ in the opposite direction compared to Sections 6.2 and 6.3 to be consistent with [4]. With an abuse of notation, we will denote by $\gamma_x(-\infty, t]$ the whole path from the opposite end of γ up until $\gamma_x(t)$. Define the regularised term and the error term as

$$\begin{aligned} \bar{F}_x(\gamma_x, t) &:= \bar{W}(\gamma_x(-\infty, t], x) + \bar{\Psi}_x/2, \\ \bar{e}_x(t) &:= \bar{W}(\gamma_x, x) - \bar{W}(\gamma_x(-\infty, t], x) = \bar{F}_x(\gamma_x) - \bar{F}_x(\gamma_x, t). \end{aligned}$$

When we want to emphasise the role of z, w , we write $\gamma_{zw}, \Psi_{zw}, F_z(\gamma_{zw})$ in place of the above. We start with a general moment bound for the truncated winding.

Lemma 6.11. *For all m , there exist constants $c = c(\tilde{K}, m)$, $\alpha = \alpha(m)$ such that, for all $z, w, \delta < \delta_{\tilde{K}}, t \geq 0$,*

$$\mathbb{E}_{\text{Wils}} (|\bar{F}_z(\gamma_{zw}, t)|^m) \leq c(1 + t + |\log|z - w||) \wedge \log(1/\delta)^\alpha.$$

Proof. The proof is immediate from Lemmas 6.9 and 6.10 and the fact that Ψ_{zw} is bounded by a constant $c(K)$ times the number of non-contractible loops in the CRSF. Indeed, this is clear from the fact that for a fixed compact set $\tilde{K} \subset \tilde{M}$ there

exists a constant $c(\tilde{K})$ such that any curve $P \subset \tilde{K}$ will cross at most $c(\tilde{K})$ copies of a given spine (or lift of a loop). Since the number of non-contractible loops has superexponential tail by Theorem 6.3, the moment bound of Ψ_{zw} is immediate. ■

By compactness, choose $r_{\tilde{K}}$ small enough so that p is injective in $B(z, r_{\tilde{K}})$ for all $z \in \tilde{K}$. Now, let $r_{\mathcal{X}}$ be defined as in (5.1) but for the set of vertices $p(\mathcal{X})$ (which are all distinct by assumption on \mathcal{X}). We observe that $r_{\mathcal{X}} \geq c(\tilde{K}) \min_{x \neq y \in \mathcal{X}} |x - y|$ for some constant $c(\tilde{K})$. We set

$$\Delta = \Delta(\mathcal{X}) = \frac{1}{10}(r_{\mathcal{X}} \wedge r_{\tilde{K}}). \tag{6.15}$$

Lemma 6.12. *There exist constants $c, c' > 0$ such that for all $m, m' \geq 1$ there exists $\alpha > 0$ such that the following holds. Let $S \subset \mathcal{X}, S' \subset \mathcal{X}$ be disjoint. Also, assume that $|S| = m, |S'| = m' \geq 0$. Let Δ be as in (6.15). Then, for all $\delta < \delta_{\tilde{K}}$ and $t \geq \log(r_{\tilde{K}}/\Delta)$,*

$$\left| \mathbb{E}_{\text{Wils}} \left(\prod_{x \in S} \bar{e}_x(t) \prod_{x \in S'} \bar{F}_x(t) \right) \right| \leq c |1 + t|^\alpha (e^{-c(t - \log(r_{\tilde{K}}/\Delta))} + \delta^{c'}).$$

Proof. For simplicity, we first assume that S' is empty. Recall that we can sample a CRSF from \mathbb{P}_{Wils} by first sampling the branches \mathfrak{s} and then the rest by Wilson’s algorithm. Note that Lemma 5.6 tells us that the isolation radius corresponding to each vertex in S with respect to the skeleton has polynomial tail. Since after sampling the skeleton the rest of the branches are sampled simply by Wilson’s algorithm, we can conclude that the application of Proposition 5.10 is valid with Δ in place of r_i .

We perform the coupling in Section 5.2 (in particular Proposition 5.10) with points $p(\mathcal{X})$ and their lift \mathcal{X} and a compact domain $D \subset \tilde{M}'$ containing all points in \mathcal{X} so that the minimal Euclidean distance between any point in \mathcal{X} and ∂D is at least $r_{\mathcal{X}}$. Note that we can choose $r_{\mathcal{X}}$ for the r in (5.1) there. Call $(\mathcal{T}_x^D)_{x \in \mathcal{X}}$ the resulting independent UST in D . For $x \in \mathcal{X}$, let γ_x^D be the branch in the UST \mathcal{T}_x^D joining x to ∂D . Let $\gamma_x^D(t)$ be parametrised so that γ_x^D enters $B(x, e^{-t} r_{\tilde{K}})$ for the first time (going from the outside to x) at time t .

Let R_x be the isolation radius of x in the application of Proposition 5.10 and write $R_x = e^{-I_x} r_{\tilde{K}}$. Let $I = \max_{x \in \mathcal{X}} I_x$. (Note that, for notational clarity, I here is shifted by $\log(r_{\tilde{K}}/\Delta)$ from the one in Section 5.2.) We now decompose

$$e_x(t) = \underbrace{\bar{W}(\gamma_x^D(t, \infty), x)}_{\alpha_x} + \underbrace{(\bar{W}(\gamma_x(t, \infty), x) - \bar{W}(\gamma_x^D(t, \infty), x))}_{\xi_x}.$$

Therefore, we need to deal with expectation of products of α_x, ξ_x for different indices x . Let \mathfrak{C} be the sigma algebra generated by the cutset exploration. We will first compute the conditional expectation and then take the overall expectation. Note that

\mathcal{T}_x^D are independent for different x and also independent of \mathbb{C} , and hence, any term involving α_x is 0. Thus, we only have to deal with terms involving $\xi_x, x \in S$.

