



Free boundary dimers: random walk representation and scaling limit

Nathanaël Berestycki¹ · Marcin Lis² · Wei Qian³

Received: 11 May 2021 / Revised: 11 January 2023 / Accepted: 30 March 2023 /

Published online: 16 May 2023

© The Author(s) 2023

Abstract

We study the dimer model on subgraphs of the square lattice in which vertices on a prescribed part of the boundary (the free boundary) are possibly unmatched. Each such unmatched vertex is called a monomer and contributes a fixed multiplicative weight $z > 0$ to the total weight of the configuration. A bijection described by Giuliani et al. (J Stat Phys 163(2):211–238, 2016) relates this model to a standard dimer model but on a non-bipartite graph. The Kasteleyn matrix of this dimer model describes a walk with transition weights that are negative along the free boundary. Yet under certain assumptions, which are in particular satisfied in the infinite volume limit in the upper half-plane, we prove an effective, true random walk representation for the inverse Kasteleyn matrix. In this case we further show that, independently of the value of $z > 0$, the scaling limit of the centered height function is the Gaussian free field with Neumann (or free) boundary conditions. It is the first example of a discrete model where such boundary conditions arise in the continuum scaling limit.

Mathematics Subject Classification 82B20

Contents

1	Introduction	736
1.1	Free boundary dimers	736
1.2	Boundary conditions	738
1.3	Statement of main results	738
1.4	Heuristics: reflection and even/odd decomposition	743
1.5	A conjecture on the boundary-touching level lines	745

✉ Marcin Lis
marcin.lis@tuwien.ac.at

¹ Universität Wien, Vienna, Austria

² Technische Universität Wien, Vienna, Austria

³ City University of Hong Kong, Kowloon Tong, Hong Kong

1.6	Folding the dimer model onto itself	746
1.7	Connection with isoradial random walk with critical weights	747
1.8	Outline of the paper and structure of proof	747
2	(Inverse) Kasteleyn matrix	752
2.1	Dimer representation	752
2.2	Kasteleyn orientation, Kasteleyn matrix and gauge change	752
2.3	Towards the inverse Kasteleyn matrix	753
2.4	An intermediate limit	755
2.5	An auxiliary walk on \mathbb{Z}	757
2.6	Random walk representation of D^{-1}	760
2.7	Convergence to potential kernel of the auxiliary walk	766
3	Infinite volume limit	767
3.1	Construction of the potential kernel in the infinite volume setting	768
3.2	Coordinatewise mirror coupling	771
3.3	Suitability of coupling (proof of (3.2))	774
3.4	A priori estimate on the gradient of the potential kernel	777
3.5	Convergence of Green's function differences to gradient of potential kernel	778
4	Scaling limit of discrete derivative of potential kernel	782
4.1	Proof of Proposition 4.3	785
4.2	Proof of Proposition 4.4	790
5	Convergence to the Neumann Gaussian Free field	793
5.1	Infinite volume coupling function and its scaling limit	793
5.2	Pfaffians and Kasteleyn theory	797
5.3	Matchings and permutations	799
5.4	Moments of the height function	801
5.5	Convergence of the height function	807
	References	811

1 Introduction

1.1 Free boundary dimers

Let $\mathcal{G} = (V, E)$ be a finite, connected, planar bipartite graph (in our analysis we will actually only consider subgraphs of the square lattice \mathbb{Z}^2). Let $\partial\mathcal{G}$ be the set of boundary vertices, i.e., vertices adjacent to the unique unbounded external face, and let $\partial_{\text{free}}\mathcal{G} \subseteq \partial\mathcal{G}$ be a fixed set called the **free boundary**. A **boundary monomer-dimer cover** of \mathcal{G} is a set $M \subseteq E$ such that

- each vertex in $V \setminus \partial_{\text{free}}\mathcal{G}$ belongs to *exactly* one edge in M ,
- each vertex in $\partial_{\text{free}}\mathcal{G}$ belongs to *at most* one edge in M .

We write $\text{mon}(M) \subseteq \partial_{\text{free}}\mathcal{G}$ for the set of vertices that do not belong to any edge in M , and call its elements **monomers**. Let $\mathcal{MD}(\mathcal{G})$ be the set of all boundary monomer-dimer covers of \mathcal{G} . We will often call such configurations simply monomer-dimer covers, keeping in mind that monomers are only allowed on the free boundary. Finally let $\mathcal{D}(\mathcal{G})$ be the set of all **dimer covers**, i.e. monomer-dimer covers M such that $\text{mon}(M) = \emptyset$.

We assign to each edge $e \in E$ a weight $w_e \geq 0$, and to each vertex $v \in \partial_{\text{free}}\mathcal{G}$ a weight $z_v \geq 0$. The **dimer model with a free boundary** (or **free boundary dimer model**) is a random choice of a boundary monomer-dimer cover from $\mathcal{MD}(\mathcal{G})$ according to the following probability measure:

$$\mathbf{P}(M) = \frac{1}{\mathcal{Z}} \prod_{e \in M} w_e \prod_{v \in \text{mon}(M)} z_v, \tag{1.1}$$

where \mathcal{Z} is the normalizing constant called the **partition function**. For convenience we will always assume that the partition function $\mathcal{Z} > 0$, i.e., that $\mathcal{MD}(\mathcal{G}) \neq \emptyset$. In this work we will only focus on the homogeneous case $w_e = 1$, for all $e \in E$, and $z_v = z > 0$ for all $v \in \partial_{\text{free}}\mathcal{G}$, with the exception of the technical assumption on the weight of corner monomers described in the next section. If we did not make the exception for the corner monomers, i.e., if $z_v = z$ for all $v \in \partial_{\text{free}}\mathcal{G}$, then

$$\mathbf{P}(M) = \frac{1}{\mathcal{Z}} z^{|\text{mon}(M)|}. \tag{1.2}$$

The **dimer model** on \mathcal{G} can be now defined as the free boundary dimer model conditioned on $\mathcal{D}(\mathcal{G})$, i.e., the event that there are no monomers.

The main observable of interest for us will be the **height function** of a boundary monomer-dimer cover which is an integer-valued function defined (up to a constant) on the bounded faces of \mathcal{G} . Its definition is identical to the one in the dimer model (see [39]). We simply note that the presence of monomers on the boundary does not lead to any topological complication (i.e., the height function is not multivalued): if u and u' are two faces of the graph, and γ and γ' are two distinct paths in the dual graph connecting u and u' , the loop formed by connecting γ and γ' (in the reverse direction) does not enclose any monomer. More precisely, we view a configuration $M \in \mathcal{MD}(\mathcal{G})$ as an antisymmetric flow (in other words a 1-form) ω_M on the directed edges of \mathcal{G} in the following manner: if $e = \{w, b\} \in M$, then $\omega_M(w, b) = 1$ and $\omega_M(b, w) = -1$ where b is the black vertex of e and w its white vertex (since \mathcal{G} is bipartite, a choice of black and white vertices can be made in advance). Otherwise, we set $\omega_M(e) = 0$. Equivalently, we may view ω_M as an antisymmetric flow on the directed dual edges, where if e^\dagger is the dual edge of e (obtained by a counterclockwise $\pi/2$ rotation of e), then $\omega_M(e^\dagger) = \omega_M(e)$. To define the height function we still need to fix a **reference flow** ω_0 which we define to be $\omega_0 = \mathbf{E}[\omega_M]$, i.e., the expected flow of M under the free boundary dimer measure. Now, if u and u' are two distinct (bounded) faces of \mathcal{G} , we simply define

$$h(u) - h(u') = \sum_{e^\dagger \in \gamma} (\omega_M(e^\dagger) - \omega_0(e^\dagger))$$

where γ is any path (of dual edges) connecting u to u' . This definition does not depend on the choice of the path since the flow $\omega_M(e^\dagger) - \omega_0(e^\dagger)$ is closed (sums over closed dual paths vanish), and hence yields a function h up to an additive constant, as desired. Note that our choice of the reference flow automatically guarantees that the height function is centered, i.e., $\mathbf{E}(h(u) - h(u')) = 0$ for all faces u and u' .

We finish this short introduction to the free boundary dimer model with a few words on its history and the nomenclature. In the original model studied in [18, 19] monomers could occupy any vertex of the graph, and hence the name **monomer-dimer model**. This generalization poses two major complications from our point of view. First of all,

the height function is not well defined, and secondly the model does not admit a Kasteleyn solution as was shown in [20]. From this point of view, it would therefore be natural if the version of the model studied here was called the boundary-monomer-dimer model. However we choose to use the less cumbersome name of free boundary dimers.

1.2 Boundary conditions

We now state conditions on the graph $\mathcal{G} = (V, E)$ which will be enforced throughout this paper. First, we assume that \mathcal{G} is a subgraph of the square lattice \mathbb{Z}^2 , and without loss of generality that $0 \in V$ and is a black vertex. This fixes a unique black/white bipartite partition of V . We also assume that

- V is contained in the upper half plane $\mathbb{H} = \{z \in \mathbb{C} : \Im(z) \geq 0\}$.
- $\partial_{\text{free}}\mathcal{G} = V \cap \mathbb{R} (\neq \emptyset)$, so the monomers are allowed only on the real line. Furthermore, we assume $\partial_{\text{free}}\mathcal{G}$ is a connected set (interval) of an even number of vertices. The leftmost and rightmost vertices of $V \cap \mathbb{R} = \partial_{\text{free}}\mathcal{G}$ will be referred to as the **monomer-corners** of \mathcal{G} .
- \mathcal{G} has at least one black dimer-corner and one white dimer-corner (where a **dimer-corner** is a vertex $v \in V$ that is not a monomer-corner, and is adjacent to the outer face of \mathcal{G} , and has degree either 2 or 4 in \mathcal{G}).

See Fig. 3 for an example of a domain satisfying these assumptions (ignore the bottom row of triangles for now, which will be described later). We make a few comments on the role of the last assumption that there are corners of both colours. For this it is useful to make a parallel with Kenyon's definition of **Temperleyan domain** [22, 23]. In that case, this condition ensured that the associated random walk on one of the four possible sublattices of \mathbb{Z}^2 (the two types of black and the two types of white vertices) was killed somewhere on the boundary. As we will see, in our case the random walk may change the lattice from black to white when it is near the real line, resulting in only two different types of walks. Then the role of the third assumption (at least one dimer-corner of each type) is to ensure that each of the two walks is killed on at least some portion of the boundary (possibly a single vertex). This follows from an observation that the boundary condition of a walk on a black (resp. white) sublattice changes from Neumann to Dirichlet (and vice-versa) at a white (resp. black) corner. See Fig. 5 for an example of a vertex with Neumann and Dirichlet boundary conditions.

1.3 Statement of main results

The free boundary dimer model as defined above was discussed in a paper of Giuliani et al. [16]. It was shown there that the partition function \mathcal{Z} can be computed as a Pfaffian of a certain (signed) adjacency matrix. Furthermore, a bijection was provided to a non-bipartite dimer model (the authors indicate that this bijection was suggested by an anonymous referee). Hence using Kasteleyn theory the correlation functions can be expressed as Pfaffians of the inverse Kasteleyn matrix K^{-1} . The bijection, which is a central tool of our analysis, will be defined in Sect. 2 where we will also recall the precise definition of the Kasteleyn matrix K .

We will now state our first main result which gives a full random walk representation for K^{-1} . Suppose that \mathcal{G} is a graph satisfying the assumptions from the previous

section. Fix $z > 0$ and assign weight z to every monomer on $\partial_{\text{free}}\mathcal{G}$ except at either monomer-corner, where (for technical reasons which will become clear in the proof) we choose the weight to be

$$z' = \frac{z}{2} + \sqrt{1 + \frac{z^2}{4}}. \tag{1.3}$$

Thus formally the free boundary dimer model we consider is of type (1.1) rather than of type (1.2). For $k \in \mathbb{N} = \{0, 1, \dots\}$, let us call $V_k = V_k(\mathcal{G}) = \{v \in V : \Im(v) = k\}$, so $\partial_{\text{free}}\mathcal{G} = V_0$, where $\Im(v)$ denotes the imaginary part of the vertex v seen as a complex number given by the embedding of the graph. Let us call $V_{\text{even}} = V_{\text{even}}(\mathcal{G}) = V_0 \cup V_2 \cup \dots$ and $V_{\text{odd}} = V_{\text{odd}}(\mathcal{G}) = V_1 \cup V_3 \cup \dots$.

Theorem 1.1 (Random walk representation of the inverse Kasteleyn matrix) *There exist two random walks Z_{even} and Z_{odd} on the state spaces $V_{\text{even}}(\mathcal{G})$ and $V_{\text{odd}}(\mathcal{G})$ respectively, whose transition probabilities will be described in Sect. 2.6 (see (2.26) and (2.30)), such that the following holds. Consider the monomer-dimer model on \mathcal{G} where the monomer weight is $z > 0$ on $V_0(\mathcal{G})$ except at its monomer-corners where the monomer weight is z' , as defined in (1.3). Let K be the associated Kasteleyn matrix, and $D = K^*K$, so that $K^{-1} = D^{-1}K^*$. Then for all $u, v \in V$, we have*

$$D^{-1}(u, v) = \begin{cases} G_{\text{odd}}(u, v) & \text{if } u, v \in V_{\text{odd}}, \\ (-1)^{\Re(u-v)} G_{\text{even}}(u, v) & \text{if } u, v \in V_{\text{even}}, \\ 0 & \text{otherwise.} \end{cases} \tag{1.4}$$

where $G_{\text{even}}, G_{\text{odd}}$ are the Green's functions of Z_{even} and Z_{odd} respectively, normalised by $D(v, v)$.

Here, by normalised Green's function of a random walk (with at least one absorbing state), we mean

$$G(u, v) = \frac{1}{D(v, v)} \mathbb{E}_u \left(\sum_{k=0}^{\infty} \mathbf{1}_{\{Z_k=v\}} \right),$$

where Z is the corresponding random walk. We now specify a few properties of the random walks Z_{even} and Z_{odd} which may be interesting to the reader already, even though the exact definition is postponed until Sect. 2.6. Both Z_{even} and Z_{odd} behave like simple random walk away (at distance more than 2) from the boundary vertices, but with jumps of size ± 2 , so the parity of the walk does not change. Both have nontrivial boundary conditions, including some reflecting and absorbing boundary arcs along the non-monomer part of the boundary $\partial\mathcal{G} \setminus \partial_{\text{free}}\mathcal{G}$. Furthermore, both walks are allowed to make additional jumps (i.e., not necessarily equal to ± 2) along their bottommost rows of vertices (V_0 for Z_{ev} and V_1 for Z_{odd}). These jumps are symmetric, bounded in the even case but not in the odd case (although they do have exponentially decaying tail). Hence in the scaling limit, these walks would converge to Brownian motion in the upper half plane \mathbb{H} with reflecting (or equivalently Neumann) boundary

conditions on the real axis, and with whatever boundary conditions are inherited from the Neumann/Dirichlet parts of the other boundary arcs.

An important property of these random walks that highlights the difference with the setup of [22], is that they can change colour of the vertex (in a bipartite coloring of $\mathbb{H} \cap \mathbb{Z}^2$). However, this can happen only when the walker visits the real line. This in turn means that the entries of the inverse Kasteleyn matrix indexed by two vertices of the same colour (which automatically vanish in Kenyon’s work) have a natural interpretation in terms of walks that go through the real line (the free boundary). This is a clear analogy with the construction of reflected random walks via the reflection principle for a walk in a reflected domain. Remarkably, this exact correspondence with reflected random walks is present already at the discrete level of the dimer model with free boundary conditions, and is the reason why the reflected Brownian motion appears in the correlation kernel of the scaling limit of the height function.

To illustrate this we explain here briefly a simple computation using Kasteleyn theory (for more details see Sect. 5.2) where this phenomenon is apparent. Let $e = \{w, b\}$ and $e' = \{w', b'\}$ be two edges of $\mathbb{Z}^2 \cap \mathbb{H}$ with w, w' white and b, b' black vertices in a fixed chessboard coloring of the lattice. Then, writing \mathcal{M} for a random boundary monomer-dimer cover, and using Kasteleyn theory for the dimer representation described in Sect. 2.1, we have

$$\begin{aligned} \mathbf{P}(e, e' \in \mathcal{M}) &= a \text{Pf} \begin{pmatrix} 0 & K^{-1}(w, b) & K^{-1}(w, w') & K^{-1}(w, b') \\ 0 & & K^{-1}(b, w') & K^{-1}(b, b') \\ & & 0 & K^{-1}(w', b') \\ & & & 0 \end{pmatrix} \\ &= a(K^{-1}(w, b)K^{-1}(w', b') + K^{-1}(b, w')K^{-1}(w, b') \\ &\quad - K^{-1}(w, w')K^{-1}(b, b')), \end{aligned}$$

where the matrix is antisymmetric and $a = K(w, b)K(w', b')$. We also have $\mathbf{P}(e \in \mathcal{M}) = K(w, b)K^{-1}(w, b)$ and $\mathbf{P}(e' \in \mathcal{M}) = K(w', b')K^{-1}(w', b')$, which leads to

$$\text{Cov}(\mathbf{1}_{e \in \mathcal{M}}, \mathbf{1}_{e' \in \mathcal{M}}) = a(K^{-1}(b, w')K^{-1}(w, b') - K^{-1}(w, w')K^{-1}(b, b')).$$

Here, the second term is new compared to Kenyon’s computation in [22] (in that case one has $K^{-1}(w, w') = K^{-1}(b, b') = 0$). Furthermore, using our random walk representation, $K^{-1}(w, w')$ and $K^{-1}(b, b')$ can be interpreted as a derivative of the Green’s function of the appropriate walks Z_{even} and Z_{odd} evaluated at pairs of vertices of different colors. Then by construction the walks which contribute to these Green’s functions must visit the boundary.

This intuition is what guides us to the next result, which however requires us to first take an infinite volume (thermodynamic) limit where an increasing sequence of graphs eventually covers $\mathbb{H} \cap \mathbb{Z}^2$. We first show that the monomer-dimer model converges in such a limit. For this we need to specify a topology: we view a monomer-dimer configuration on $\mathbb{H} \cap \mathbb{Z}^2$ as an element of $\{0, 1\}^{E(\mathbb{H})}$ where $E(\mathbb{H})$ is the edge set of $\mathbb{Z}^2 \cap \mathbb{H}$, and equip this space with the product topology (so convergence in this space corresponds to convergence of local observables).

To state the result we will fix a sequence \mathcal{G}_n of graphs such that \mathcal{G}_n satisfies the assumptions of Sect. 1.2, and moreover $\mathcal{G}_n \uparrow \mathbb{Z}^2 \cap \mathbb{H}$. For simplicity of the arguments and ease of presentation, we have chosen \mathcal{G}_n to be a concrete approximation of rectangles, although the result is in fact true much more generally; we have not tried to find the most general setting in which this applies.

Theorem 1.2 (Infinite volume limit) *Let \mathcal{G}_n be rectangles of diverging even sidelengths (number of vertices on a side) whose ratio remains bounded away from zero and infinity as $n \rightarrow \infty$, and such that in the top row the right-hand side half of the vertices are removed (if this number is even, otherwise we remove one less vertex to keep the graph dimerable). Let μ_n denote the law of the free boundary dimer model on \mathcal{G}_n with monomer weight $z > 0$ except at the monomer-corners where the weight is z' , as in (1.3). Then μ_n converges weakly as $n \rightarrow \infty$ to a law μ which describes a.s. a random boundary monomer-dimer configuration on $\mathbb{Z}^2 \cap \mathbb{H}$.*

We note that the particular type of domains chosen in this statement (see e.g. Fig. 3 for an illustration) guarantees that both the odd and even walks mentioned above are killed on a macroscopic part of the upper rows of \mathcal{G}_n (the odd walk is killed on the right-hand side half and the even walk on the left-hand side half of its uppermost row); importantly this killing can occur with positive probability without touching the real line starting from any position away from the real line. The other key requirement is that the domain grows to infinity in a “homogeneous” way: for instance, \mathcal{G}_n both contains a ball of radius $c_1 n$ and is contained in a ball of radius $c_2 n$ for some suitable constants $c_1, c_2 > 0$. Subject to these two requirements (both walks may be killed with positive probability without touching the real line once they are far away from it, and the growth is homogeneous in all directions), there should be no difficulty extending the result in Theorem 1.2; see immediately above Proposition 3.7 for the kind of domains which are explicitly allowed in the proof. For concreteness and simplicity we opted to choose the sequence of domains as in the statement: removing (approximately) half vertices in the top row makes it possible for walks of all types to be killed without touching the real line.

We stress the fact that the limiting law μ depends on the monomer weight $z > 0$. As mentioned before, we can associate to the monomer-dimer configuration in the infinite half-plane a height function which is defined on the faces of $\mathbb{H} \cap \mathbb{Z}^2$, up to a global additive constant. The last main result of this paper shows that in the scaling limit, this (centered) height function converges to a **Gaussian free field with Neumann** (or free) **boundary conditions**, denoted by Φ^{Neu} . We will not define this in complete generality here (see [8] for a comprehensive treatment). We will simply point out what is concretely relevant for the theorem below to make sense. Given a simply connected domain Ω with a smooth boundary, $\Phi_{\Omega}^{\text{Neu}}$ may be viewed as a stochastic process indexed by the space $\mathcal{D}_0(\Omega)$ of smooth test functions $f : \Omega \rightarrow \mathbb{R}$ with compact support and with zero average (meaning $\int_{\mathbb{H}} f(z) dz = 0$). The latter requirement corresponds to the fact that Φ is only defined modulo a global additive constant. The law of this stochastic process is characterised by a requirement of linearity (i.e. $(\Phi_{\Omega}^{\text{Neu}}, af + bg) = a(\Phi_{\Omega}^{\text{Neu}}, f) + b(\Phi_{\Omega}^{\text{Neu}}, g)$ a.s. for any $f, g \in \mathcal{D}_0(\Omega)$ and $a, b \in \mathbb{R}$), and moreover $(\Phi_{\Omega}^{\text{Neu}}, f), (\Phi_{\Omega}^{\text{Neu}}, g)$ follow centered Gaussian distributions with

covariance

$$\text{Cov}((\Phi_{\Omega}^{\text{Neu}}, f), (\Phi_{\Omega}^{\text{Neu}}, g)) = \iint_{\Omega^2} f(x)g(y)G_{\Omega}^{\text{Neu}}(x, y)dx dy,$$

where $G_{\Omega}^{\text{Neu}}(x, y)$ is a Green’s function in Ω with Neumann boundary conditions. (Note that by contrast to the Dirichlet case, such Green’s functions are not unique and are defined only up to a constant.) In the case of the upper-half plane $\Omega = \mathbb{H}$, the Green’s function is given explicitly by

$$G_{\mathbb{H}}^{\text{Neu}}(x, y) = -\log|x - y| - \log|x - \bar{y}|.$$

Informally, pointwise differences $\Phi_{\mathbb{H}}^{\text{Neu}}(a) - \Phi_{\mathbb{H}}^{\text{Neu}}(b)$ for $a, b \in \mathbb{H}$ (which do not depend on the choice of the global additive constant) are centered Gaussian random variables with covariances

$$\begin{aligned} &\mathbb{E}[(\Phi_{\mathbb{H}}^{\text{Neu}}(a_i) - \Phi_{\mathbb{H}}^{\text{Neu}}(b_i))(\Phi_{\mathbb{H}}^{\text{Neu}}(a_j) - \Phi_{\mathbb{H}}^{\text{Neu}}(b_j))] \\ &= -\log \left| \frac{(a_i - a_j)(b_i - b_j)(\bar{a}_i - a_j)(\bar{b}_i - b_j)}{(a_i - b_j)(b_i - a_j)(\bar{a}_i - b_j)(\bar{b}_i - a_j)} \right|. \end{aligned} \tag{1.5}$$

Note that our Green’s function is normalised so that it behaves like $1 \times \log(1/|x - y|)$ as $y - x \rightarrow 0$. Naturally, (1.5) must be understood in an integrated way since pointwise differences are not actually defined.

We may now state the announced result. For $\delta > 0$ (the mesh size), let h^{δ} denote the height function (defined up to a constant, and by definition centered) of the free boundary dimer model μ with weight z in the infinite half-plane $\mathbb{H} \cap \delta\mathbb{Z}^2$ (rescaled by δ). We identify h^{δ} with a function defined almost everywhere on \mathbb{H} by taking the value of h^{δ} to be constant on each face, and view h^{δ} as a random distribution (also called a random generalized function) acting on smooth compactly supported functions f on \mathbb{H} with zero average, i.e., satisfying $\int_{\mathbb{H}} f(a)da = 0$ (see Sect. 5.5 for details).

Theorem 1.3 (Scaling limit) *Let $f_1, \dots, f_k \in \mathcal{D}_0(\mathbb{H})$ be arbitrary test functions. Then for all $z > 0$, as $\delta \rightarrow 0$,*

$$(h^{\delta}, f_i)_{i=1}^k \rightarrow \left(\frac{1}{\sqrt{2\pi}} \Phi_{\mathbb{H}}^{\text{Neu}}, f_i \right)_{i=1}^k$$

in distribution.

Note that, maybe surprisingly, the scaling limit does not depend on the value of $z > 0$ (we discuss this in more detail in Sect. 1.4). We also wish to call the attention of the reader to the normalising factor $1/(\sqrt{2\pi})$ in front of Φ on the right-hand side of Theorem 1.3. It is equal to the one appearing in the usual dimer model in which the centered height function has zero (Dirichlet) boundary conditions. We note that comparisons with other works such as [6, 23] should be made carefully, since the normalisation of the Green’s function and of the height function may not be the same:

for instance, Kenyon takes the Green’s function to be normalised so that $G(x, y) \sim 1/(2\pi) \log 1/|x - y|$ as $y \rightarrow x$, so his GFF is $1/\sqrt{2\pi}$ ours (ignoring different boundary conditions). Also, in Kenyon’s work [22], the height function is such that the total flow out of a vertex is 4 instead of 1 here (so his height function is 4 times ours), while it is 2π in [6] (so their height function is 2π times ours). Adjusting for these differences, there is no discrepancy between the constant $1/(\sqrt{2\pi})$ on the right-hand side of Theorem 1.3 and the one in [22, 23].

1.4 Heuristics: reflection and even/odd decomposition

As noted before, Theorem 1.3 may be surprising at first sight, when we consider the behaviour of the model in the two extreme cases $z = 0$ and $z = \infty$. Indeed, when $z = 0$, the free boundary dimer model obviously reduces to the dimer model on \mathbb{H} , in which case the limit is a Dirichlet GFF. When $z = \infty$, all vertices of V_0 are monomers, so the model reduces to a dimer model on $(V_1 \cup V_2 \cup \dots) \simeq \mathbb{H} \cap \mathbb{Z}^2$. Hence, the limit is also a Dirichlet GFF in this case. However, the result above says that for any z strictly in between these two extremes, the limit is a Neumann GFF.

The result (and the reason for this arguably surprising behaviour) may be heuristically understood through the following **reflection argument**. Let \mathcal{G} be a large finite graph approximating \mathbb{H} and satisfying the assumptions of Sect. 1.2. Let $\tilde{\mathcal{G}}$ be a copy of \mathcal{G} shifted by $i/2$, so with a small abuse of notation, $\tilde{\mathcal{G}} = \mathcal{G} + i/2$ (here $i = \sqrt{-1}$), and let $\bar{\mathcal{G}}$ be the same graph to which we add its conjugate (reflection through the real axis). We also add vertical edges crossing the real axis of the form $(k - i/2, k + i/2)$ for each $k \in V_0$; note that the resulting graph is then clearly bipartite. Given a monomer-dimer configuration on \mathcal{G} , we can readily associate a monomer-dimer configuration on $\bar{\mathcal{G}}$ by reflecting it in the same manner. In this way, a monomer in $k + i/2$ necessarily sits across another monomer in $k - i/2$ for any $k \in V_0$. Such a pair of monomers can be interpreted as a dimer on the edge $(k - i/2, k + i/2)$ and once we have phrased it this way the resulting configuration is just an *ordinary dimer configuration* on $\bar{\mathcal{G}}$ (which however has the property that it is reflection symmetric). It follows that its height function (defined on the faces of $\bar{\mathcal{G}}$) is even, i.e., $h(f) = h(\bar{f})$ for every face f (where \bar{f} is the symmetric image of f about the real axis). Moreover, a moment of thought shows that monomer-dimer configurations on \mathcal{G} are in bijection in this manner with the set $\mathcal{ED}(\bar{\mathcal{G}})$ of even (symmetric) dimer configurations on $\bar{\mathcal{G}}$, and that under this bijection the image of the law (1.2) is given by

$$\mathbb{P}(M) = \frac{1}{\bar{Z}} z^{|\text{mon}(M)|} \tag{1.6}$$

(where for a dimer configuration $M \in \mathcal{D}(\bar{\mathcal{G}})$, $\text{mon}(M)$ is the set of vertical edges of M crossing the real axis), conditioned on the event $\mathcal{ED}(\bar{\mathcal{G}})$ of being even, where \bar{Z} is the partition function of the dimer model on $\bar{\mathcal{G}}$.

Now, suppose e.g. that \mathcal{G} is such that $\bar{\mathcal{G}}$ is piecewise Temperleyan [36] (meaning that $\bar{\mathcal{G}}$ has two more white convex corners than white concave corners, see [36] for precise definitions). This happens for instance if \mathcal{G} is a large rectangle with appropriate

dimensions. By a result of Russkikh [36], in this case and if $z = 1$, the unconditional (centered) height function associated with the dimer model (1.6) converges to a Gaussian free field with Dirichlet boundary condition in the scaling limit.

It is reasonable to believe that this convergence holds true even when $z \neq 1$. For instance, when the monomer weights alternate between z and 1 every second vertex, then whatever the value of z , the dimer model has a Temperleyan representation (see [5, 27]): indeed, in that case the weighted graph is obtained as a superposition of a planar graph and its dual for which the dual edges all have weight one.

Then by considerations related to the imaginary geometry approach (see [6]), this convergence to the Dirichlet GFF is universal provided that the underlying random walk converges to Brownian motion (this will be rigorously proved in the forthcoming work [7]). In particular, given these results, we should get convergence to the Dirichlet GFF for the height function even when $z \neq 1$: indeed, when we modify the weight of all the edges crossing the real line, random walk will still converge to Brownian motion. So far, this discussion concerned the (unconditioned) dimer model on $\tilde{\mathcal{G}}$ defined in (1.6). Once we start conditioning on $\mathcal{ED}(\tilde{\mathcal{G}})$ it might be natural to expect that the scaling limit should be a ‘‘Dirichlet GFF conditioned to be even’’, though this is a highly degenerate conditioning. Nevertheless, this conditioning makes sense in the continuum, and in fact its restriction to the upper half plane gives the Neumann GFF, as we are about to argue. Indeed, for a full plane GFF $\Phi_{\mathbb{C}}$ restricted to \mathbb{H} , it is easy to check that one has the decomposition

$$\Phi_{\mathbb{C}} = \frac{1}{\sqrt{2}}(\Phi_{\mathbb{H}}^{\text{Neu}} + \Phi_{\mathbb{H}}^{\text{Dir}}) \tag{1.7}$$

where $\Phi_{\mathbb{H}}^{\text{Neu}}$, $\Phi_{\mathbb{H}}^{\text{Dir}}$ are independent fields on \mathbb{H} with Neumann and Dirichlet boundary conditions on \mathbb{R} respectively. This follows immediately from the fact that any test function can be written as the sum of an even and odd functions, and this decomposition is orthogonal for the Dirichlet inner product $(\cdot, \cdot)_{\nabla}$ on $\mathcal{D}_0(\mathbb{C})$. Therefore, conditioning $\Phi_{\mathbb{C}}$ to be even amounts to conditioning on $\Phi_{\mathbb{H}}^{\text{Dir}}$ to vanish everywhere, meaning that $\Phi_{\mathbb{C}}$ (restricted to the upper half plane) is exactly equal to $\Phi_{\mathbb{H}}^{\text{Neu}}/\sqrt{2}$. (See Exercise 1 of Chapter 5 in [8] for details.)

We note that while this argument correctly predicts the Neumann GFF as a scaling limit of the height function, it is however also somewhat misleading as it suggests that the limit of h^δ is not $(1/\sqrt{2\pi})\Phi_{\mathbb{H}}^{\text{Neu}}$ as in Theorem 1.3, but is smaller by a factor $1/\sqrt{2}$, i.e., $1/(2\pi)\Phi_{\mathbb{H}}^{\text{Neu}}$.