Let Λ be the *last* time γ_x enters $B(x, e^{-I}r_{\tilde{K}})$. In the coupling, γ_x, γ_x^D agree inside $B(x, e^{-I}r_{\tilde{K}})$, so in particular $\gamma_x(\Lambda, \infty) = \gamma_x^D(\Lambda, \infty)$, so

$$\xi_x = \bar{W}(\gamma_x(t, \Lambda), x) - \bar{W}(\gamma_x^D(t, \Lambda), x).$$

Let G be the event that ξ_x without the bar (meaning without the expectation terms) is 0. Observe that if G does not occur, then one of the following events happen. Either

$$I - \log(r_{\tilde{K}}/\Delta) > (t - \log(r_{\tilde{K}}/\Delta))/2$$

which has probability at most $e^{-c(t-\log(r_{\tilde{K}}/\Delta))} \vee \delta^c \leq (e^{-c(t-\log(r_{\tilde{K}}/\Delta))} + \delta^c)$ (Proposition 5.10). Otherwise, γ_x has to exit $B(x, e^{-I}r_{\tilde{K}})$ after hitting $e^{-t}r_{\tilde{K}}$. This also has probability at most $e^{-c(t-\log(r_{\tilde{K}}/\Delta))}$ (Lemma 5.8). Now, we bound the moments of ξ_x . Notice that we can write

$$\begin{aligned} \mathbb{E}_{\text{Wils}}(|\xi_x|^k) &\leq \sum_{j=t/2+\log(r_{\tilde{K}}/\Delta)/2}^{\infty} \mathbb{E}(|\xi_x|^k 1_{j \leq I \leq j+1}) + ce^{-c(t-\log(r_{\tilde{K}}/\Delta))} \\ &\leq \sum_{j=t/2+\log(r_{\tilde{K}}/\Delta)/2}^{\infty} |1+j|^k (e^{-c(j-\log(r_{\tilde{K}}/\Delta))} + \delta^c), \end{aligned}$$

where we first used Cauchy–Schwarz and then used Lemma 6.9 to bound the moment and Proposition 5.10. Thus, overall,

$$\mathbb{E}_{\text{Wils}}(|\xi_x|^k) \leq |1+t|^k (e^{-c(t-\log(r_{\tilde{K}}/\Delta))} + \delta^c).$$

We have a product of at most m terms, and so, using Hölder’s inequality, taking $\alpha = m$ works.

Finally, if S' is non-empty, the proof is exactly the same as the vertices in S' are distinct from S , and hence, local independence from the coupling still holds. We then use Lemma 6.11 to conclude using Cauchy–Schwarz. Details are left to the reader. ■

Corollary 6.13. *Let $S \subset \mathcal{X}$ containing vertices with distinct indices such that $|S| = m$. There exists a constant $c = c(\tilde{K}, m), \alpha = \alpha(m)$ such that, for all $\delta < \delta_{\tilde{K}}$,*

$$\left| \mathbb{E}_{\text{Wils}} \left(\prod_{x \in S} \bar{F}_x(\gamma_x) \right) \right| \leq c |1 + \log^\alpha(\Delta)|.$$

Proof. Decompose

$$\mathbb{E}_{\text{Wils}} \left(\prod_{x \in S} \bar{F}_x(\gamma_x) \right) = \mathbb{E}_{\text{Wils}} \left(\prod_{x \in S} (\bar{F}_x(\gamma_x, t) + \bar{e}_x(t)) \right)$$

with $t = 2 \log(r_{\tilde{K}}/\Delta)$. Then, this is a straightforward application of Lemmas 6.11 and 6.12. ■

Now, we prove a convergence of the height function integrated against f in the sense of moments (still for the Wilson law \mathbb{P}_{Wils}). Recall that X_δ converges in the sense of moments as $\delta \rightarrow 0$ if, for all k , $\mathbb{E}(X_\delta^k)$ converges as $\delta \rightarrow 0$.

Lemma 6.14. $\int_{\tilde{K}} \bar{h}_{\text{ext}}^{\#\delta}(z) \tilde{f}(z) d\mu(z)$ converges in the sense of moments under the law \mathbb{P}_{Wils} . Furthermore, the limit does not depend on the sequence $(G')^{\#\delta}$.

Proof. Using (6.11), (6.13), and (6.14), we need to prove the convergence as $\delta \rightarrow 0$ of

$$\int_{\tilde{K}^{2k}} \mathbb{E}_{\text{Wils}} \left(\prod_{x \in S} \bar{F}_x(\gamma_x^{\#\delta}) \right) \frac{f^+(z_i) f^-(w_i)}{Z(\tilde{f})} d\mu(z_i) d\mu(w_i). \tag{6.16}$$

We use Fubini to bring the expectation inside the integral in the above display. We first observe that integrating (6.16) over all sets of vertices S so that $p(x) \neq p(y)$ for all $x \neq y \in S$ is enough. (Recall that $p : \tilde{M}' \mapsto M'$ is the covering map.) Indeed, using Lemma 6.11 and the fact that $\|f\|_\infty < \infty$, we see that the integrand can be bounded by $O(\log^\alpha(1/\delta))$. Since the volume of $\{S \subset \tilde{K}^{2k} : p(x) = p(y) \text{ for some } x, y \in S\}$ is $O(\delta)$ (indeed the number of pre-images of any vertex in \tilde{K} is bounded by a constant depending only on \tilde{K}), we see that the integral over this set is $O(\delta \log^\alpha(1/\delta))$.