To understand this discrepancy, we now explain why the additional factor turns out to be an artifact of a Gaussian computation and does not arise in the discrete setup. A convincing one-dimensional parallel can be that of Gaussian and simple random walk bridges. Indeed, consider bridges of $2n$ steps starting and ending at 0, with symmetric Bernoulli and Gaussian jump distributions with variance one. Now condition the walks to be symmetric around time n , i.e. $X(n \pm k) = X(n \mp k)$. Again, the Gaussian conditioning is singular but can be easily made sense of using Gaussian integrals. Restricted to the time interval $[0, n]$, the conditioned simple random walk bridge is just a simple random walk with the same step distribution as the original bridge. However, the conditioned Gaussian walk has step distribution with variance

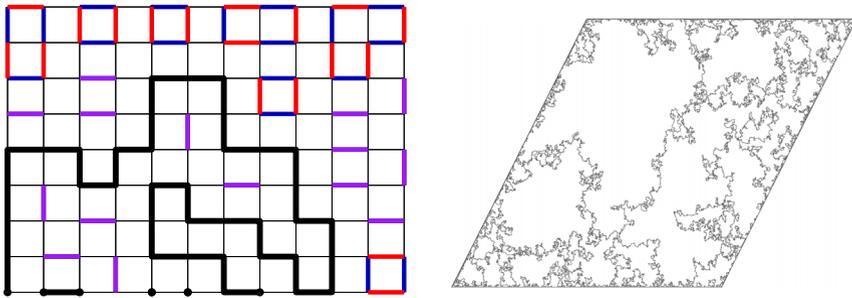


Fig. 1 Left: A superposition of two monomer-dimer configurations, respectively blue and red. Double edges are in purple. The boundary-touching level lines of the height-function is the collection of arcs joining monomers to monomers marked in bald black. Right: A simulation of ALE by B. Werness (color figure online)

$1/2$ as a result of the conditioning. In particular, in the diffusive scaling limit, the former walk converges to standard Brownian motion whereas the latter to $1/\sqrt{2}$ times the standard Brownian motion. The framework of the current paper is more similar to the simple random walk case as discrete height functions are its “two-dimensional-time” analogs. This concludes the discussion giving the heuristics for Theorem 1.3.

1.5 A conjecture on the boundary-touching level lines

In the study of the dimer model, a well known conjecture of Kenyon concerns the superposition of two independent dimer configurations. It is easy to check that such a superposition results in a collection of loops (including double edges) covering every vertex. This observation is attributed to Percus [33]. These loops are the level lines of the difference of the two corresponding dimer height functions. Kenyon’s conjecture (stated somewhat informally in [25] for instance) is that the loops converge in the scaling limit to CLE_4 , the conformal loop ensemble with parameter $\kappa = 4$ (defined in [37], see also [38]). This is strongly supported by the fact that in the continuum, CLE_4 can be viewed as the level lines of a (Dirichlet) GFF with a specified variance (a consequence of a well known coupling between the GFF and CLE_4 of Miller and Sheffield, see [1] for a complete statement and proof). Major progress has been made recently on this conjecture through the work of Kenyon [25], Dubédat [12] and Basok and Chelkak [4], and the only remaining ingredient of the full proof is to show precompactness of the family of loops in a suitable metric space. *Updated during the revision process:* in combination with the work of Bai and Wan [2] the above results now show that the so-called cylindrical probabilities of the double dimer model converge to those of CLE_4 . This gives the proof of a weak form of Kenyon’s conjecture.

It is natural to ask if any similar phenomenon occurs when we superpose two independent monomer-dimer configurations sampled according to the free boundary dimer model, say in the upper half-plane. For topological reasons, this gives rise to a gas of loops as above but also a collection of curves connecting monomers to monomers (and hence the real line to the real line). See Fig. 1 for an example. In fact, the

superposition of two free boundary dimer configurations is related to the superposition of two dimer configurations with different boundary conditions considered in [26], with the difference being that in that model the monomers occur in fixed locations. We note that these authors already establish connections between their results and multiple-strand SLE.

An obvious question is to describe the law of this collection of curves in the scaling limit. By analogy with the above, and in view of our result (Theorem 1.3), it is natural to expect that these curves converge in the scaling limit to the level lines of a GFF with Neumann boundary conditions on the upper-half plane. The law of these curves was determined by Qian and Werner [35] to be the ALE process (ALE stands for Arc Loop Ensemble. It is a collection of arcs that can be connected into loops, but here we will not be interested in this aspect and will only see them as arcs.). ALE is one possible name for this set, but more precisely it is equal to the branching $SLE_4(-1, -1)$ exploration tree targeting all boundary points, and is also equal to the (gasket of) $BCLE_4(-1)$ in [32] and $\mathbb{A}_{-\lambda, \lambda}$ in [1].

This leads us to the following conjecture:

Conjecture 1.4 *For any $z > 0$, in the scaling limit, the collection of boundary-touching curves resulting from superimposing two independent free boundary dimer models converges to the **Arc Loop Ensemble** ALE in the upper half-plane.*

1.6 Folding the dimer model onto itself

The discussion in Sects. 1.4 and 1.5 lead naturally to another conjecture which we now spell out. In Sect. 1.5 we explained a conjecture pertaining to the superposition of two *independent* monomer-dimer configurations sampled according to the free boundary dimer model. But there is at least one other natural way to superpose two such configurations that are not independent: namely, when they come from the *same* full plane dimer model. In fact, there are two ways to do the folding, depending on whether we shift by $i/2$ or not.

Let us explain this more precisely. Let us define the graph $\hat{\mathcal{G}}$ which is obtained by adding to \mathcal{G} its reflection with respect to the real axis. The vertices of \mathcal{G} on the real axis (i.e., V_0) are not reflected: we only keep one copy of them in $\hat{\mathcal{G}}$. (By contrast, in the graph $\tilde{\mathcal{G}}$, where \mathcal{G} is shifted by $i/2$ prior to reflection, these vertices are duplicated).

Now, consider an (infinite volume) dimer cover M on $\hat{\mathcal{G}}$, viewed as a subset of edges where every vertex has degree 1, and consider the superposition $\hat{\Sigma}$ obtained by superposing M with itself via a reflection through the real line: thus,

$$\hat{\Sigma} = M|_{\mathbb{H}} \cup (-M)|_{-\mathbb{H}}.$$

Then $\hat{\Sigma}$ is a subgraph of degree two (including double edges), except for vertices on $V_0 \subset \mathbb{R}$ which in M are connected to a vertical edge. Thus $\hat{\Sigma}$ is exactly of the same nature as the graph in Fig. 1. It is not hard to see that the “height function” $h_{\hat{\Sigma}}$ (really defined only up to a global additive constant) naturally associated with $\hat{\Sigma}$ converges in the fine mesh size limit to $(1/\pi)\Phi_{\mathbb{H}}^{\text{Neu}}$: this is because at the discrete level, the corresponding height function $h_{\hat{\Sigma}}(f)$ at a face $f \subset \mathbb{H}$ can be viewed as

$h_M(f) + h_M(\bar{f})$ (where h_M is the height function associated with M), and h_M is known to converge to $(1/\sqrt{2\pi})\Phi_{\mathbb{C}}$ [10]. These considerations lead us to the following conjecture:

Conjecture 1.5 *In the scaling limit, the collection of boundary-touching curves in $\hat{\Sigma}$ converges to the Arc Loop Ensemble ALE in the upper half-plane.*

We remark that it is also meaningful to fold a dimer configuration on $\bar{\mathcal{G}}$ (rather than $\hat{\mathcal{G}}$ above) onto itself via reflection through the real line. In that case, one must erase the vertical edges straddling the real line and view the corresponding dimers as pairs of monomers. The resulting superposition $\bar{\Sigma}$ is a subgraph of degree two, including multiple edges and double points (on $V_0 \subset \mathbb{R} + i/2$). In particular there are no boundary arcs in $\bar{\Sigma}$, except for degenerate lines connecting every monomer to itself. For the same reason as above, the height function $h_{\bar{\Sigma}}$ associated to $\bar{\Sigma}$ may be viewed as $h_M(f) - h_M(\bar{f})$ and so converges in the scaling limit towards $(1/\pi)\Phi_{\mathbb{H}}^{\text{Dir}}$. Analogously to Conjecture 1.5, we conjecture that the loops of $\bar{\Sigma}$ converge to CLE_4 .

1.7 Connection with isoradial random walk with critical weights

The following remark was suggested by an anonymous referee. There is a special value of the fugacity parameter z , namely

$$z^2 = \tan(\pi/8) \tag{1.8}$$

such that the *even* walk Z_{even} coincides (after a small change in the embedding) with the random walk on isoradial graphs with critical weights considered in the work of Kenyon [24]. To see this, one can notice that the even walk Z_{even} is equivalent to a random walk on two upper-half planes (or more precisely, square lattices on these half planes) welded together via a row of triangles. See Fig. 2. Such a graph has an isoradial embedding and the corresponding critical weights have weight 1 on the square lattice edges, and weight z given by (1.8) on the remaining triangle edges, as follows from elementary calculations. In that case, convergence of the derivative of the potential kernel (i.e., part of the inverse Kasteleyn matrix) would follow from Theorem 4.3 in [24].

1.8 Outline of the paper and structure of proof

This paper is organized as follows: In Sect. 2 we first describe the measure preserving mapping from [16] between the free boundary dimer model on \mathcal{G} and the standard dimer model on an augmented (nonbipartite) graph \mathcal{G}^0 (as in Fig. 3), which is the starting point of our paper. We then define a Kasteleyn orientation on \mathcal{G}^0 , the associated Kasteleyn matrix \tilde{K} , and finally we choose a convenient complex-valued gauge changed Kasteleyn matrix K (this gauge is closely related to the one of Kenyon [22] and allows one to interpret K as a discrete Dirac operator). Kasteleyn theory (which we recall later on in the paper) says that the correlations of the dimer model on \mathcal{G}^0

Fig. 2 After reflecting the white sublattice of V_{even} (dashed lines), a graph composed of two square lattice half-planes glued together by a row of interlacing isosceles triangles is formed (solid lines). If the angle between the legs of the triangle is equal to $\pi/4$, the graph is isoradial and the walk Z_{even} for $z^2 = \tan(\pi/8)$ is the same as the walk studied by Kenyon in [24]

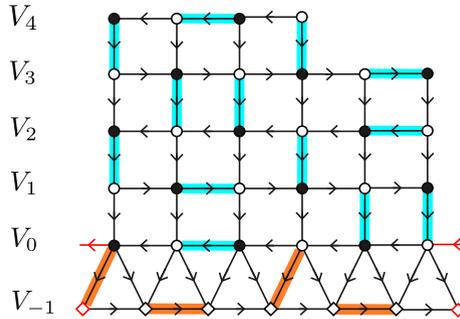
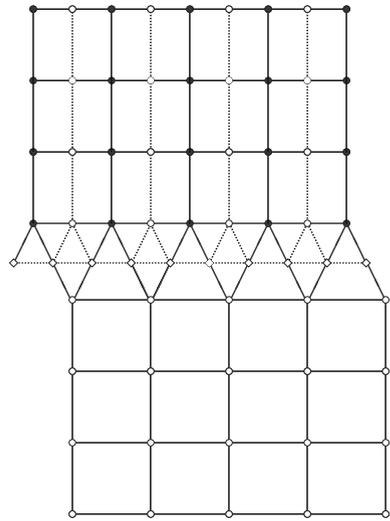


Fig. 3 An augmented non-bipartite graph \mathcal{G}^0 and its Kasteleyn orientation. The graph is constructed from a piece of the square lattice \mathcal{G} with $\partial_m \mathcal{G} = V_0$ by adding the bottom circuit of triangles. The two red edges and vertices in the figure denote a single edge and vertex in \mathcal{G}^0 . Diamonds represent vertices of $\mathcal{G}^0 \setminus \mathcal{G}$. We assume that $|V_0|$ even (here equal 6) so that the mapping of [16] can be directly applied. In this case \mathcal{G} has one black and one white monomer-corner, three black dimer-corners, and one white dimer-corner. The additional circuit of triangles simulates the presence of monomers in the free boundary dimer model by means of a standard dimer model. This is expressed as a measure-preserving two-to-one map between $\mathcal{D}(\mathcal{G}^0)$ and $\mathcal{MD}(G)$ with a proper choice of weights [16]. One of the two dimer covers of \mathcal{G}^0 corresponding to the fixed monomer-dimer cover of \mathcal{G} is depicted in orange (color figure online)

(and hence also of the free boundary dimer model on \mathcal{G}) can be computed from the inverse Kasteleyn matrix K^{-1} .

With an intention of developing its random walk representation, we therefore begin analyzing the inverse Kasteleyn matrix when \mathcal{G} is a subgraph of the square lattice with appropriate boundary conditions described in Sect. 1.2. To this end we look at the matrix $D = K^*K$, whose off-diagonal entries we interpret as (signed) transition weights. These weights away from $\partial_{\text{free}} \mathcal{G}$ (which is a subset of the real line) are positive and hence define proper random walks as in [22]. However, the description of D as a Laplacian matrix associated to a random walk breaks down completely for vertices

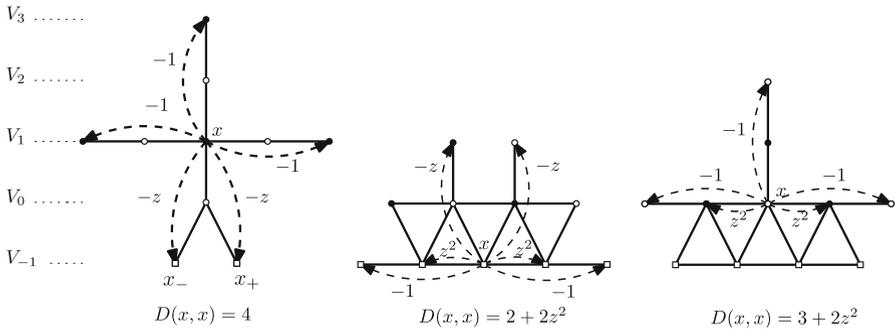


Fig. 4 Three types of vertices x where the transition weights of $D = K^*K$ are signed. The arrows indicate the corresponding value of $D(x, y)$. Note the following crucial observations: First, in the rightmost case (when $x \in V_0$), the absolute values of the transition weights sum up to the diagonal term. Moreover, the transition weight is negative if and only if the size of the step is odd (more precisely equal to one). A similar observation holds in the central picture (when $x \in V_{-1}$) if one ignores the transition weights that lead back to V_1 . This is the basis for the construction of Sect. 2.3 and the definition of the auxiliary random walk on \mathbb{Z} from (2.7). Our approach is to “forget” what the walk does when it stays in V_{-1} and resum over all trajectories contained in V_{-1} and with the same endpoints in V_1

on the three bottommost rows of \mathcal{G}^0 (as in Fig. 4). We stress the fact that the level of complication is considerably higher for transitions between odd rows (that will lead to the definition of the walk Z_{odd}). Indeed, as mentioned in Fig. 4, for even rows the arising walk Z_{even} can be relatively easily understood as a proper random walk reflected on the real line after taking into account a *global* sign factor appearing in D (which leads to the formula in the second line of (1.4)).

Therefore the remainder of Sect. 2 is devoted to the random walk representation for K^{-1} , which is one of the main contributions of this paper. The main idea is to “forget” the steps of the signed walk induced by D taken along the row V_{-1} , or more precisely to only specify the trajectory of a path away from V_{-1} and combine together all paths that agree with this choice. The hope is that the resulting projected signed measure on trajectories contained in $V_{\text{odd}} = V_1 \cup V_3 \cup \dots$ with (unbounded) steps from V_1 to V_1 is actually a true probability measure. Remarkably (in our opinion), we show that this is indeed the case; this phenomenon is what really lies behind the random walk representation of Theorem 1.1. To achieve this, an additional (intermediate) limiting procedure is required. To be precise, we first pretend that the rows V_0 and V_{-1} of \mathcal{G}^0 are infinite. This is done by defining graphs \mathcal{G}^N , where a circuit of additional triangles is appended to \mathcal{G} , and then taking the limit while the size of the circuit diverges. This allows us to perform exact computations for the transition weights from V_1 to V_1 by analysing the potential kernel of the auxiliary one-dimensional walk on \mathbb{Z} defined in Sect. 2.5. The required positivity of the combined weights and the identity stating that these weights sum to one as we sum over all possible jump locations, stated in Lemma 2.3, is the result of an exact (and rather long) computation involving the potential kernel of this auxiliary walk.

This intermediary limit is also the technical reason for the introduction of the modified monomer weight z' , which arises as the limiting weight of the peripheral monomers on \mathcal{G} . Finally, in Sect. 2.6 we use the notion of **Schur complement** of a

matrix as a convenient tool to implement the idea of combining all the walks with given excursions away from (the now infinite) row V_{-1} . All in all, at the end of Sect. 2 a random walk representation of K^{-1} is developed and Theorem 1.1 is proved.