Thus, we now concentrate on the integral (6.16) over

$$\mathcal{A}(\tilde{K}) = \{\text{sets of vertices } S \text{ so that } p(x) \neq p(y) \text{ for all } x \neq y \in S\}. \tag{6.17}$$

We now use Corollary 6.13 and dominated convergence theorem. Since $\|f\|_\infty < \infty$ and $\log^m(\Delta)$ is integrable for any $m > 0$, we need to prove the convergence of the expectation inside the integral above. Now, we claim that the regularised part

$$\mathbb{E}_{\text{Wils}} \left(\prod_{x \in S} \bar{F}_x(\gamma_x^{\#\delta}, t) \right) \tag{6.18}$$

converges as $\delta \rightarrow 0$. This follows from our assumption of a.s. convergence of Temperleyan CRSF and because the term inside the expectation is a.s. continuous function of the Temperleyan CRSF. The same can be said about the skeleton using the absolute continuity statement given by Proposition 5.3. Indeed, this follows from a.s. continuity properties of SLE_2 and the fact that the CRSF branches are made a.s. from finitely many chunks of SLE_2 ; in particular, for a fixed t , the SLE curve is a.s. not a tangent to the boundary of circle of radius e^{-t} . (We refer to [5, Section 3] for details.) Lemma 6.11 tells us that the random variable in (6.18) is uniformly integrable which completes the proof of convergence of the regularised term (6.18).

Now, given a fixed S , choose $t \geq \log(r_{\tilde{K}}/\Delta)$. Now, we claim that the error term satisfies

$$\left| \mathbb{E}_{\text{Wils}} \left(\prod_{x \in S} \bar{F}_x(\gamma_x) \right) - \mathbb{E}_{\text{Wils}} \left(\prod_{x \in S} \bar{F}_x(\gamma_x, t) \right) \right| < c|1 + t|^\alpha (e^{-c(t - \log(r_{\tilde{K}}/\Delta))} + \delta^{c'}). \tag{6.19}$$

Indeed, (6.19) follows by writing $\bar{F}_x(\gamma_x) = \bar{F}_x(\gamma_x, t) + \bar{e}_x(t)$ and then expanding and using the bounds in Lemmas 6.11 and 6.12. To finish the proof of the lemma, fix $\varepsilon > 0$. First, choose t large enough and then δ small enough so that the right-hand side of (6.19) is less than ε . Next, choose a smaller δ if needed so that, for all $\delta' < \delta$,

$$\left| \mathbb{E}_{\text{Wils}} \left(\prod_x \bar{F}_x(\gamma_x^{\#\delta'}, t) \right) - \mathbb{E}_{\text{Wils}} \left(\prod_x \bar{F}_x(\gamma_x^{\#\delta}, t) \right) \right| < \varepsilon$$

via the convergence of the regularised term. This completes the proof as ε is arbitrary.

Observe that (6.19) completes the proof that the limit does not depend on the sequence $((G')^{\#\delta})_{\delta > 0}$ since the main term (6.18) is measurable with respect to the limiting continuum Temperleyan CRSF which is universal by Theorem 6.3. ■

It remains to prove convergence in law (still for \mathbb{P}_{Wils} for now). For this, we need to alter the definition of truncation slightly. Given $t > 0$, take a cover $\{B_{\text{euc}}(x, r_{\tilde{K}}e^{-10t}) : x \in \tilde{K}\}$ of \tilde{K} and take a finite subcover. Let $c(z), c(w)$ denote the centre of one of the balls (chosen arbitrarily) in the finite subcover to which z, w belong. Define $F_x(t), e_x(t)$ to be the same as $F_x(t)$ and $e_x(t)$, but we cut off the first time γ_x enters $B(c(x), e^{-t})$. (So compared to the above, the centre of the cutoff ball is shifted by an amount which is at most e^{-10t} .) However, in this definition, we still measure winding around x , not $c(x)$.

Lemma 6.15. *The statements of Lemmas 6.11 and 6.12 still hold if we replace F, e by F, e everywhere.*

Proof. For Lemma 6.11, the proof identically follows from Lemmas 6.9 and 6.10 by noticing that the supremum over all continuous subpaths contains the portion of the branch until it first hits the shifted ball. Notice also that we shift only by an additive term in the exponential scale which is $O(t)$.

For Lemma 6.12, the proof is also identical; in particular, we still consider the coupling around the points in S . Because we still shift only by $O(t)$ in the exponential scale, the proof readily follows. ■

We can now complete the proof of Theorem 6.1.

Proof of Theorem 6.1. Again, we first prove this under \mathbb{P}_{Wils} before explaining how to extend this and concluding the proof of Theorem 6.1 to \mathbb{P}_{Temp} . We write $r = r_{\tilde{K}}$ and

pick a $t > 0$ to be taken large later. Recall from (6.10) and (6.12) that

$$\begin{aligned} & \int_{\tilde{K}} \bar{h}_{\text{ext}}^{\#\delta}(z) \tilde{f}(z) d\mu(z) \\ &= \int_{\tilde{K} \times \tilde{K}} (\bar{W}(\tilde{\gamma}_{zw}^{\#\delta}, z) + \bar{W}(\tilde{\gamma}_{zw}^{\#\delta}, w) + \bar{\Psi}_{zw}) \frac{\tilde{f}^+(z) \tilde{f}^-(w)}{Z(\tilde{f})} d\mu(z) d\mu(w). \end{aligned}$$

Here, $Z(\tilde{f})$ is a deterministic constant, so we can assume it to be 1 from now on without loss of generality. Firstly, we recall that we only need to integrate over the set $\mathcal{A}(\tilde{K})$ as in (6.17) as the integral over the remaining part is $O(\delta \log^c(1/\delta))$. Let us denote by $Y = Y(\delta)$ the above integral over $\mathcal{A}(\tilde{K})$. Introduce

$$\begin{aligned} X(\delta, t) &:= \int_{\mathcal{A}(\tilde{K})} (\bar{F}_{c(z)c(w)}(t) + \bar{F}_{c(w)c(z)}(t)) 1_{|c(z)-c(w)| > re^{-t/10}} \\ &\quad \times \tilde{f}^+(z) \tilde{f}^-(w) d\mu(z) d\mu(w). \end{aligned}$$

Note that $X(\delta, t)$ is a sum of regularised winding of finitely many branches and hence converges in law (as $\delta \rightarrow 0$ and t is fixed) by our assumption.

Now, we show that, for all $t > 0$ and $\delta < C'\delta_0 e^{-t}$,

$$\mathbb{E}(Y(\delta) - X(\delta, t))^2 \leq c(1+t)^\alpha e^{-c't}.$$

We expand the above square to get an integral over $(z, w, z', w') \in \mathcal{A}(\tilde{K})^2$. We can again without loss of generality restrict to the set $\mathcal{A}_2(\tilde{K})$ such that the projection p of all four points (z, w, z', w') maps to pairwise distinct points on M . In that case, let $\Delta = \Delta(z, w, z', w')$ be as in (6.15). We now argue that this integral over the set $\Delta \leq e^{-t/11}$ is exponentially small in t . Indeed, we are integrating over a set which has exponentially small measure with respect to μ^4 , and furthermore, the moments are integrable using Corollary 6.13 and Lemma 6.15.

For the rest of the integral, we write

$$(\bar{W}(\tilde{\gamma}_{zw}^{\#\delta}, z) + \bar{W}(\tilde{\gamma}_{zw}^{\#\delta}, w) + \bar{\Psi}_{zw}) = \bar{F}_{zw}(t) + \bar{F}_{wz}(t) + \bar{e}_{zw}(t) + \bar{e}_{wz}(t),$$

and we expand again the product inside the integral to get products of terms of the form

$$\bar{e}_{zw}(t) + \bar{e}_{wz}(t); \quad \bar{F}_{zw}(t) + \bar{F}_{wz}(t) - \bar{F}_{c(z)c(w)}(t) - \bar{F}_{c(w)c(z)}(t).$$

Note here that the indicator over $|c(z) - c(w)| > e^{-t/10}$ is included in $\Delta > e^{-t/11}$, so we can get rid of the indicator. Products containing at least one \bar{e} are small because of Lemma 6.15. Also, note that on $\Delta > e^{-t/11}$, $|z - w| > c(\tilde{K})e^{-t/11}$. Thus, we need to show that

$$\mathbb{E}(|\bar{F}_{zw}(t) - \bar{F}_{c(w)c(z)}(t)|^2 1_{|z-w| > c(\tilde{K})e^{-t/11}}) \leq c(1+t)^\alpha e^{-c't}.$$

Observe that

$$\begin{aligned} & (\bar{F}_{zw}(t) - \bar{F}_{c(z)c(w)}(t))1_{|z-w|>c(\tilde{K})}e^{-t/11} \\ &= (\bar{F}_{zw}(t) - \bar{F}_{c(z)w}(t))1_{|z-w|>c(\tilde{K})}e^{-t/11} \\ &+ (\bar{F}_{c(z)w}(t) - \bar{F}_{c(z)c(w)}(t))1_{|z-w|>c(\tilde{K})}e^{-t/11}. \end{aligned}$$

If $p(z)$ and $p(c(z))$ merge before exiting $B(c(z), e^{-5t})$, then the paths we need to consider are identical outside $B(c(z), e^{-t})$; only the centres around which we measure the topological winding are different. This difference in winding is therefore deterministically $O(e^{-9t})$. On the other hand, the paths we need to consider are not identical outside $B(c(z), e^{-5t})$ if and only if a simple random walk from z does not hit $p(c(z))$ before exiting $B(c(z), e^{-5t})$ which has probability exponentially small in t by a Beurling-type lemma (Lemma 5.1). Hence, we can conclude using Cauchy–Schwarz and Lemma 6.15. (Note that here the shifted cutoff is particularly useful since the cutoff point is the same for both z and $c(z)$.)

The proof of the rest of the statements is a simple consequence of the fact that the main term $X(\delta, t)$ is an a.s. continuous function of $\mathcal{T}^{\#\delta}$ for a fixed t . So, trivially, $(X(\delta, t), \mathcal{T}^{\#\delta})$ converges jointly in law; the limit does not depend on the sequence $(G')^{\#\delta}$ and is conformally invariant. It is a simple exercise to show that the L^2 bound on the error term $Y(\delta) - X(\delta, t)$ as shown above is enough to conclude the proof of the theorem (i.e., the proof of the same claim as in the previous sentence for the pair $(\bar{h}_{\text{ext}}^{\#\delta}(z), \mathcal{T}^{\#\delta})$).

We have proved convergence in law of $(\bar{h}_{\text{ext}}^{\#\delta}(z), \mathcal{T}^{\#\delta})$ for the \mathbb{P}_{Wils} law. We now explain how to convert the result so that it holds under the Temperleyan law \mathbb{P}_{Temp} . Recall that if \mathcal{T} is sampled by performing Wilson’s algorithm (after sampling the skeleton) and \mathcal{T}^\dagger is sampled from the uniform distribution among all oriented duals of \mathcal{T} , then the Radon–Nikodym derivative of $(\mathcal{T}, \mathcal{T}^\dagger)$ with respect to \mathbb{P}_{Temp} is $Z = 2^{K^\dagger} / \mathbb{E}_{\text{Wils}}(2^{K^\dagger})$, where K^\dagger is the number of nontrivial cycles of \mathcal{T}^\dagger by Lemma 3.8. Now, observe that Z is measurable with respect to $\mathcal{T}^{\#\delta}$, so $(Z, (\bar{h}_{\text{ext}}, \tilde{f}))$ jointly converges in law under \mathbb{P}_{Wils} by the above. Furthermore, all moments of Z and of $(\bar{h}_{\text{ext}}, \tilde{f})$ are bounded under \mathbb{P}_{Wils} , and therefore, we conclude that