The goal of this section is to prove Theorem 1.2, i.e., to establish the infinite volume limit of the model when a sequence of graphs exhausts $\mathbb{H} \cap \mathbb{Z}^2$. By Kasteleyn theory, it is enough that the inverse Kasteleyn matrix has a limit. This will be shown using the random walk representation established in Sect. 2. Essentially, the main goal is to show that in the infinite volume limit, the difference of the Green function associated to the random walk Z_{even} or Z_{odd} at two fixed vertices x, y converge to the difference of the **potential kernel** of the corresponding infinite volume walk. In fact, the very definition of this potential kernel is far from clear and occupies us for a sizeable part of this section. For the usual simple random walk on the square lattice, the definition of the potential kernel (see e.g. [30]) relies on precise estimates for the random walk coming from the exact computation of the Fourier transform of the law of random walk. Such an exact computation is clearly impossible here, since the effective walks cannot be viewed as a sum of i.i.d. random variables. We overcome this obstacle by developing a general method (which we think may be of independent interest) to define the potential kernel of a recurrent random walk and prove convergence of Green function differences towards it. The main idea is to proceed by **coupling**. We note that a similar idea has also been recently advocated by Popov (see Section 3.2 of [34]); but the approach in [34] also takes advantage of some properties and symmetries which are not available here. Instead, our starting point is the robust estimate of Nash (see e.g. [3]) characterising the heat kernel decay. With our approach, only a weak (polynomial of any order) bound for the probability of non coupling suffices to show the existence of the potential kernel. An immediate byproduct of our quantitative approach (which is crucial for us) is the proof of the desired convergence of Green function differences towards the differences of the potential kernel, obtained in Proposition 3.7.

In Sect. 4 we move on to describe the scaling limit (now in the limit of fine mesh size) for the potential kernel of the effective walks Z_{even} and Z_{odd} . A key idea is to say that when such a walk hits the real line, it will hit it many times and therefore has a probability roughly 1/2 to end up at a vertex with even (resp. odd) horizontal coordinate once it is reasonably far away from the real line. This idea eventually leads us to asymptotic formulae for the potential kernel which depends on the parity of the horizontal coordinate of a point (see Theorem 4.1). To achieve this, we introduce an intermediary process which we call **coloured random walk**, which is a random walk on (twice) the usual square lattice, but which can also carry a colour (representing, roughly speaking, the actual parity of the effective walk). This colour may change only when the walk hits the real line, and then does so with a fixed probability p . The proof of Theorem 4.1 relies on first comparing our effective walk to the coloured random walk (Proposition 4.3) and then from the coloured walk to half of the potential kernel of the usual simple random walk (Proposition 4.4).

We are now finally in a position to start the proof of Theorem 1.3. From Theorem 4.1 we obtain a scaling limit for the inverse Kasteleyn matrix of the (infinite volume) free boundary dimer model. After recalling Kasteleyn theory in the nonbipartite setting, we then compute the scaling limit of the pointwise moments of height function differences on \mathbb{H} in Sect. 5.4. The argument is based on Kenyon's original computation [22] but with substantial modifications coming from the fact that we use Pfaffian formulas instead of the determinantal formulas for bipartite graphs. This leads to different expressions which fortunately simplify asymptotically (for reasons that are related but distinct from those in [22]). This leads to the formula in Proposition 5.6, which is an asymptotic expression for the limiting joint moments of pointwise height differences, with an explicit quantification of the validity of the limiting formula (needed in the following). To finish the proof of the result, we transfer this result in Sect. 5.5 into one about the scaling limit of the height function as a random distribution. This is essentially obtained by integrating the result of Proposition 5.6, but extra arguments are needed for the case when some of the variables of integration are close to one another.

Remark 1.6 An alternative strategy for establishing the scaling limit of the inverse Kasteleyn matrix, suggested by an anonymous referee, would be the following. It would suffice to concentrate on one of the two types of walks (say Z_{even} , which is simpler to define than its counterpart Z_{odd}) and analyse its potential kernel in the manner indicated above in order to derive the scaling limit of $K^{-1}(u, \cdot)$ where $u \in V_{\text{even}}$ is a given even vertex. Once this is done, discrete holomorphicity and antisymmetry can be invoked to obtain asymptotics of K^{-1} on the remaining vertices, and with the same error bounds (using a discrete version of the Poisson formula for the derivative of harmonic functions).

We have chosen not to implement this strategy for the following reasons. On the one hand, the asymptotic analysis of Z_{even} (and in particular its potential kernel) is as difficult as it is for Z_{odd} . As this is probably the most challenging part of the analysis, there would be no real simplification in considering Z_{even} only. On the other hand, the exact random walk representation of K^{-1} seems interesting in its own right, especially since it shows a connection with reflected random walks even at the discrete level.

We end the introduction by mentioning the following problem. The dimer model on special families of bipartite planar graphs is famously related, through various measure preserving maps, to other classical models of statistical mechanics like spanning trees (see e.g. [28]), the double Ising model [9, 11] or the closely related double random current model [13]. This indicates the following direction of study.

Problem 1.7 *Analyse the boundary conditions in these classical lattice models induced by the presence of monomers in their dimer model representations.*

2 (Inverse) Kasteleyn matrix

2.1 Dimer representation

In [16] a representation of the free boundary dimer model was given in terms of a dimer model on an augmented (nonbipartite) graph where a circuit of triangles is appended to $\partial_{\text{free}}\mathcal{G}$. By circuit we mean here that the additional triangles form a cycle. See Fig. 3 for a picture.

Here we state a lemma that conveniently fits into our framework but the result holds in much bigger generality.

Lemma 2.1 ([16]) *Let \mathcal{G} be a finite subgraph of the upper-half plane square lattice such that $\partial_m\mathcal{G}$ is contained in the real line and forms an interval of an even number of vertices. Let \mathcal{G}^0 be \mathcal{G} with an appended bottom circuit of triangles as in Fig. 3. Assign weight z_v to each edge of the triangle that is incident to a vertex $v \in \partial_m\mathcal{G}$. Then, for each monomer-dimer cover in $\mathcal{MD}(\mathcal{G})$, there exist exactly two dimer covers in $\mathcal{D}(\mathcal{G}^0)$. Moreover, this is a weight-preserving mapping.*

In other words, there is a measure preserving two-to-one map between $\mathcal{D}(\mathcal{G}^0)$ and $\mathcal{MD}(\mathcal{G})$.

2.2 Kasteleyn orientation, Kasteleyn matrix and gauge change

A Kasteleyn orientation of a planar graph is an assignment of orientations to its edges such that for each face of the graph, as we traverse the edges surrounding this face one by one in a counterclockwise direction, we encounter an odd number of edges in the opposite direction (see e.g. [40]). For graphs as defined in Sect. 1.2 (but with an extra row of triangles, as explained in the previous section) we make the following choice (see Fig. 3): every vertical line is oriented downwards (including the non-horizontal sides of triangular faces at the bottom). The orientation of horizontal edges alternates: in odd rows (starting at row -1): edges are oriented from left to right, whereas in even rows (starting at row 0) they are oriented from right to left.

Given a Kasteleyn orientation, the standard Kasteleyn matrix $\tilde{K}(x, y)$ is taken to be the signed, weighted adjacency matrix: that is, $\tilde{K}(x, y) = \pm 1_{x \sim y} w_{(x, y)}$ where the sign is $+$ if and only if the edge is oriented from x to y , and the weight $w_{(x, y)}$ is 1 for horizontal and vertical edges (including on V_{-1}), and z for the nonhorizontal sides of triangular faces. However, it will be useful to perform a change of gauge, as follows. For every $k \geq 0$ even, and for every $x \in V_k$, we multiply by i the weight of every edge adjacent to x . In particular, every horizontal edge in V_k with k even receives a factor of i twice coming from both of its endpoints, whereas each vertical edge receives a factor of i exactly once. We define the gauge-changed Kasteleyn matrix $K(x, y)$ to be the resulting matrix. Formally,

$$K(x, y) := \tilde{K}(x, y) i^{1_{x \in V_{\text{even}}} + 1_{y \in V_{\text{even}}}}. \quad (2.1)$$

For instance, if $x \in V_0$ is not on the boundary, then x has five neighbours. Starting from the vertical edge and moving counterclockwise, the weights $K(x, y)$ are given by $-i, -1, iz, iz, 1$.

2.3 Towards the inverse Kasteleyn matrix

Let $D = K^*K$. In this section we explain the key idea involved in computing D^{-1} , and thus ultimately K^{-1} . The matrix D already played a crucial role in [22], where Kenyon observed that it reduced to the Laplacian on the four types of sublattices of the square grid.

We will follow a similar approach but, as we will see, the immediate interpretation of D as a Laplacian breaks down in the rows V_{-1} , V_0 and V_1 . Nevertheless, admitting the formal sum-over-all-paths identity (2.2), we will be able to make a guess on the structure of D^{-1} . This will ultimately lead us to the identification of D^{-1} as the Green's function of a certain **effective** random walk (or, in fact, a pair of effective random walks) which appear in the statement of Theorem 1.1.

Therefore, the purpose of this section is mostly to explain the heuristic principles guiding the proof, and to introduce the relevant objects and the notation. Once this framework is defined we will start with the actual proof in Sect. 2.4. We will complete the rigorous computation of D^{-1} (and therefore the proof of Theorem 1.1) in Sect. 2.6.

We now fix a finite arbitrary graph \mathcal{G} that satisfies the conditions of Sect. 1.2. We first compute D explicitly. Note that if $x \in V_k$ with $k \geq 2$, the entries of D are computed in a way identical to Kenyon [22]. Namely, the diagonal term is

$$D(x, x) = K^*K(x, x) = \sum_{y \sim x} K^*(x, y)K(y, x) = \sum_{y \sim x} |K(y, x)|^2 = \deg(x).$$

Moreover, the off-diagonal terms are nonzero if and only if y is at distance two from x , but not diagonally (the diagonal cancellation is a consequence of the Kasteleyn orientation), i.e., if y is a neighbour of x on one of the sublattices $2\mathbb{Z} \times 2\mathbb{Z}$, $(2\mathbb{Z} + 1) \times (2\mathbb{Z} + 1)$, $2\mathbb{Z} \times (2\mathbb{Z} + 1)$ or $(2\mathbb{Z} + 1) \times 2\mathbb{Z}$ in which case one can check as above that $D(x, y) = -1$. Therefore away from the boundary $\partial_{\text{free}}\mathcal{G}$, in the same way as in [22], D is the Laplace operator associated to a simple random walk on each of the sublattices, up to a multiplicative constant.

Complications for such an interpretation arise when $x \in V_{-1} \cup V_0 \cup V_1$. See Fig. 4 for the nonzero entries of D in these cases. Notice that now it is not necessarily true that the diagonal term $D(x, x)$ is (up to a sign) the same as the sum of the off-diagonal entries on the row corresponding to x , or in other words, the transition weights $d_{x,y}$ in (2.3) do not sum up to 1. Moreover, some of them are negative. While this seems like a very serious obstacle for describing the behaviour of the operator D^{-1} in the scaling limit, we nevertheless show in the next section how we can recover an effective random walk for which D really is the Laplacian.

More precisely, D^{-1} can be formally viewed as a sum of weights of paths of all possible lengths, where the weight of a path is the product of (signed) transition weights of individual jumps. That is, formally,

$$D^{-1}(u, v) = \frac{1}{D(v, v)} \sum_{\pi: u \rightarrow v} w(\pi), \tag{2.2}$$

where for a path $\pi : u \rightarrow v$,

$$w(\pi) = \prod_{(x,y) \in \pi} d_{x,y} \quad \text{with} \quad d_{x,y} = -\frac{D(x, y)}{D(x, x)}. \tag{2.3}$$

For x in the bulk, $d_{x,y} = 1/4$ for each y which is neighbour of x on the sublattice of twice larger mesh size containing x , and is 0 otherwise, which is the same as the transition probability of a simple random walk on that sublattice.

Let us now point out that the transition weights between an even row and an odd row are always 0. Compared to the odd rows, the construction for even rows is much simpler. As seen in Fig. 4, for $x \in V_0$, $D(x, x)$ is in fact equal to the sum of $|D(x, y)|$ for all $y \neq x$. We can therefore view $|d_{x,y}|$ for $x \in V_0$ as the transition weights of a random walk that is reflected on row V_0 (and can make jumps of size one and two on that row). When we take into account the signs of $d(x, y)$ in (2.3), this gives rise to a global sign factor which depends only on u and v can be seen in the second line of (1.4).

The rest of this section is devoted to the more complicated task of giving a random walk representation to D^{-1} restricted to the vertices in odd rows V_{odd} . We now describe the main idea. We will manage to give a meaning to the right hand side of (2.2) by fixing a specific order of summation. We will later on prove that this definition really does give us the inverse of D , and we will also find a random walk interpretation to this definition. We emphasise this because the signs are not constant, and hence the order of summation is a priori relevant to the value of the sum. Essentially we will compute the sum in (2.2) by ignoring the details of what the path does when it visits V_{-1} . That is, we will identify two paths if they enter V_{-1} at the same place in V_1 and leave V_{-1} at the same places in V_1 for each visit to V_{-1} , and we will be able to estimate contributions to (2.2) coming from each such equivalence class.

An important observation (see Fig. 4) here is that for each $x \in V_{-1}$, the diagonal term $D(x, x)$ is equal to the sum of $|D(x, y)|$ for all $y \in V_{-1}$ not equal to x . Note that $D(x, y)$ is non zero for $y = x \pm 1$ or $y = x \pm 2$ (understood cyclically). This allows us to express the weight of the paths which stay in V_{-1} as the weight of a random walk with steps ± 1 and ± 2 on V_{-1} . One can therefore associate a Green's function $g(\cdot, \cdot)$ (or more formally a potential kernel since the walk is recurrent on the cycle V_{-1}) with the random walk on V_{-1} with transition probabilities

$$p_{x,y} = |d_{x,y}|. \tag{2.4}$$

For $x \in V_1 = V_1(\mathcal{G}^0) = V_1(\mathcal{G})$, let x_- and x_+ be the left and right vertex in $V_{-1} = V_{-1}(\mathcal{G}^0)$ two steps away from x (see Fig. 4). We fix $u, v \in V_1$ and let $u_\bullet \in \{u_-, u_+\}$ and $v_\bullet \in \{v_-, v_+\}$. We define $\mathcal{P}_{u_\bullet, v_\bullet}^1$ to be the set of paths from u_\bullet to v_\bullet which are contained in V_{-1} . Observe that if $\pi \in \mathcal{P}_{u_\bullet, v_\bullet}^1$, then π makes jumps of size ± 1 or ± 2 , and that each odd jump contributes a negative weight to (2.2) whereas each even jump

contributes a positive weight. Since π goes from u_\bullet to v_\bullet the parity of the number of odd jumps is fixed and depends only on the distance between u_\bullet and v_\bullet in V_{-1} . Hence

$$w(\pi) = (-1)^{\Re(v_\bullet - u_\bullet)} \prod_{(x,y) \in \pi} |d_{x,y}|,$$

where $d_{x,y}$ is defined in (2.3).

Going further: if $\mathcal{P}_{u_\bullet, v}^1$ is the set of paths going from u_\bullet to v and staying in V_{-1} (except for the last step, which must be from v_\pm to v), then

$$\sum_{\pi \in \mathcal{P}_{u_\bullet, v}^1} w(\pi) = (-1)^{\Re(v_+ - u_\bullet)} (g(u_\bullet, v_+) - g(u_\bullet, v_-)) \frac{z}{2 + 2z^2} \tag{2.5}$$

where the last term accounts for the weight $-D(v_\pm, v)/D(v_\pm, v_\pm)$ of the last step from V_{-1} to V_1 . Finally, let $\mathcal{P}_{u, v}^1$ be the set of paths from u to v which stay in V_{-1} except for the first and last step (which necessarily are from V_1 to V_{-1} and vice versa). Using (2.5) we have

$$\begin{aligned} \sum_{\pi \in \mathcal{P}_{u, v}^1} w(\pi) &= \frac{z^2}{8 + 8z^2} (-1)^{\Re(v - u)} (g(u_+, v_+) - g(u_+, v_-) \\ &\quad - g(u_-, v_+) + g(u_-, v_-)) =: \frac{1}{4} q_{u, v}, \end{aligned} \tag{2.6}$$

where the additional term $\frac{z}{4}$ compared to (2.5) accounts for the weight $-D(u, u_\pm)/D(u, u)$ of the first step from V_1 to V_{-1} . The factor $\frac{1}{4}$ in the definition of $q_{u, v}$ is included for later convenience. As mentioned, the Green’s function for the walk on V_{-1} is not directly defined, but its gradient (as in the expression above) will be defined with the use of the potential kernel.