$$\left(\int_{\tilde{M}} (h_{\text{ext}}(z) - \mathbb{E}_{\text{Wils}}(h_{\text{ext}}(z))) \tilde{f}(z) \mu(dz); \mathcal{T}^{\#\delta} \right)$$

converges in law and in the sense of moments under \mathbb{P}_{Temp} . In particular, taking expectation (under \mathbb{P}_{Temp}) of the first quantity, we also deduce that $\int_{\tilde{M}} (\mathbb{E}_{\text{Temp}}(h_{\text{ext}}(z)) - \mathbb{E}_{\text{Wils}}(h_{\text{ext}}(z))) \tilde{f}(z) \mu(dz)$ converges. Taking the difference, it follows that

$$\left(\int_{\tilde{M}} (h_{\text{ext}}(z) - \mathbb{E}_{\text{Temp}}(h_{\text{ext}}(z))) \tilde{f}(z) \mu(dz); \mathcal{T}^{\#\delta} \right)$$

converges in law and in the sense of moment under \mathbb{P}_{Temp} , as desired. Furthermore, since the limit of $\mathcal{T}^{\#\delta}$ is independent of the chosen graph sequence (subject to the assumptions in Section 2.4), so is the limit of $\int \tilde{f}(x) \bar{h}_{\text{ext}}^{\#\delta}(x) d\mu(x)$. ■

Remark 6.16. In fact, with some more work, the convergence in Theorem 6.1 can be upgraded to specifically include information about the instanton component. More precisely, using Lemma 4.13, we obtain the following.

Take an ordered finite set of continuous simple loops which forms the basis of the first homology group of M' , all endowed with a fixed orientation. Let $H^{\#\delta} \in \mathbb{R}^{4g+2b-4}$ denote the vector of height differences along these loops (i.e., for each such loop, record the height accumulated by going along the loop once in the prescribed orientation). Consider the one-form

$$d\bar{h}^{\#\delta}(e) := \bar{h}^{\#\delta}(\tilde{e}^+) - \bar{h}^{\#\delta}(\tilde{e}^-).$$

It is well known (see Theorem 2.9) that the instanton component of the above one-form is completely determined by $H^{\#\delta}$. Then, we have

$$\left(\int \tilde{f}(x) \bar{h}_{\text{ext}}^{\#\delta}(x) d\mu(x), H^{\#\delta}, \mathcal{T}^{\#\delta} \right)$$

converging jointly in law as $\delta \rightarrow 0$. The first two coordinates also jointly converge in the sense of all moments. Furthermore,

$$\lim_{\delta \rightarrow 0} \left(\int \tilde{f}(x) \bar{h}_{\text{ext}}^{\#\delta}(x) d\mu(x), H^{\#\delta} \right)$$

is measurable with respect to the limit \mathcal{T} of $\mathcal{T}^{\#\delta}$. We do not go into the details of these claims for the sake of brevity.

A. Geometry of spines

In this section, we prove Lemma 4.9 and Lemma 6.5. But before getting into the proofs, we remind the readers of certain basic facts from the classical theory of Riemann surfaces.

By the uniformisation theorem of Riemann surfaces (Section 2.2), recall that there exists a conformal map from the Riemann surface \mathbb{D}/g to M' , where g is a Fuchsian group which is a discrete subgroup of the group of Möbius transforms. Furthermore, such a conformal map is unique up to conformal automorphisms (i.e., Möbius transforms) of the unit disc. In other words, if g, g' are two Fuchsian groups such that M' is conformally equivalent to both \mathbb{D}/g and \mathbb{D}/g' , then there exists a Möbius map

$$\phi : \mathbb{D} \mapsto \mathbb{D}$$

such that

$$g' = \phi^{-1} \circ g \circ \phi.$$

Since we have fixed a canonical lift, we have defined Γ uniquely.

It is further known that g is isomorphic to the fundamental group $\pi_1(M', x)$. (Topologically, this is also known as the group of deck transformations.) This connection is described as follows: choose a base point in the manifold x_0 and a particular lift of \tilde{x}_0 ; then any simple loop ℓ based in x_0 can be lifted to a simple curve $\tilde{\ell}$ in \mathbb{D} starting from \tilde{x}_0 , with the endpoint \tilde{x}_1 depending only on the homotopy class of the loop. Then, to ℓ we associate the map $\phi_{\ell, \tilde{x}_0} \in g$ that sends \tilde{x}_0 to \tilde{x}_1 . Note then that $\phi_{\ell, \tilde{x}_0}(\tilde{\ell})$ is a curve which projects via p to the same curve ℓ in M' since M' is conformally equivalent to \mathbb{D}/g (and in particular is homeomorphic) and it does not intersect $\tilde{\ell}$ since ℓ is a simple loop. Furthermore, since the endpoint of $\phi_{\ell, \tilde{x}_0}(\tilde{\ell})$ is $\phi_{\ell, \tilde{x}_0} \circ \phi_{\ell, \tilde{x}_0}(\tilde{x}_0)$, by the unique path lifting property, the curve $\tilde{\ell} \cup \phi_{\ell, \tilde{x}_0}(\tilde{\ell})$ is the unique lift obtained by going around ℓ twice in the same direction. Iterating, we obtain that the infinite path obtained by going around ℓ in the same direction is given by $\bigcup_{n=0}^{\infty} \phi_{\ell, \tilde{x}_0}^{(n)}(\tilde{\ell})$, where $\phi_{\ell, \tilde{x}_0}^{(n)}$ is the n -fold composition of ϕ_{ℓ, \tilde{x}_0} and that the union is a disjoint union.

We can actually say more using the classification of Möbius maps according to their trace. Recall that Möbius maps preserving the unit disc have the form

$$\phi(z) = e^{i\theta} \frac{z - a}{\bar{a}z - 1}; \quad |a| < 1; \theta \in [0, 2\pi)$$

and can be classified (up to conjugation with Möbius transforms) depending upon the behaviour of the trace as follows.

- If $|e^{i\theta} - 1|^2 = 4(|a|^2 - 1)^2$, then $\phi(z)$ is conjugate to either $z + 1$ or $z - 1$, seen as maps from \mathbb{H} to \mathbb{H} .
- If $|e^{i\theta} - 1|^2 > 4(|a|^2 - 1)^2$, then $\phi(z)$ is conjugate to λz , where $\lambda > 1$, seen as maps from \mathbb{H} to \mathbb{H} .
- If $|e^{i\theta} - 1|^2 < 4(|a|^2 - 1)^2$, then ϕ is conjugate to a rotation of the unit disc.

Thus, there is a Möbius map Φ from \mathbb{D} to \mathbb{D} or \mathbb{H} such that

$$\phi = \Phi^{-1} \circ \mu_\phi \circ \Phi,$$

where μ_ϕ is either a translation by ± 1 or a scaling by λ (in case Φ maps \mathbb{D} to \mathbb{H}) or a rotation (in case Φ maps \mathbb{D} to \mathbb{D}).

We argue that, for any loop ℓ , $\mu_{\phi_{\ell, \tilde{x}_0}}$ cannot be a rotation. Indeed, if a map ϕ_{ℓ, \tilde{x}_0} was conjugate to a rotation, then its iterates would be either periodic or a dense set (on the image of a circle). Being periodic is forbidden because of path lifting property, and being dense is forbidden because g is discrete. From the two remaining cases, we can complete the proof of Lemma 4.9.

Proof of Lemma 4.9. Let S be a spine; we can choose a point x_0 on $p(S)$ and a lift \tilde{x}_0 on S . Then, we see that the previous general theory applies, so we can find an open path $\tilde{\ell}$ (a certain lift of the path going once around a non-contractible loop) and a map $\phi_{p(S),\tilde{x}_0}$ such that

$$S = \bigcup_{n \in \mathbb{Z}} \phi_{p(S),\tilde{x}_0}^{(n)}(\tilde{\ell}),$$

and the map $\phi_{p(S),\tilde{x}_0}$ is conjugate to either a scaling or a translation. The case of a scaling clearly gives a simple path between two different points (the image of 0 and ∞), while the case of a translation gives a simple loop where the image of ∞ is the unique point of $\partial\mathbb{D}$ on the loop.

In the case of the torus, the map $\phi_{p(S)}$ is simply a translation in \mathbb{C} , so the result is trivially true. ■

We now prove Lemma 6.5 which provides bounds on the macroscopic winding of the spines. We repeat the statement of the lemma here for convenience.

Lemma (Restating Lemma 6.5). *Fix a compact set $\tilde{K} \subset \tilde{M}'$. There exist constants $c, c' > 0$ such that for all $\tilde{v} \in \tilde{K}$, for all $\delta < \delta_{p(\tilde{B}_0)}$, $n \geq 1$, $j \geq 1$, and $u \in B_1$ such that $\mathbb{P}(X_{\tau_{j+1}} = u) > 0$,*

$$\mathbb{P}\left(\sup_{\mathcal{Y} \subset \tilde{X}[\tau_j, \tau_{j+1}]} |W(\mathcal{Y}, \tilde{v})| > n \mid X_{\tau_{j+1}} = u\right) \leq ce^{-c'n},$$

$$\mathbb{P}\left(\sup_{\mathcal{Y} \subset \tilde{X}[\tau_j, \tau_{j+1}]} |W(\mathcal{Y}, \tilde{v})| > n \mid X_{\tau_{j+1}} \in \partial(G')^{\#\delta} \cup \mathfrak{s}\right) \leq ce^{-c'n}.$$

Here, the supremum is over all continuous paths obtained by erasing portions of $\tilde{X}[\tau_j, \tau_{j+1}]$.

Proof. We are going to borrow the notations from Section 6.2. Take a non-contractible simple loop ℓ in M' through v and find a continuous path $\ell^{\#\delta}$ in $G^{\#\delta}$ which approximates this loop in the sense that it stays within distance $c\delta/\delta_0$ from ℓ . (This is guaranteed to exist by the uniform crossing estimate for small enough δ depending on ℓ .) Notice that the lift starting from \tilde{v} of a curve which winds infinitely many times in the clockwise (resp., anti-clockwise) direction of $\ell^{\#\delta}$ defines an infinite path $\tilde{\ell}_+$ (resp., $\tilde{\ell}_-$) which converges to the boundary in the hyperbolic case (Lemma 4.9) or goes to infinity in an asymptotic direction in case of the torus. Let $\tilde{\ell}'_{\pm}$ denote the portion of $\tilde{\ell}_{\pm}$ from the last exit of B_1 to infinity (given an arbitrary parametrisation starting from \tilde{v}). Notice that $(\tilde{\ell}'_- \cup \tilde{\ell}'_+)$ divides the annulus $\tilde{M}' \setminus B_1$ into two simply connected domains. By compactness, we can find a constant C such that the winding of $\tilde{\ell}'_+$ around \tilde{v} is bounded by C , uniformly over all points $\tilde{v} \in \tilde{K}$ for a suitable choice

of loop ℓ . Let $\tau_{+,1}$ be the first hitting time of $\tilde{\ell}'_+$ in the interval $[\tau_{i_j}, \tau_{i_j+1}]$, and by induction, define $\tau_{-,i}$ to be the first hitting time of $\tilde{\ell}'_-$ after $\tau_{+,i}$ and define $\tau_{+,i}$ to be the hitting time of $\tilde{\ell}'_+$ after $\tau_{-,i-1}$. Let I_+ be the number of $\tau_{+,i}$ before τ_{i_j+1} , i.e.,