Recall that our intention is to interpret the quantities $q_{u, v}$ as transition probabilities between vertices in V_1 . In particular we would wish $q_{u, v}$ to be positive and sum up to (something less than) one (since the other three transition weights induced by D from a vertex in the bulk of V_1 to V_1 and V_3 are equal to $3/4$). Unfortunately, in the setting described so far, we were unable to do so. However, a nice solution to this problem is the following construction. We note that this construction is the reason for the appearance of the special monomer weight z' at the monomer-corners in the statement of our results.

2.4 An intermediate limit

To overcome the issue raised above, we introduce an intermediate limiting procedure in our model. To this end, let $2m$ denote the number of triangles in \mathcal{G}^0 , and let \mathcal{G}^N be the graph where we add a circuit of $2N \wedge 2m$ triangles (instead of $2m$ triangles as in \mathcal{G}^0), see Fig. 5 for an example. We assign weights 1 to every edge except if it belongs to a triangle and is not horizontal, in which case we assign weight z . Since we assumed

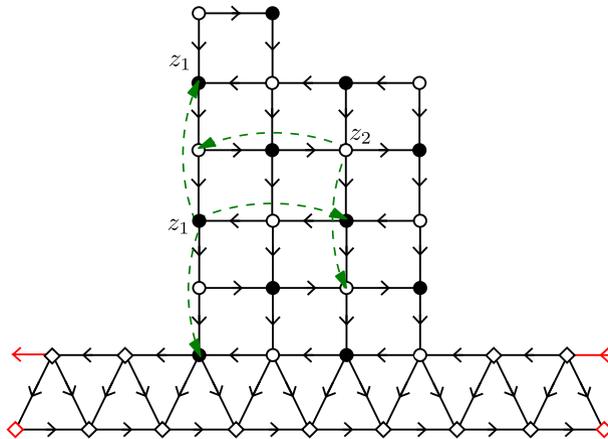


Fig. 5 A graph \mathcal{G} (drawn with round vertices), and its extension \mathcal{G}^8 (there is a circuit of 16 additional triangles appended to the bottom of \mathcal{G}). The two red edges and vertices in the figure denote a single edge and vertex in \mathcal{G}^8 . The weights of transitions denoted by arrows are $D(\cdot, \cdot) = D_N(\cdot, \cdot) = -1$. The diagonal terms are $D(z_1, z_1) = 3$, $D(z_2, z_2) = 4$. The black vertex z_1 has Neumann boundary conditions for the associated walk, since the total weight of outgoing transitions is also 3. The white vertex z_2 has Dirichlet boundary conditions since the total outgoing weight is $2 < 4$ (color figure online)

that \mathcal{G} has a dimer cover, it is easy to see that \mathcal{G}^N also has at least one dimer cover. We can hence talk about the dimer model on \mathcal{G}^N with the specified weights.

Let us call $\tilde{\mathcal{G}}^N$ the graph \mathcal{G}^N where the bottom row $V_{-1}(\mathcal{G}^N)$ is removed. In other words, $\tilde{\mathcal{G}}^N$ is \mathcal{G} with $\partial_{\text{free}}\mathcal{G}$ replaced by the N -cycle. Using Lemma 2.1, we can rephrase the dimer model on \mathcal{G}^N as a free boundary dimer model on $\tilde{\mathcal{G}}^N$. We claim that the monomer-dimer configuration on $\tilde{\mathcal{G}}^N$ restricted to \mathcal{G} , in the limit $N \rightarrow \infty$ has the law of the free boundary dimer model with weight z' from (1.3) at the monomer-corners. To see this, let \mathcal{Z}_k be the partition function of the free boundary dimer model on a segment of \mathbb{Z} of length k with monomer weight z and edge weight 1 (monomers are allowed anywhere on the segment). Moreover, let L be the number of vertices in the row $V_0(\mathcal{G}^N)$ that are not in \mathcal{G} (so $L = (N - m)_+$). Then, it is not difficult to see that the effective weight for the presence of two (resp. one and zero) corner-monomers in \mathcal{G} is \mathcal{Z}_{L+2} (resp. \mathcal{Z}_{L+1} and \mathcal{Z}_L). The claim is therefore a consequence of the following elementary lemma.

Lemma 2.2 As $k \rightarrow \infty$,

$$\frac{\mathcal{Z}_{k+1}}{\mathcal{Z}_k} \rightarrow z', \quad \text{where } z' = \frac{z}{2} + \sqrt{1 + \frac{z^2}{4}}.$$

Proof It is enough to solve the recursion $\mathcal{Z}_{k+1} = z\mathcal{Z}_k + \mathcal{Z}_{k-1}$ to get that

$$\mathcal{Z}_k = \left(\frac{1}{2} - \frac{z}{4\beta}\right)\left(\frac{z}{2} - \beta\right)^k + \left(\frac{1}{2} + \frac{z}{4\beta}\right)\left(\frac{z}{2} + \beta\right)^k,$$

where $\beta = \sqrt{1 + \frac{z^2}{4}}$. □

Let K_N be the Kasteleyn matrix of \mathcal{G}^N with Kasteleyn weights as discussed above, see e.g. Fig. 5, and let $D_N = (K_N)^* K_N$. The statement above and Kasteleyn theory imply that the inverse Kasteleyn matrix K_N^{-1} restricted to \mathcal{G}^0 converges as $N \rightarrow \infty$ to the inverse Kasteleyn matrix $(K')^{-1}$ for the free boundary dimer model on \mathcal{G}^0 with monomer weights z' at the monomer-corners.

2.5 An auxiliary walk on \mathbb{Z}

It will be convenient to consider a random walk on $V_{-1}(\mathbb{Z}^2) \simeq \mathbb{Z}$ with transition probabilities given by

$$p_{x,x\pm 1}^\infty = \frac{z^2}{2 + 2z^2} =: 1/2 - p, \quad p_{x,x\pm 2}^\infty = \frac{1}{2 + 2z^2} =: p. \tag{2.7}$$

In other words, this is the infinite volume version of the walk from (2.4). Now, while the Green’s function of this walk is infinite since the walk is recurrent, its differences makes sense in the form of the **potential kernel** (see [30], Section 4.4.3) given by

$$\alpha_k = \sum_{n=0}^\infty (p_n(0) - p_n(k)) = \lim_{N \rightarrow \infty} \left(\sum_{n=0}^N p_n(0) - \sum_{n=0}^N p_n(k) \right), \tag{2.8}$$

where $p_n(k) = \sum_{i=0}^n \mathbb{P}_0(X_i = k)$ with X being the random walk with jump distribution (2.7). Using the potential kernel, for $u, v \in V_1(\mathbb{Z}^2) \simeq \mathbb{Z}$, we can now define the infinite volume version of the transition weight $q_{u,v}$ from (2.6) by

$$q_{u,v}^\infty = \frac{z^2}{2 + 2z^2} (-1)^{k+1} (2\alpha_k - \alpha_{k+1} - \alpha_{k-1}), \tag{2.9}$$

where $k = \mathfrak{N}(v - u)$. Note that the sign is opposite to that in (2.6). To is due to the $-p_n(k)$ term in the definition of the potential kernel.

The next result is one of the crucial observations in this work.

Lemma 2.3 (Effective transition probabilities) *For all $z > 0$, and any pair of vertices $u, v \in V_1(\mathbb{Z}^2)$, we have*

$$q_{u,v}^\infty \geq 0 \quad \text{and} \quad \sum_{v \in V_1(\mathbb{Z}^2)} q_{u,v}^\infty = 1.$$

Moreover, $q_{u,v}^\infty \rightarrow 0$ exponentially fast as $|u - v| \rightarrow \infty$.

Before we give the proof note that neither of these two facts is at all clear from the definitions of $q_{u,v}^\infty$. Together they imply that we can think of $q_{u,v}^\infty$ as the step distribution of some **effective random walk** on V_1 . Later in Lemma 2.6, we will prove that $q_{u,v}^\infty$ is the limit of $q_{u,v}$ from (2.6) on \mathcal{G}^N when $N \rightarrow \infty$.

Proof of Lemma 2.3 The proof is based on an exact formula for the potential kernel α of the walk on \mathbb{Z} defined by (2.7). To start with, by Theorem 4.4.8 from [30] we know that

$$\alpha_k = \frac{|k|}{\sigma^2} + A + O(e^{-\beta|k|})$$

for some constants $\beta > 0$ and $A \in \mathbb{R}$, and where $\sigma^2 = 1 + 6p$ is the variance of the walk with p as in (2.7). Moreover, α is harmonic (except at $k = 0$) with respect to the Laplacian of the walk (2.7). The recursion for the sequence α_k has the following characteristic equation

$$1 = (1/2 - p)(\gamma + \gamma^{-1}) + p(\gamma^2 + \gamma^{-2}).$$

This equation has one double root at 1, and two other roots

$$\gamma_{1,2} = \pm \sqrt{\left(\frac{1}{2} + \frac{1}{4p}\right)^2 - 1} - \frac{1}{2} - \frac{1}{4p}.$$

Note that since $z > 0$, we have $p < 1/2$, so the two roots above are distinct. This implies that the $O(e^{-\beta|k|})$ term is of the form $B_1\gamma_1^{|k|} + B_2\gamma_2^{|k|}$ for some constants B_1 and B_2 . However, we have $|\gamma_1| < 1$ and $|\gamma_2| > 1$, so we must have $B_2 = 0$. This implies that

$$\alpha_k = \frac{|k|}{\sigma^2} + A + B\gamma^{|k|}$$

for some constants A and B , and

$$\gamma = \sqrt{\left(\frac{1}{2} + \frac{1}{4p}\right)^2 - 1} - \frac{1}{2} - \frac{1}{4p}. \tag{2.10}$$

Using that $\alpha_0 = 0$ by definition, we get $A = -B$ and hence

$$\alpha_k = \frac{|k|}{1 + 6p} - B + B\gamma^{|k|}. \tag{2.11}$$

We still need to compute B which is equivalent to computing α_1 . Let X be the walk with transition probabilities (2.7). Let $\tau = \inf\{n > 0 : X_n > 0\}$, and

$$q = \mathbb{P}_0(X_\tau = 1) \quad \text{and} \quad 1 - q = \mathbb{P}_0(X_\tau = 2).$$

Then, by considering the possible four different first steps $(+1, -1, +2, -2)$ of X and using translation invariance and the strong Markov property, we get that

$$q = \left(\frac{1}{2} - p\right) + \left(\frac{1}{2} - p\right)((1 - q) + q^2) + p(q(1 - q) + q^3 + (1 - q)q),$$

which simplifies to

$$pq^2 + \left(\frac{1}{2} - p\right)(q - 1) = 0. \tag{2.12}$$

One can check that $q = \gamma + 1$. Moreover, using the symmetry of jumps of X and the Markov property for the walk, we get the equation (again considering the first four steps in the same order)

$$\begin{aligned} \alpha_1 = & 1 + \left(\frac{1}{2} - p\right)(-\alpha_1) + \left(\frac{1}{2} - p\right)[q\alpha_1 + (1 - q)(-\alpha_1)] \\ & + p[(1 - q)\alpha_1 + q(-\alpha_1)] + p[(q^2 + (1 - q))\alpha_1 + q(1 - q)(-\alpha_1)]. \end{aligned} \tag{2.13}$$

To justify (2.13), one starts from the definition of α_1 in (2.8) as the limit as $N \rightarrow \infty$ of the expected difference of number of visits by time N to the sites 0 and 1. We first apply the simple Markov property at the first step, and depending on the outcome of the first step, apply the strong Markov property at the next time τ (after time 1) that the walk returns to 0 or 1, taking care of the contribution coming from the event $\{\tau > N\}$. We then let $N \rightarrow \infty$. There is no problem in doing so, first because the sequence $\{\alpha_n\}_{n \geq 0}$ is bounded, which lets us use the dominated convergence theorem, and second because the contribution coming from the event $\{\tau > N\}$ to the difference between the number of visits at 0 and 1 by time N is bounded by 1. Details are left to the reader.

Together with (2.12), (2.13) gives

$$\alpha_1 = \frac{1}{(1 + 2p(-1 + q))(2 - q)} = \frac{1}{(1 + 2p\gamma)(1 - \gamma)}. \tag{2.14}$$

and hence from (2.11) we obtain

$$B = \frac{4p}{(\gamma - 1)(6p + 1)(2p\gamma + 1)} \leq 0. \tag{2.15}$$

We can now define

$$q_k = (-1)^{k+1} \frac{z^2}{2 + 2z^2} \Delta\alpha_k = (-1)^{k+1} \left(\frac{1}{2} - p\right) \Delta\alpha_k,$$

where $\Delta\alpha_k = 2\alpha_k - \alpha_{k+1} - \alpha_{k-1}$ is the Laplacian of simple random walk. Then $q_k = q_{u,v}^\infty$ whenever $|u - v| = k$. Using (2.11), we have

$$(-1)^{k+1} \Delta\alpha_k = \begin{cases} -B|\gamma|^{|k|}(2 - \gamma - \gamma^{-1}) \geq 0 & \text{for } k \neq 0, \\ \frac{2}{1+6p} - 2B(1 - \gamma) \geq 0 & \text{for } k = 0. \end{cases} \tag{2.16}$$

and hence the total transition weight is

$$\sum_{k \in \mathbb{Z}} q_{|k|} = \left(\frac{1}{2} - p\right) \left(-2B(2 - \gamma - \gamma^{-1}) \frac{-\gamma}{1 + \gamma} - 2B(1 - \gamma) + \frac{2}{1 + 6p} \right)$$

$$= \left(\frac{1}{2} - p\right) \left(-4B \frac{1 - \gamma}{1 + \gamma} + \frac{2}{1 + 6p}\right). \tag{2.17}$$

Using (2.10) and (2.15), it can be checked that the last expression is equal to one for all $0 < p < 1/2$ (equivalently all $z > 0$). Exponential decay of q_k is clear from (2.16).

2.6 Random walk representation of D^{-1}

Here we finally establish a rigorous version of (2.2) using the ingredients from the previous sections. Recall that K_N is the Kasteleyn matrix of the graph \mathcal{G}^N and $D_N = (K_N)^* K_N$. We will be mostly interested in the restriction of D_N^{-1} to the vertices of \mathcal{G} . Observe that D_N can be written as a block-diagonal matrix if we consider vertices respectively in the odd or even rows (this is one advantage of taking periodic boundary conditions in the bottom rows). Hence to invert D_N it will suffice to invert each of these blocks separately. We call D_N^{odd} (resp. D_N^{even}) the matrix D_N restricted to $V_{\text{odd}}(\mathcal{G}^N) \cup V_{-1}(\mathcal{G}^N)$ (resp. $V_{\text{even}}(\mathcal{G}^N)$).

We first focus on the odd case (the even case is much easier as explained before), and for now we will write D_N for D_N^{odd} . The key idea will be to use the **Schur complement formula**. To be more precise, we observe that D_N has the block structure

$$D_N = \begin{pmatrix} \mathcal{A} & \mathcal{B} \\ \mathcal{B}^T & \mathcal{C} \end{pmatrix},$$

where \mathcal{A} is indexed by the special row V_{-1} , and \mathcal{C} is indexed by all the other rows V_{odd} . Hence \mathcal{B} and \mathcal{B}^T can be thought of as a “transition matrices” between V_{-1} and V_{odd} . Note that these matrices depend on N but we don’t write this explicitly to lighten the notation. We define the Schur complement of \mathcal{A} to be the matrix

$$D_N/A := \mathcal{C} - \mathcal{B}^T \mathcal{A}^{-1} \mathcal{B}. \tag{2.18}$$

With this definition, the restriction of D_N^{-1} to V_{odd} is simply given by

$$D_N^{-1}|_{V_{\text{odd}}} = (D_N/A)^{-1}. \tag{2.19}$$

One issue for us is that we will see A is not directly invertible, making the use of the Schur complement formula not immediately possible. Instead, we will consider a limiting procedure: we will introduce a modification $D_{N,\varepsilon}$ of D_N . While $D_{N,\varepsilon}$ does not have any direct dimer interpretation, we will be able to compute its inverse $D_{N,\varepsilon}^{-1}$ using the Schur complement formula, and taking a limit as $\varepsilon \rightarrow 0$, deduce an expression for D_N^{-1} . (The key observation is that, while \mathcal{A}^{-1} is not well-defined, $\mathcal{B}^T \mathcal{A}^{-1} \mathcal{B}$ is.)

We now outline how we proceed.