$$I_+ = |\{i \mid \tau_{+,i} \leq \tau_{i_j+1}\}|.$$

We now use the deterministic bound

$$\sup_{\mathcal{Y} \subset \tilde{X}[\tau_{i_j}, \tau_{i_j+1}]} |W(\mathcal{Y}, \tilde{v})| \leq (C + 2\pi)(I_+ + 1)$$

and observe that conditionally, on either $\{X_{\tau_{i_j+1}} = u\}$ or $\{X_{\tau_{i_j+1}} \in \partial(G')^{\#\delta} \cup \mathfrak{s}\}$, I_+ has an exponential tail. The proof of this fact is essentially the same as the second item of [4, Lemma 4.8]. Indeed, we need to show that once the random walk intersects $\tilde{\ell}'_+$ outside B_1 , there is a positive probability (uniform in δ) for the walk to create a non-contractible loop without intersecting $\tilde{\ell}'_-$ (as in the proof of Lemma 5.2). This is intuitively clear, but a complete proof of this needs an input from Riemannian geometry. The issue at hand is that, too close to the boundary, the uniform crossing ceases to hold uniformly in δ . We control the winding of this portion of this walk as follows.

Consider a compact set $A \subset M'$ containing ℓ which forms an approximation of ℓ in the sense that topologically K is an annulus with ℓ being non-contractible in K . Recall from the proof of Lemma 5.2 that the simple random walk has a uniform positive probability to create a non-contractible loop inside K by winding around exactly once, and we stop the simple random walk if this happens. But, in this process, we can assume without loss of generality on ℓ that the lift of the walk stays inside at most four consecutive copies of K (where these four copies are mapped to each other by Möbius transforms). So, if the walk is on $\tilde{\ell}'_+$ at a copy of \tilde{A} which is more than four copies away from \tilde{v} and $\tilde{\ell}'_-$, the lift of that walk cannot intersect ℓ^- in this process. On the other hand, applying four copies of the corresponding Möbius map to an arbitrary $\tilde{v} \in \tilde{K}$ yields by compactness of \tilde{K} a slightly bigger compact \tilde{K}' . Applying the uniform crossing property for the given choice of δ to \tilde{K}' , we now simply apply the argument of [4, Lemma 4.8]. ■

Acknowledgements. The authors are grateful to A. Dahlqvist and J. Dubédat for a number of useful discussions. Finally, they thank the anonymous referee for their valuable feedback on an earlier version of the manuscript.

Funding. NB’s research is supported by FWF Grant P33083, “Scaling limits in random conformal geometry”. This paper was finished while NB was in residence at the Mathematical Sciences Research Institute in Berkeley, California, USA, during the Spring 2022 semester on *Analysis and Geometry of Random Spaces*, which was

supported by the National Science Foundation under Grant no. DMS-1928930. GR's research is supported by NSERC 50311-57400. BL's research is supported by ANR-18-CE40-0033 "Dimers".

References

- [1] A. Ahn, M. Russkikh, and R. Van Peski, [Lozenge tilings and the Gaussian free field on a cylinder](#). *Comm. Math. Phys.* **396** (2022), no. 3, 1221–1275 Zbl [7620657](#) MR [4507921](#)
- [2] J. Aru, *The geometry of the Gaussian free field combined with SLE processes and the KPZ relation*. Ph.D. thesis, Ecole normale supérieure de Lyon, 2015
- [3] N. Berestycki and L. Haunschmid-Sibitz, [Near-critical dimers and massive SLE](#). 2022, arXiv:[2203.15717](#)
- [4] N. Berestycki, B. Laslier, and G. Ray, [Dimers and imaginary geometry](#). *Ann. Probab.* **48** (2020), no. 1, 1–52 Zbl [1445.60014](#) MR [4079430](#)
- [5] N. Berestycki, B. Laslier, and G. Ray, [Dimers on Riemann surfaces II: Conformal invariance and scaling limits](#). 2022, arXiv:[2207.09875](#)
- [6] N. Berestycki and M. Liu, [Piecwise Temperleyan dimers and a multiple SLE_g](#). 2023, arXiv:[2301.08513](#)
- [7] N. Berestycki and E. Powell, [Gaussian free field, Gaussian multiplicative chaos and Liouville quantum gravity](#). 2015, book in preparation
- [8] R. Bott and L. W. Tu, *Differential forms in algebraic topology*. Grad. Texts in Math. 82, Springer, New York, 1982 Zbl [0496.55001](#) MR [658304](#)
- [9] C. Boutillier and B. de Tilière, [Loop statistics in the toroidal honeycomb dimer model](#). *Ann. Probab.* **37** (2009), no. 5, 1747–1777 Zbl [1179.60065](#) MR [2561433](#)
- [10] D. Cimasoni, [Discrete Dirac operators on Riemann surfaces and Kasteleyn matrices](#). *J. Eur. Math. Soc. (JEMS)* **14** (2012), no. 4, 1209–1244 Zbl [1243.82017](#) MR [2928849](#)
- [11] D. Cimasoni and N. Reshetikhin, [Dimers on surface graphs and spin structures. I](#). *Comm. Math. Phys.* **275** (2007), no. 1, 187–208 Zbl [1135.82006](#) MR [2335773](#)
- [12] D. Cimasoni and N. Reshetikhin, [Dimers on surface graphs and spin structures. II](#). *Comm. Math. Phys.* **281** (2008), no. 2, 445–468 Zbl [1168.82012](#) MR [2410902](#)
- [13] M. Ciucu and C. Krattenthaler, [A factorization theorem for lozenge tilings of a hexagon with triangular holes](#). *Trans. Amer. Math. Soc.* **369** (2017), no. 5, 3655–3672 Zbl [1355.05072](#) MR [3605983](#)
- [14] H. Cohn, R. Kenyon, and J. Propp, [A variational principle for domino tilings](#). *J. Amer. Math. Soc.* **14** (2001), no. 2, 297–346 Zbl [1037.82016](#) MR [1815214](#)
- [15] R. Costa-Santos and B. M. McCoy, [Dimers and the critical Ising model on lattices of genus \$> 1\$](#) . *Nuclear Phys. B* **623** (2002), no. 3, 439–473 Zbl [1097.82518](#) MR [1883057](#)
- [16] F. David, A. Kupiainen, R. Rhodes, and V. Vargas, [Liouville quantum gravity on the Riemann sphere](#). *Comm. Math. Phys.* **342** (2016), no. 3, 869–907 Zbl [1336.83042](#) MR [3465434](#)
- [17] P. Di Francesco, H. Saleur, and J.-B. Zuber, [Relations between the Coulomb gas picture and conformal invariance of two-dimensional critical models](#). *J. Statist. Phys.* **49** (1987), no. 1-2, 57–79 Zbl [0960.82507](#) MR [0923852](#)