- We introduce the matrix $D_{N,\varepsilon}$, which is a version of D_N but with an additional small killing probability at each vertex of V_{-1} .
- We then write the inverse $\mathcal{A}_\varepsilon^{-1}$ appearing in the block decomposition of $D_{N,\varepsilon}$ in terms of the potential kernel (i.e., difference of Green’s functions) for the random walk on $V_{-1}(\mathcal{G}^N)$ with transition probabilities as in (2.4) and with killing.

- As $\varepsilon \rightarrow 0$, and N is fixed, we show that the potential kernel of the periodic, killed walk converges to the potential kernel of the periodic eternal (unkilled) walk, and give a formula for the latter. From this formula it also follows that as the period size N diverges to ∞ , the potential kernel of the periodic eternal walk converges to the potential kernel of the infinite walk.
- This gives us a formula for the Schur complement D_N/\mathcal{A} via (2.18). We then use that for N sufficiently large, this Schur complement can be viewed as a (genuine) Laplacian for a random walk. The proof of this statement is postponed until Sect. 2.7.
- As a consequence of (2.19), this gives a formula for the inverse of D_N as a Green’s function of a genuine random walk.
- Finally, as the number N of triangles appended to \mathcal{G}^0 tends to infinity, on the one hand, the above analysis shows that the inverse Kasteleyn matrix (restricted to V_{odd}) can be written in terms of the potential kernel of a random walk with jumps along the boundary. On the other hand as mentioned before, the free boundary dimer model becomes equivalent to the same model on \mathcal{G}^0 with modified monomer weights z' as in (1.3) at the monomer-corners.
- The results of this section are summarised below as Corollary 2.8.

Let

$$D_{N,\varepsilon}(u, v) = D_N(u, v) + \varepsilon \mathbf{1}_{\{u=v \in V_{-1}(\mathcal{G}^N)\}}.$$

In words, we have added ε to the diagonal values on V_{-1} , which correspond to killing the walk with probability $p_\varepsilon = \varepsilon/D_{N,\varepsilon}(z, z) = O(\varepsilon)$ at each step in V_{-1} . $D_{N,\varepsilon}$ can still be written in block form as

$$D_{N,\varepsilon} = \begin{pmatrix} \mathcal{A}_\varepsilon & \mathcal{B} \\ \mathcal{B}^T & \mathcal{C} \end{pmatrix},$$

We now start with the computation of $\mathcal{A}_\varepsilon^{-1}$. To this end let

$$g^{N,\varepsilon}(u, v) = \sum_{\substack{\gamma: u \rightarrow v \\ \gamma \subseteq V_{-1}(\mathcal{G}^N)}} \prod_{e=(x,y) \in \gamma} p_{x,y}^{N,\varepsilon} \tag{2.20}$$

be the Green’s function of the random walk on $V_{-1}(\mathcal{G}^N)$ with transition probabilities $p_{x,y}^{N,\varepsilon}$ defined for \mathcal{G}^N as in (2.4) with killing probability p_ε ; that is,

$$p_{x,y}^{N,\varepsilon} = (1 - p_\varepsilon)p_{x,y}^N$$

with $p_{x,y}^N$ as in (2.7) (except here x, y are points on the N -cycle $\mathbb{Z}/(N\mathbb{Z})$, identified with $V_{-1}(\mathcal{G}^N)$: that is, $p_{x,y}^N$ is translation invariant on the cycle, $p_{x,x\pm 1}^N = z^2/(2+2z^2)$, $p_{x,x\pm 2}^N = 1/(2+2z^2)$, and the addition is mod N). Note that $g^{N,\varepsilon}(u, v)$ is well defined because of the killing.

Lemma 2.4 *Let $u, v \in V_{-1} = V_{-1}(\mathcal{G}^N)$. Then*

$$\mathcal{A}_\varepsilon^{-1}(u, v) = \frac{1}{\mathcal{A}(v, v)} (-1)^{\Re(u-v)} g^{N,\varepsilon}(u, v).$$

Proof This follows from the fact that $|\mathcal{A}_\varepsilon|$ is the Laplacian for the random walk described above, and moreover (as mentioned before) the sign of the transition weights induced by \mathcal{A}_ε is negative if the step is of size ± 1 and positive otherwise (step size ± 2). This follows from the definition of D_N and the Kasteleyn matrix.

We now explain how this yields an interpretation for the Schur complement $D_{N,\varepsilon}/\mathcal{A}_\varepsilon$ as a (genuine) Laplacian for a random walk in the bulk $V_{\text{odd}}(\mathcal{G})$ with jumps along the boundary $V_1(\mathcal{G})$. For $u, v \in V_1 = V_1(\mathcal{G}^N) = V_1(\mathcal{G})$, we define

$$q_{u,v}^{N,\varepsilon} = (\mathcal{B}^T \mathcal{A}_\varepsilon^{-1} \mathcal{B})(u, v). \tag{2.21}$$

Recalling that $D_N(v, v) = \mathcal{A}(v, v) = 2 + 2z^2$ for $v \in V_{-1}(\mathcal{G}^0)$ and $N \geq 1$, a straightforward computation using Lemma 2.4 shows that

$$q_{u,v}^{N,\varepsilon} = \frac{z^2}{2 + 2z^2} (-1)^{\Re(v-u)} \left((g^{N,\varepsilon}(u_+, v_+) - g^{N,\varepsilon}(u_+, v_-)) - (g^{N,\varepsilon}(u_-, v_+) - g^{N,\varepsilon}(u_-, v_-)) \right), \tag{2.22}$$

where again u_\pm, v_\pm are the left and right vertices in V_{-1} at distance two from u and v respectively.

We now show that as $\varepsilon \rightarrow 0$, the differences of Green’s functions on the right-hand side converge to a quantity, which could be viewed as the potential kernel of the walk on the N -cycle with transition probabilities $p_{x,y}^N$ from (2.4) (i.e., without killing). We note that the existence of this potential kernel is not obvious, as the walk is recurrent on a finite graph (so we cannot directly apply the results of Lawler and Limic [30]). Nevertheless, since the transition probabilities converge exponentially fast to their equilibrium distribution on this finite graph, it is also not hard to see that the series defining this potential kernel converges. However, we will not make any use of this definition, and will not verify the equality of this series with our limit, although this would not be very hard to establish.

Lemma 2.5 *Fix $u_\pm, v_\pm \in V_{-1}(\mathcal{G}^N)$. Then, as $\varepsilon \rightarrow 0$,*

$$(g^{N,\varepsilon}(u_+, v_+) - g^{N,\varepsilon}(u_+, v_-)) - (g^{N,\varepsilon}(u_-, v_+) - g^{N,\varepsilon}(u_-, v_-)) \rightarrow F^N(u_\pm, v_\pm)$$

where, by definition:

$$F^N(u_\pm, v_\pm) = \sum_{k \in \mathbb{Z}} \alpha(u_+, v_+ + kN) - \alpha(u_+, v_- + kN) - [\alpha(u_-, v_+ + kN) - \alpha(u_-, v_- + kN)], \tag{2.23}$$

$\alpha(x, o) = \alpha_{x-o}$ is the potential kernel defined in (2.8), and we identify each vertex u_{\pm}, v_{\pm} with a position in \mathbb{Z} (i.e., with its real part).

The proof of this lemma (including the finiteness of F^N) will be given in Sect. 2.7.

From the above lemma and (2.22) it follows directly that $q_{u,v}^{N,\varepsilon}$ has a limit as $\varepsilon \rightarrow 0$, which we call $q_{u,v}^N$. The next lemma allows us to take a second limit, now as $N \rightarrow \infty$: in fact, the results imply that for N sufficiently large, $q_{u,v}^N$ can really be viewed as transition probabilities, in the sense that they are positive and sum up to a quantity which is less than one. Recall the definition of $q_{u,v}^\infty$ from (2.9).

Lemma 2.6 *Let $u, v \in V_1 = V_1(\mathcal{G}^N) = V_1(\mathcal{G}^0)$. Then $q_{u,v}^N \rightarrow q_{u,v}^\infty$ as $N \rightarrow \infty$ pointwise. In particular, for N sufficiently large,*

$$q_{u,v}^N > 0 \quad \text{and} \quad \sum_{v \in V_1} q_{u,v}^N < 1. \tag{2.24}$$

Proof By the definition of $q_{u,v}^N$ and of $q_{u,v}^\infty$, it suffices to show that only the term corresponding to $k = 0$ survives in the limit. This follows at once from the exponential decay of differences of gradients of α established in (2.11), which shows every summand (except $k \neq 0$) tends to zero, and gives a uniform exponential bound on the summands, allowing the use of the dominated convergence theorem.

For the same reason, we can fix N large enough that (2.24) holds, and then ε small enough that

$$q_{u,v}^{N,\varepsilon} > 0 \quad \text{and} \quad \sum_{v \in V_1} q_{u,v}^{N,\varepsilon} < 1. \tag{2.25}$$

Note that the second inequality is strict since the sum is taken over $V_1(\mathcal{G}) \subsetneq V_1(\mathbb{Z}^2 \cap \mathbb{H})$. Now let N be sufficiently large that (2.24) holds true and ε sufficiently small that (2.25) also holds true, and consider a transition matrix between vertices in $u, v \in V_{\text{odd}}$ given by

$$R_{N,\varepsilon}(u, v) = I(u, v) - \frac{1}{\mathcal{C}(u, u)} \left(\mathcal{C}(u, v) - q_{u,v}^{N,\varepsilon} \mathbf{1}_{\{u,v \in V_1\}} \right), \tag{2.26}$$

where I is the identity. Note that

$$R_{N,\varepsilon}(u, v) \geq 0 \quad \text{and} \quad \sum_v R_{N,\varepsilon}(u, v) \leq 1$$

so that $R_{N,\varepsilon}$ is a substochastic matrix. Indeed, this follows from the definition of $\mathcal{C} = D_N|_{V_{\text{odd}}}$ and (2.25). In other words, we may add a cemetery absorbing point ∂ to the state space and declare $R_{N,\varepsilon}(x, \partial) = 1 - \sum_y R_{N,\varepsilon}(x, y) \geq 0$. This turns $R_{N,\varepsilon}$ into the transition matrix of a proper random walk on the augmented state space $V_{\text{odd}} \cup \{\partial\}$, which is absorbed at ∂ . We let $Z^{N,\varepsilon}$ be the random walk on $V_{\text{odd}} \cup \{\partial\}$ whose transition probabilities are given by $R_{N,\varepsilon}(x, y)$. We call this random walk the **effective (odd) bulk random walk**.

The interest of introducing the transition matrix $R_{N,\varepsilon}$ of this effective bulk random walk is that its associated Laplacian gives us the Schur complement D_N/A_ε : that is, for $u, v \in V_{\text{odd}}$, we have

$$(D_N/A_\varepsilon)(u, v) = \mathcal{C}(u, u)(I(u, v) - R_{N,\varepsilon}(u, v)), \tag{2.27}$$

which follows from the definition of the Schur complement (2.18), (2.21) and the definition of R_N .

From this formula and the Schur complement formula (2.19), it is immediate to deduce the following proposition. This implies that the inverse of $D_N^{\text{odd}} = D_N$ (i.e., the inverse of $(K_N)^*K_N$ restricted to bulk odd vertices) is given by the Green’s function of the effective bulk random walk. Recall that $\mathcal{C}(v, v) = D_N(v, v)$.

Proposition 2.7 *Let $u, v \in V_{\text{odd}}(\mathcal{G})$. Then for all N sufficiently large and ε sufficiently small, we have*

$$(D_{N,\varepsilon}^{\text{odd}})^{-1}(u, v) = G_{\text{odd}}^{N,\varepsilon}(u, v), \tag{2.28}$$

where $G_{\text{odd}}^{N,\varepsilon}$ is the (normalised) Green’s function associated to R_N , i.e.,

$$G_{\text{odd}}^{N,\varepsilon}(u, v) = \frac{1}{D_{N,\varepsilon}(v, v)} \mathbb{E}_u \left(\sum_{t=0}^{\infty} \mathbf{1}_{\{Z_t^{N,\varepsilon}=v\}} \right). \tag{2.29}$$

Taking a limit as $\varepsilon \rightarrow 0$, both (2.28) and (2.29) hold without ε .

When we take the limit as $\varepsilon \rightarrow 0$ (with N fixed sufficiently large so that (2.24) holds) in the above proposition, no convergence issue arises: indeed D_N is invertible (since K_N is invertible, as the partition function is assumed to be nonzero) so the inverse function is continuous at D_N . Likewise the convergence of the Green function of $Z^{N,\varepsilon}$ is a straightforward application of the dominated convergence theorem (since the walk can be absorbed somewhere on the boundary of \mathcal{G}).

We now address the even case, and write $D_N = D_N^{\text{even}}$. We introduce a “sign” diagonal matrix $S(x, x) = (-1)^{\mathfrak{N}(x)}$. Then, the matrix

$$\tilde{D}_N := S^{-1}D_N S$$

is positive on the diagonal and negative off-diagonal. Moreover, we have

$$\tilde{D}_N^{-1}(u, v) = G_{\text{even}}^N(u, v)$$

where

$$G_{\text{even}}^N(u, v) = \frac{1}{D_N(v, v)} \mathbb{E}_u \left(\sum_{t=0}^{\infty} \mathbf{1}_{\{\tilde{Z}_t=v\}} \right),$$

where \tilde{Z} is a random walk on $V_{\text{even}}(\mathcal{G}^N)$ with the transition probabilities:

$$\tilde{R}^N(x, y) = \frac{|D_N(x, y)|}{D_N(x, x)} \mathbf{1}_{x \neq y}. \tag{2.30}$$

The fact that the even case is much simpler than the odd one can be seen here since $\tilde{R}^N(x, y)$ is actually a transition matrix of a true random walk on $V_{\text{even}}(\mathcal{G}^N)$. Indeed, (see Fig. 5 for an illustration):

- in the bulk of $V_{\text{even}}(\mathcal{G}^N) \setminus V_0(\mathcal{G}^N)$, the walk jumps by ± 2 in each direction with probability $1/4$ each,
- On the boundary $\partial \mathcal{G} \cap V_{\text{even}}(\mathcal{G}^N)$, the walk makes jumps according to the local boundary conditions which are either Dirichlet or Neumann,
- On $V_0(\mathcal{G}^N) \cap V_0(\mathcal{G})$ it may jump horizontally by ± 1 with probability $z^2/(3 + 2z^2)$ or by ± 2 with probability $1/(3 + 2z^2)$, and vertically by $+2$ also with probability $1/(3 + 2z^2)$. This is consistent with the fact that $D(x, x) = 3 + 2z^2$ for $x \in V_0(\mathcal{G})$,
- On $V_0(\mathcal{G}^N) \setminus V_0(\mathcal{G})$, it may jump horizontally by ± 1 (interpreted cyclically) with probability $z^2/(2 + 2z^2)$ or by ± 2 with probability $1/(2 + 2z^2)$. This is consistent with the fact that $D(x, x) = 2 + 2z^2$ for $x \in V_0(\mathcal{G}^N) \setminus V_0(\mathcal{G})$,

As \tilde{R}^N is directly a transition matrix, we do not need to introduce an ε -regularisation or use the Schur complement formula to compute its inverse.

All in all we obtain that

$$D_N^{-1}(u, v) = (-1)^{\Re(v-u)} G_{\text{even}}^N(v, u). \tag{2.31}$$

Now a moment of thought shows that there is no problem in letting $N \rightarrow \infty$ in this expression. This is because the random walk associated with R_N is absorbed on some portion of the boundary $\partial \mathcal{G} \setminus \partial_{\text{free}} \mathcal{G}$, as described in Sect. 1.2.

Hence we deduce that

$$\lim_{N \rightarrow \infty} D_N^{-1}(u, v) = (-1)^{\Re(u-v)} G_{\text{even}}(u, v). \tag{2.32}$$

where $G_{\text{even}}(u, v)$ is the Green’s function on \mathcal{G}^∞ (that is, the graph \mathcal{G}^0 to which infinitely many triangles have been added on either side of V_0) associated with the random walk on \mathcal{G}^∞ whose transition probabilities are given by (2.30).

At the same time, when $N \rightarrow \infty$, the free boundary dimer model on \mathcal{G}^N , restricted to \mathcal{G}^0 , becomes equivalent to a free boundary dimer model on \mathcal{G}^0 where the monomer weights on the extreme vertices (corners) of V_0 have been given the weight $z' > 0$ as in (1.3).

We now summarise the results obtained in this section.

Corollary 2.8 *Consider the free boundary dimer model on \mathcal{G} where the monomer weight $z > 0$ on $V_0(\mathcal{G})$ except at its monomer-corners where the monomer weight is z' as in (1.3). Let K be the associated Kasteleyn matrix, and $D = K^* K$. Then for all $u, v \in V(\mathcal{G})$, we have*

$$D^{-1}(u, v) = \begin{cases} G_{\text{odd}}(u, v) & \text{if } u, v \in V_{\text{odd}}(\mathcal{G}), \\ (-1)^{\mathfrak{N}(v-u)} G_{\text{even}}(u, v) & \text{if } u, v \in V_{\text{even}}(\mathcal{G}), \\ 0 & \text{otherwise.} \end{cases}$$

where $G_{\text{odd}}, G_{\text{even}}$ are the normalised Green’s functions associated with the effective (odd and even) bulk random walks described in (2.26) and (2.30) respectively, normalised by $D(v, v)$.

In particular, the inverse Kasteleyn matrix is given by $K^{-1} = D^{-1}K^*$.

This result implies Theorem 1.1 with the walks Z_{even} and Z_{odd} explicitly defined as above.

2.7 Convergence to potential kernel of the auxiliary walk

In this section we prove the convergence statement from Lemma 2.5. Let us fix $y_1, y_2, z \in \mathbb{Z}$ and their corresponding residue classes $\bar{y}_1, \bar{y}_2, \bar{z} \pmod N$. We introduce the function f defined on the N -cycle, identified with $V_{-1}(\mathcal{G}^N)$,

$$f(\bar{x}) = \sum_{k \in \mathbb{Z}} \alpha(x, y_1 + kN) - \alpha(x, y_2 + kN) - [\alpha(z, y_1 + kN) - \alpha(z, y_2 + kN)],$$

where $x \in \mathbb{Z}$ is chosen so that x has \bar{x} as its residue mod N , and α is the potential kernel defined in (2.11) associated with the random walk (2.7) on \mathbb{Z} . First observe that this function is indeed well defined, in the sense that the series actually converges (this follows from the asymptotics of α , whose gradient is near infinity a constant plus an exponentially decaying function), and in the sense that it does not depend on the choice of the representative x of \bar{x} , by translation invariance of that potential kernel. The last two terms have only been added in order to make the series converge, and they will play no role (and disappear from the statement) once we consider differences of f , i.e. quantities of the form $f(\bar{x}) - f(\bar{x}')$ as required for Lemma 2.5.

Recall the transition probabilities p^N on the N -cycle already described above Lemma 2.4.

Lemma 2.9 *The function f defined above is discrete p^N -harmonic on the N -cycle, except at $\bar{x} = \bar{y}_1$ and $\bar{x} = \bar{y}_2$: that is,*

$$p^N f(\bar{x}) - f(\bar{x}) = \varphi(\bar{x}) = -\mathbf{1}_{\{\bar{x}=\bar{y}_1\}} + \mathbf{1}_{\{\bar{x}=\bar{y}_2\}}. \tag{2.33}$$

It is in the sense of Lemma 2.9 that f is a (difference of) potential kernels.

Proof The proof comes from the fact that α is a potential kernel on \mathbb{Z} for p^∞ , and so satisfies the same equation as (2.33) but on \mathbb{Z} : that is, if $g(x) = \alpha(x, y_1) - \alpha(x, y_2) - [\alpha(z, y_1) - \alpha(z, y_2)]$, $x \in \mathbb{Z}$, then

$$p^\infty g(x) - g(x) = -\mathbf{1}_{\{x=y_1\}} + \mathbf{1}_{\{x=y_2\}}.$$

To actually deduce (2.33) from this, we truncate the infinite series defining f and use the exponential decay of the difference of gradients of α to show that the boundary contributions tend to zero. Details are left to the reader. \square

Proof of Lemma 2.5 As a consequence, if X_n is a random walk on the N -cycle with transition probabilities p^N , and if $A_n = \sum_{i=1}^n \varphi(X_n)$, then

$$M_n = f(X_n) - A_n \tag{2.34}$$

is a $\mathbb{P}_{\bar{x}}$ -martingale. Let T_ε denote an independent geometric random variable with success probability p_ε (which corresponds to the time at which the walk would be killed, if it has probability p_ε to be killed at each step). Then $(M_{n \wedge T_\varepsilon})_{n \geq 0}$ is also a $\mathbb{P}_{\bar{x}}$ -martingale, and hence, applying the optional stopping theorem (which is allowed by combining dominated and monotone convergence theorems, since the state space is finite)

$$\begin{aligned} f(\bar{x}) &= \mathbb{E}_{\bar{x}}(M_0) = \mathbb{E}_{\bar{x}}(M_{T_\varepsilon}) \\ &= \mathbb{E}_{\bar{x}}[f(X_{T_\varepsilon})] + \left(\mathbb{E}_{\bar{x}}[L_{T_\varepsilon}(\bar{y}_1)] - \mathbb{E}_{\bar{x}}[L_{T_\varepsilon}(\bar{y}_2)] \right) \end{aligned}$$

where $L_n(\bar{y}) = \sum_{i=1}^n 1_{\{X_n = \bar{y}\}}$ denotes the local time at $\bar{y} \in \mathbb{Z}/N\mathbb{Z}$ of the random walk X . Hence in particular, $\mathbb{E}_{\bar{x}}[L_{T_\varepsilon}(\bar{y})] = g^{N,\varepsilon}(\bar{x}, \bar{y})$ is the Green's function defined in (2.20) and appearing in the statement of Lemma 2.5.

Now take another starting point, \bar{x}' , and consider the difference $f(\bar{x}) - f(\bar{x}')$. Then

$$\begin{aligned} f(\bar{x}) - f(\bar{x}') &= \mathbb{E}_{\bar{x}}[f(X_{T_\varepsilon})] - \mathbb{E}_{\bar{x}'}[f(X_{T_\varepsilon})] \\ &\quad + \left(g^{N,\varepsilon}(\bar{x}, \bar{y}_1) - g^{N,\varepsilon}(\bar{x}, \bar{y}_2) - [g^{N,\varepsilon}(\bar{x}', \bar{y}_1) - g^{N,\varepsilon}(\bar{x}', \bar{y}_2)] \right) \end{aligned}$$

The third term (in brackets) on the right-hand side is precisely the quantity we wish to take the limit of as $\varepsilon \rightarrow 0$. The left hand side is precisely the conclusion of Lemma 2.5 (note that the terms involving z cancel). Hence in order to conclude, it remains to show that for arbitrary \bar{x}, \bar{x}' ,

$$\mathbb{E}_{\bar{x}}[f(X_{T_\varepsilon})] - \mathbb{E}_{\bar{x}'}[f(X_{T_\varepsilon})] \rightarrow 0 \tag{2.35}$$

as $\varepsilon \rightarrow 0$. However, this is immediate: the walks started from \bar{x} and \bar{x}' can be coupled in finite time almost surely. Since $T_\varepsilon \rightarrow \infty$ in probability as $\varepsilon \rightarrow 0$, and the state space of these walks (i.e., the N -cycle) is finite, we see that (2.35) holds. This concludes the proof of Lemma 2.5. \square

3 Infinite volume limit

In the previous section we showed that D_N^{-1} (and hence K_N^{-1}) has a limit as $N \rightarrow \infty$ which is given in terms of two Green's functions G_{odd} and G_{even} associated to random walks on $V_{\text{odd}}(\mathcal{G})$ and $V_{\text{even}}(\mathcal{G})$ which may jump along $V_1(\mathcal{G})$ and $V_0(\mathcal{G})$,

and with various boundary conditions (Dirichlet or mixed Neumann–Dirichlet) on $\partial\mathcal{G} \setminus (V_0(\mathcal{G}) \cup V_1(\mathcal{G}))$. Let us also denote these Green's functions by $G_{\text{odd}}^{\mathcal{G}}$ and $G_{\text{even}}^{\mathcal{G}}$ to emphasize their dependence on \mathcal{G} .

The purpose of this section is to take an infinite volume limit as \mathcal{G} tends to the upper half-plane. In this limit the Green's functions $G_{\text{odd}}^{\mathcal{G}}$ and $G_{\text{even}}^{\mathcal{G}}$ diverge (corresponding to the fact that the limiting bulk effective random walk is recurrent). However, we can still make sense of its potential kernel. Hence the inverse Kasteleyn matrix, which is obtained as a derivative of these Green's functions, has a well defined pointwise limit.

The argument for this convergence as \mathcal{G} increases to the upper half plane are essentially the same for both the odd and even walks. As will be clear from the proof below, the arguments rely only on the fact that (a) the two walks coincide with the usual simple random walk (with jumps of size 2) away from the real line, (b) they are reflected on the real line with some jump probabilities that decay exponentially fast with the jump size (in fact, in the even case the jumps are bounded), and (c) they can 'switch colour' with positive probability along the real line. This terminology will be explained below. For these reasons, and in order to avoid unnecessarily cumbersome notation, we focus in this section solely on the *odd walk* (the argument works literally in the same way for the even case, and can in fact be made a little easier since the jumps on the real line have bounded size).

3.1 Construction of the potential kernel in the infinite volume setting

We write Γ for the weighted graph corresponding to the odd effective random walk. Thus, the vertex set V of Γ can be identified (after translation so that $V_1 \subset \mathbb{R}$) with $(\mathbb{Z} \times 2\mathbb{Z}) \cap \mathbb{H}$ and its edges E are those of $(2\mathbb{Z})^2$, plus those of $(2\mathbb{Z} + 1) \times (2\mathbb{Z})$, plus additional edges connecting these two lattices along the real lines. In reality, it will be easier to consider a symmetrised version of Γ obtained by taking the vertex set to be $V \cup \bar{V}$ and the edges to be $E \cup \bar{E}$, where \bar{V} and \bar{E} are the complex conjugates of V and E . We will still denote this graph by Γ . Throughout this and the next section the random walks we will consider will take values in this symmetrised graph. Note that Γ is not locally finite: any vertex on the real line has infinite degree, but the total weight out of every vertex is finite (and is equal to 1). We recall that when away from the real line, the random walk on Γ looks like simple random walk on the square lattice *up to factor 2*: the transitions from a point $x \in \mathbb{Z}^2$ away from \mathbb{R} are to the four points $x \pm 2e_1$ or $x \pm 2e_2$, where (e_1, e_2) is the standard basis of \mathbb{Z}^2 . On the real line, the effective random walk can make jumps of any size, but the jumps are symmetric and the transition probabilities have an exponential tail. Note that the odd effective random walk only jumps between vertices of the same colour in the bulk, and can possibly change colour only on the real line. In the current section, we will also use the word **class** to denote the notion of colour. Finally, we say that two vertices in Γ have the same **parity** (or **periodicity**) if the differences of their vertical and horizontal coordinates are multiples of 4.

Our first goal will be to show that differences of Green's functions evaluated at two different vertices of the same class for the walk killed when leaving a large box, converge (when the box tends to infinity) to differences of the **potential kernel** of the

walk on the infinite graph Γ . Our first task will be to define this potential kernel. For the usual simple random walk on \mathbb{Z}^2 this is an easy task because the asymptotics of the transition probabilities are known with great precision. In turn this is because simple random walk can be written as a sum of i.i.d. random variables making it possible to use tools from Fourier analysis: see Chapter 4 of [30] for a thorough introduction. The walk on Γ obviously does not have this structure, and in fact it seems that there are few general tools for the construction of the potential kernel for walks on a planar graph beyond the i.i.d. case. The coupling arguments we introduce below may therefore be of independent interest.

Let P denote the transition matrix of simple random walk on Γ , and let $\tilde{P} = (I + P)/2$ be that of the associated lazy chain (from now on \sim will be used to denote objects related to the lazy chain). The rationale for considering this version is that, on the one hand, it gets rid of periodicity issues, while on the other hand, it only modifies the Green's function by a constant factor: e.g., on a transient graph, $\tilde{G}(x, y) = 2G(x, y)$ for any x, y , if G and \tilde{G} are the corresponding Green's functions (this is because the jump chains are the same, and the lazy chain stays on average twice as long at any vertex as the non-lazy chain).

The basic idea for the definition of the potential kernel will be the following. Let X and X' denote (lazy) random walks started respectively from two vertices x and x' of the same class, and suppose that they are coupled in a certain way so that after a random time T (which may be infinite), X and X' remain equal forever on the event that $T < \infty$: that is,

$$X_{T+s} = X'_{T+s}, \quad s \geq 0. \quad (3.1)$$

We will define a coupling (its precise definition will be given below) that depends on a time-parameter t such that for this particular value of t ,

$$\mathbb{P}(T > t) \lesssim (\log t)^w t^{-1/2} \quad (3.2)$$

for some $w > 0$ whose value will not be relevant. (Note that this inequality should not be understood as saying something about the tail of T , since T depends on t ; indeed T might be infinite with positive probability). In fact, a much weaker control of the form $\mathbb{P}(T > t) \lesssim t^{-\varepsilon}$, for some $\varepsilon > 0$, would be sufficient for the definition of the potential kernel alone, as will be apparent from the argument below. We however insist on (3.2) in order to get good a priori bound on the potential kernel (see Proposition 3.4). As we will see, the goal of this coupling will be to compare $\tilde{p}_t(x, o)$ to $\tilde{p}_t(x', o)$ which is why T is allowed to depend on t , and why we only require T to be less than t with high probability (but we do not care what happens on the event $\{T > t\}$). Here and later on, $\tilde{p}_t(\cdot, \cdot)$ is the transition probability of the lazy chain whose transition matrix is \tilde{P} defined above, and o denotes an arbitrary fixed vertex.

We first argue that we can get a good a priori control on the transition probabilities $\tilde{p}_t(x, o)$. Let $A \subset \mathbb{Z} \times 2\mathbb{Z}$ be a finite set. By ignoring the long range edges which may leave A through the real line, and using the standard discrete isoperimetric inequality on \mathbb{Z}^2 (Loomis-Whitney inequality, Theorem 6.22 in [31]) it is clear that

$$\sum_{x \in A, y \in A^c} w_{x,y} \gtrsim |A|^{1/2}$$

where $w_{x,y}$ is the weight of the edge (x, y) in Γ . This means that Γ satisfies the two-dimensional isoperimetric inequality (I_2) (we here use the notation of [3]). Consequently, by Theorem 3.7, Lemma 3.9 and Theorem 3.14 of [3], Γ satisfies the two-dimensional **Nash inequality**, (N_2) (see [3] for a precise statement of that inequality). Therefore, if $q_s^x(\cdot)$ denote the transition probabilities of the continuous time walk on Γ , normalised by its invariant measure, we have by Theorem 4.3 in [3] that

$$q_s^x(x) \lesssim 1/s,$$

and since q_s^x is maximised on the diagonal, we deduce that

$$\tilde{p}_s(x, o) \lesssim 1/s, \tag{3.3}$$

where the implied constant is uniform in x, o and $s \geq 1$.

Now suppose we have a coupling satisfying (3.1) and (3.2). We will explain why this implies that

$$\sum_{t=0}^{\infty} (\tilde{p}_t(x, o) - \tilde{p}_t(x', o)) \tag{3.4}$$

converges. We couple the walks starting from x, x' according to (3.1). Obviously, on the event $\{T \leq t/2\}$, $X_t = o$ if and only if $X'_t = o$, and thus

$$\begin{aligned} |\tilde{p}_t(x, o) - \tilde{p}_t(x', o)| &\leq 2\mathbb{P}(T \geq t/2) \max_y \tilde{p}_{t/2}(y, o) \\ &\lesssim t^{-3/2}(\log t)^w \end{aligned} \tag{3.5}$$

which is summable, whence the series (3.4) converges.

Definition 3.1 We set

$$\tilde{a}(x, o) - \tilde{a}(x', o) = - \sum_{t=0}^{\infty} (\tilde{p}_t(x, o) - \tilde{p}_t(x', o)),$$

where $\tilde{p}_t(\cdot, \cdot)$ is the transition probability of the lazy chain. By convention we define $\tilde{a}(o, o) = 0$ and so this recipe may be used to define $\tilde{a}(x, o)$ provided that x and o are of the same class (by summing increments along a given path from x to o). (As the choice of a path from x to o does not matter before the limit in the series is taken, this is well defined.) Since x and o are arbitrary vertices of the same class, this defines $\tilde{a}(\cdot, \cdot)$ everywhere on this class.¹

If also (3.2) holds for *one* pair x, x' not of the same class, then this defines $\tilde{a}(\cdot, \cdot)$ over the entire graph.

¹ The arguments in this section rely on thinking of $\tilde{a}(\cdot, \cdot)$ as a function of the first variable while the second is frozen, which is why we prefer to use x for the first variable and o for the second. In the next section, both variables will start playing a more symmetric role and we will switch to x and y .