- [18] S. Donaldson, *Riemann surfaces*. Oxf. Grad. Texts Math. 22, Oxford University Press, Oxford, 2011 Zbl 1235.30001 MR 2856237
- [19] J. Dubédat, *Dimers and families of Cauchy–Riemann operators I*. *J. Amer. Math. Soc.* **28** (2015), no. 4, 1063–1167 Zbl 1321.82011 MR 3369909
- [20] J. Dubédat and R. Gheissari, *Asymptotics of height change on toroidal Temperleyan dimer models*. *J. Stat. Phys.* **159** (2015), no. 1, 75–100 Zbl 1319.82009 MR 3320952
- [21] C. Guillarmou, A. Kupiainen, R. Rhodes, and V. Vargas, *Segal’s axioms and bootstrap for Liouville theory*. 2021, arXiv:2112.14859
- [22] E. P. Hsu, *Stochastic analysis on manifolds*. Grad. Stud. Math. 38, American Mathematical Society, Providence, RI, 2002 Zbl 0994.58019 MR 1882015
- [23] J. Jost, *Compact Riemann surfaces*. 3rd edn., Universitext, Springer, Berlin, 2006 Zbl 1125.30033 MR 2247485
- [24] A. Kassel and R. Kenyon, *Random curves on surfaces induced from the Laplacian determinant*. *Ann. Probab.* **45** (2017), no. 2, 932–964 Zbl 1377.82037 MR 3630290
- [25] R. Kenyon, *Conformal invariance of domino tiling*. *Ann. Probab.* **28** (2000), no. 2, 759–795 Zbl 1043.52014 MR 1782431
- [26] R. Kenyon, *Dominos and the Gaussian free field*. *Ann. Probab.* **29** (2001), no. 3, 1128–1137 Zbl 1034.82021 MR 1872739
- [27] R. Kenyon and A. Okounkov, *Limit shapes and the complex Burgers equation*. *Acta Math.* **199** (2007), no. 2, 263–302 Zbl 1156.14029 MR 2358053
- [28] R. W. Kenyon, J. G. Propp, and D. B. Wilson, *Trees and matchings*. *Electron. J. Combin.* **7** (2000), article no. 25 Zbl 0939.05066 MR 1756162
- [29] H. Lacoïn, R. Rhodes, and V. Vargas, *Semiclassical limit of Liouville field theory*. *J. Funct. Anal.* **273** (2017), no. 3, 875–916 Zbl 1367.81068 MR 3653942
- [30] G. F. Lawler, O. Schramm, and W. Werner, *Conformal invariance of planar loop-erased random walks and uniform spanning trees*. *Ann. Probab.* **32** (2004), no. 1B, 939–995 Zbl 1126.82011 MR 2044671
- [31] C. Mercat, *Discrete Riemann surfaces*. In *Handbook of Teichmüller theory. Vol. I*, pp. 541–575, IRMA Lect. Math. Theor. Phys. 11, European Mathematical Society, Zürich, 2007 Zbl 1136.30315 MR 2349680
- [32] J. Miller and S. Sheffield, *Imaginary geometry I: Interacting SLEs*. *Probab. Theory Related Fields* **164** (2016), no. 3-4, 553–705 Zbl 1336.60162 MR 3477777
- [33] J. Miller and S. Sheffield, *Imaginary geometry IV: Interior rays, whole-plane reversibility, and space-filling trees*. *Probab. Theory Related Fields* **169** (2017), no. 3-4, 729–869 Zbl 1378.60108 MR 3719057
- [34] L. Petrov, *Asymptotics of uniformly random lozenge tilings of polygons. Gaussian free field*. *Ann. Probab.* **43** (2015), no. 1, 1–43 Zbl 1315.60062 MR 3298467
- [35] M. Russkikh, *Dimers in piecewise Temperleyan domains*. *Comm. Math. Phys.* **359** (2018), no. 1, 189–222 Zbl 1395.82060 MR 3781449
- [36] O. Schramm, *Scaling limits of loop-erased random walks and uniform spanning trees*. *Israel J. Math.* **118** (2000), 221–288 Zbl 0968.60093 MR 1776084
- [37] S. Sheffield, *Gaussian free fields for mathematicians*. *Probab. Theory Related Fields* **139** (2007), no. 3-4, 521–541 Zbl 1132.60072 MR 2322706

[38] W. Sun, Toroidal dimer model and Temperley's bijection. 2016, arXiv:[1603.00690](https://arxiv.org/abs/1603.00690)

Communicated by Adrian Tanasă

Received 27 July 2022.

Nathanaël Berestycki

Fakultät für Mathematik, University of Vienna, Oskar Morgenstern Platz 1, 1090 Wien, Austria; nathanael.berestycki@univie.ac.at

Benoit Laslier

Laboratoire de Probabilités, Statistique et Modélisation, Université Paris Cité and Sorbonne Université, CNRS, 75205 Paris, France; laslier@lpsm.paris

Gourab Ray

University of Victoria, DTB A440, Victoria, BC V8W 2Y2, Canada; gourabray@uvic.ca