Note also that due to the fact that $\pi(x) = 1$ is a *constant* reversible measure on Γ (hence $\tilde{p}_k(x, y) = \tilde{p}_k(y, x)$), the potential kernel is symmetric: $\tilde{a}(x, y) = \tilde{a}(y, x)$ for any x, y . We will not however need this property in the following.

In the next subsection we describe a concrete coupling which will be used for the construction of the potential kernel. We call this the **coordinatewise mirror coupling**, which is a variation on a classical coupling for Brownian motion in \mathbb{R}^d . We will then use this coupling again to obtain *a priori* estimates on the potential kernel.

Before describing this coupling and justifying (3.2), we first state and prove a lemma which will be useful in many places in the the following and which gives a subdiffusive estimate on the walk. Let dist denote the usual ℓ^1 distance (graph distance) on \mathbb{Z}^2 .

Lemma 3.2 *Let x be a vertex of Γ and let $T_R = \inf\{n \geq 0 : \text{dist}(X_n, x) \geq R\}$. Then for every $c_1 > 0$ there exists $c_2 > 0$ such that for any $n \geq 1$, and for any $R \geq c_1\sqrt{n} \log n$,*

$$\mathbb{P}(T_R \leq n) \lesssim \exp(-c_2(\log n)^2).$$

Proof One possibility would be to use a result of Folz [14] (based on work of Grigor'yan [17] in the continuum) which shows that an on-diagonal bound on the heat kernel $p_t(x, x)$ and $p_t(y, y)$ implies a Gaussian upper bound on the off-diagonal term $p_t(x, y)$. However, it is more elementary to use the following martingale argument. We may write $X_n = (u_n, v_n)$ in coordinate form. Since (v_n) is a lazy simple random walk on the integers, the proof is elementary in this case (and of course also follows from the more complicated estimate below). We therefore concentrate on bounding $\sum_{i=1}^n \mathbb{P}(|u_i| \geq R)$. We bound $\mathbb{P}(|u_i| \geq R)$ for $1 \leq i \leq n$ as follows: either there is one jump larger than say $K = (\log n)^2$ by time n (this has probability at most $n \exp(-c(\log n)^2)$ by a union bound and exponential tail of the jumps) or if all the jumps are less than K , then u coincides with a martingale \bar{u} such that all its jumps are bounded by K in absolute value: indeed, we simply replace every jump of u greater than K in absolute value by a jump of the same sign and of length K . Since the jump distribution (2.6) is symmetric, the resulting sum \bar{u}_n is again a martingale. Furthermore, \bar{u}_n is a martingale with bounded jumps. We may apply Freedman's inequality [15, Proposition (2.1)] to it which implies (since the quadratic variation of \bar{u} at time $1 \leq i \leq n$ is bounded by $b \lesssim n$),

$$\mathbb{P}(|\bar{u}_i| \gtrsim \sqrt{n} \log n) \lesssim \exp\left(-c \frac{n(\log n)^2}{(\log n)^2 \sqrt{n} \log n + n}\right) \lesssim \exp(-c(\log n)^2). \quad (3.6)$$

The result follows by summing over $1 \leq i \leq n$. \square

3.2 Coordinatewise mirror coupling

Let x, x' be two vertices of the graph Γ of the same class, and let \tilde{X}, \tilde{X}' be two (lazy) effective random walks started from x and x' respectively. In the coupling we will describe below, it will be important to first fix the vertical coordinate (stage 1). The coupling ends when we also fix the horizontal coordinate (stage 4). In between, we

have two short stages (possibly instantaneous), where we make sure the class is correct (stage 2) followed by a so-called “burn-in” phase where the walks get far away from the real line in parallel (stage 3). This depends on a parameter r , which is a free choice. (When we prove (3.2) we will choose r to be slightly smaller by logarithmic factors than \sqrt{t}).

We need to do so while respecting the natural **parity** (i.e., periodicity) of the coordinates we are trying to match. We will use the laziness to our advantage in order to deal with the potential issues arising from the walks not being of the same parity.

Note the following important property of \tilde{P} . At each step, the walk moves with probability $1/2$. Conditionally on moving, the horizontal coordinate moves with probability $1/2$, and otherwise the vertical coordinate moves (and in that case it is equally likely to go up or down by two); since we symmetrised Γ note also that $\tilde{p}(x, x + y)$ and $\tilde{p}(x, x - y)$ are always equal, for all $x, y \in \mathbb{Z}^2$ (i.e., the jump distribution is symmetric). We will need a fair coin C to decide which of the two Coordinates moves (if moving), and another fair coin L to decide whether the walk is Lazy or moves in this step.

Stage 1: vertical coordinate. Suppose that $\tilde{X}_t = (u_t, v_t)$, $\tilde{X}'_t = (u'_t, v'_t)$ are given. We now describe one step of the coupling. If $v_t = v'_t$ move to stage 2. If $v_t \neq v'_t$ then we consider the following two cases. In any case, we start by tossing C. If heads, then we plan for both \tilde{X} and \tilde{X}' to move their horizontal coordinates, and if tails, for both their vertical coordinates.

1. Case 1: $v_t - v'_t = 2 \pmod{4}$. Suppose C is tails so the parity of vertical coordinate has a chance to be improved. Then we toss L. Depending on the result, one stays put and the other moves, or vice versa (either way the vertical coordinates are of the same parity after, and will stay so forever after). If instead C was heads, so horizontal coordinate moves for both walks, then they move simultaneously or stay put simultaneously, and move independently of one another if at all.
2. Case 2: $v_t - v'_t = 0 \pmod{4}$. Suppose C is tails, so the vertical coordinates have a chance to be improved or even matched. Then we toss L and according to the result they both move simultaneously or stay put simultaneously. If moving at all, we declare the change in v_t and the change in v'_t to be opposite one another: thus, $v_{t+1} = v_t \pm 2$ with equal probability, whence $v'_{t+1} = v'_t \mp 2$. If however C is heads (so the horizontal coordinates move), then the walks move simultaneously or stay put simultaneously, and move independently of one another if at all.

We leave it to the reader to check that this is a valid coupling (all moves are balanced and according to the transition probabilities P if moving, and altogether each walk moves or stays put with probability $1/2$ as desired). As mentioned, once the parity of the vertical coordinates of the walks is matched (meaning the difference in vertical coordinates is even), it will remain matched forever.

Note also that once the vertical parity is matched ($v_t - v'_t = 0 \pmod{4}$), conditionally on the vertical coordinate moving (which is then the case for both walks simultaneously), the direction of movements is opposite: in other words, the positions of the vertical coordinates v_t and v'_t throughout time and until they match are mirrors of one another, with a reflection axis which is a horizontal line L_1 . This line can be

described as having a vertical coordinate equal to the average of v_t and v'_t at the first time t that the parity of v_t and v'_t matches (note that L_1 goes via $(2\mathbb{Z})^2$). In particular, the two coordinates v_t and v'_t will match after the first hitting time T_1 of the line L_1 . By the end of the first stage, the two walks sit on the same horizontal line. This will remain so forever.

Stage 2: setting class and/or periodicity. We now aim to match the horizontal coordinate. If also $u_t = u'_t$ the coupling is over and we let $\tilde{X}'_{t+1} = \tilde{X}_{t+1}$ chosen according to $\tilde{P}(\tilde{X}_t, \cdot)$. However the two walks might not be in the same class at that point, even if they started in the same class at the beginning of stage 1 (their class might change if one hits the real line but not the other during that stage). During stage 2, we will make sure the walks become of the same class if they were not at the beginning of that stage (amounting to $u_t - u'_t$ even), and we will also make sure that they become of the same “parity” or “periodicity”, meaning $u_t - u'_t = 0 \pmod{4}$. If that is the case already at the beginning of this stage, we can immediately move on to the next stage.

Otherwise, as before, suppose that $\tilde{X}_t = (u_t, v_t)$, $\tilde{X}'_t = (u'_t, v'_t)$ are given, and suppose that $v_t = v'_t$. (In particular, $v_t = 0$ if and only if $v'_t = 0$.) We proceed as follows. As before, in any case we start by tossing C . If heads, then we plan for both \tilde{X} and \tilde{X}' to move their horizontal coordinates, and if tails, for both their vertical coordinates. In the latter case, we will use the same moves for both \tilde{X} and \tilde{X}' , so we only describe what happens if the move is horizontal.

1. If $u_t - u'_t$ is odd, and $v_t = v'_t \neq 0$, then the walks move simultaneously and in parallel.
2. In all other situations, one walk will stay put while the other moves, or vice-versa, depending on the outcome of L .

We make a few comments. First, note that with every visit to the real line there is a fixed positive chance to have $u_t - u'_t = 0 \pmod{4}$ and hence to end this stage. Also, if $u_t - u'_t$ is even to begin with, then there is also a fixed positive chance to end the stage right away.

Stage 3: burn-in. In stage 3 of the coupling, we let the walks evolve in parallel (i.e., with the same jumps) until they are at distance r from the real line. We will later choose r as a function of t (see (3.7)), which explains our comment under (3.2) that T depends on t . This is a valid choice of coupling since they will hit the real line simultaneously. At the end of stage 2, the walks are on the same horizontal line and of the same “periodicity” meaning that they are $0 \pmod{4}$ apart. This will remain so until the end of stage 3.

Stage 4: horizontal coordinate. As before, suppose that $\tilde{X}_t = (u_t, v_t)$, $\tilde{X}'_t = (u'_t, v'_t)$ are given, and suppose that $v_t = v'_t$. (In particular, $v_t = 0$ if and only if $v'_t = 0$.) If also $u_t = u'_t$ we let $\tilde{X}'_{t+1} = \tilde{X}_{t+1}$ chosen according to $\tilde{P}(\tilde{X}_t, \cdot)$. Otherwise we proceed as follows; we only describe a way of coupling the walks until hitting the real line; if coupling has not occurred before then we say that $T = \infty$. As before, in any case we start by tossing C . If tails, we let both walk evolve vertically in parallel. Otherwise, the walks will move their horizontal coordinates or stay put simultaneously depending on

the result of L. If both walks move horizontally, then let u_t and u'_t move in opposite manners, i.e., $u_{t+1} - u_t = -(u'_{t+1} - u'_t)$. This is possible by symmetry of the jump distribution P (even on the real line).

Again, we leave it to the reader to check that what we have described in stages 2,3 and 4 forms a valid coupling. We note that any movement in the vertical coordinate is replicated across both walks, whatever the cases, and so the match created in stage 1 is never destroyed. Note also that once the walks are $0 \bmod 4$ apart, this remains the case until hitting the real line. Therefore the movement of the horizontal coordinates of both walks in stage 2 of this coupling will also be mirror off one another, with the mirror being a vertical line L_2 whose horizontal coordinate is the average of u_t and u'_t at the end of stage 3. We call T_1, \dots, T_4 the end of each four stage respectively (with T_4 being infinity if the walks hit the real line first).

3.3 Suitability of coupling (proof of (3.2))

In order to use the above coupling to construct the potential kernel of the walk on Γ , we need to verify two points. We will consider two cases: the main one is that x and x' are of the same class and $\text{dist}(x, x') = 2$. The other case is if x, x' are on the real line and $\text{dist}(x, x') = 1$. By Definition 3.1, these two cases allow us to define the potential kernel over the entire graph. We will focus on the first case since it is a bit more involved than the second (which can be checked in a similar manner). We will first need to verify (3.2), which requires that the two walks coincide with high probability at time t .

We will check that each stage lasts less than $t/4$ with overwhelming probability (meaning with error probability satisfying (3.2)).

Stage 1. We may assume without loss of generality that $v_0 - v'_0 = 0 \bmod 4$ since otherwise it takes a geometric number of attempts until that is the case. Note then that $\mathbb{P}(T_1 > t/4)$ is bounded by the probability that the random walk avoids the (horizontal) reflection line L_1 of stage 1 for time $t/4$. As the vertical coordinate performs a lazy simple random walk on the integers (with laziness parameter $3/4$) this is bounded by the probability that a random walk on the integers starting from 1 (or more generally a random value with geometric tails, as discussed above) avoids 0 for at least $\gtrsim t$, which is bounded by $\lesssim 1/\sqrt{t}$ by gambler's ruin arguments (see e.g. Proposition 5.1.5 in [30]).

Stage 2. Let $k = \text{dist}(x, \mathbb{R}) (= |v_0|)$. If the walks remained of the same class during the first stage, then stage 2 is over in a time which has a geometric tail so (3.2) holds trivially. On the other hand, if they did change class during the first stage, it is necessary to hit the real line again (and then wait for an extra time with geometric tail).

At the end of stage 1, the walk is on the reflection line L_1 which has vertical coordinate $v + O(1)$ and so is again at distance $k + O(1)$ from the real line. Let $T_{\mathbb{R}}$ denote the hitting time of \mathbb{R} . Then by Proposition 5.1.5 in [30] again,

$$\mathbb{P}(T_{\mathbb{R}} > t/4) \lesssim \frac{k}{\sqrt{t}}.$$

On the other hand, the probability that \tilde{X}_{T_1} and \tilde{X}'_{T_1} changed class during the first phase is bounded by $\lesssim 1/k$ again by gambler’s ruin (since it requires touching the real line before the reflection line L_1), and so

$$\mathbb{P}(T_2 - T_1 > t/4; \tilde{X}_{T_1} \approx \tilde{X}'_{T_1}) \lesssim \frac{1}{k} \times \frac{k}{\sqrt{t}} = \frac{1}{\sqrt{t}},$$

where \approx denotes being of different class. This implies (3.2) for $T_2 - T_1$.

Stage 3. Here we will need to choose the parameter r appropriately. We will take it to be

$$r = \frac{\sqrt{t}}{(\log t)^b}, \tag{3.7}$$

where $b > 0$ can be chosen as desired. Note that every r^2 units of time, if the walk starts in the strip S of width r around the real line, it has a positive probability, say p , of leaving S (where p does not depend on the starting point of the walk). Thus for $j \geq 1$,

$$\mathbb{P}(T_3 - T_2 > jr^2) \leq (1 - p)^j.$$

Hence

$$\mathbb{P}(T_3 - T_2 > t/4) \leq \exp(-ct/r^2)$$

so that if $b > 1$ is any number, the right hand side above is $\lesssim 1/\sqrt{t}$, as desired.

Stage 4. To prove the corresponding bound in stage 4, we need the following lemma which shows (up to unimportant logarithmic terms) the horizontal displacement accumulated in the first stage has a Cauchy tail. This corresponds of course to the well known fact that the density of Brownian motion when it hits a fixed line has exactly a Cauchy distribution.

Lemma 3.3 *We have*

$$\mathbb{P}(\sup_{t \leq T_1} |u_t - u_0| \geq k) \lesssim \frac{(\log k)^2}{k}. \tag{3.8}$$

(The factor of $(\log k)^2$ is not optimal in the right hand side of (3.8) but is sufficient for our purposes.) We now use this to derive a bound for $\mathbb{P}(T_4 - T_3 > t/4)$. From the construction of the coupling and Lemma 3.3, we see that at the beginning of stage 4, the walk is at a distance from the (vertical) reflection line L_4 which has the same tail as in Lemma 3.3 (this is because the additional discrepancy accumulated during stage 2 is easily shown to have geometric tail). Let us condition on everything before time T_3 , and call k the distance of the walk X at time T_3 to the reflection line. Let $T_{\mathbb{R}}$ denote the hitting time of the real line and let T_L denote the hitting time of the reflection line L_4 . Set $s = t/4$ for convenience. Then we can bound the tail of T_L in terms of the usual simple random walk on the lattice (without extra jumps on the real line). Indeed, until time $T_{\mathbb{R}}$, the walk coincides with the usual **lazy** simple random walk on the square lattice. Writing \mathbb{Q} for the law of the latter random walk, we have

$$\mathbb{P}(T_4 - T_3 > s \mid \mathcal{F}_{T_3}) \leq \mathbb{P}(T_L > s, T_{\mathbb{R}} > s \mid \mathcal{F}_{T_3}) + \mathbb{P}(T_L > T_{\mathbb{R}}, T_{\mathbb{R}} \leq s \mid \mathcal{F}_{T_3})$$

$$\begin{aligned} &\leq \mathbb{Q}(T_L > s, T_{\mathbb{R}} > s) + \mathbb{Q}(T_{\mathbb{R}} < T_L) \\ &\lesssim \mathbb{Q}(T_L > s) + \frac{k}{r} \\ &\lesssim k\left(\frac{1}{\sqrt{t}} + \frac{1}{r}\right). \end{aligned}$$

To go from the second line to the third line, we used that the walk starts at distance r from the real line and Proposition 5.1.5 in [30], and to go the last line we also used that same result. Taking expectations (we only use the above bound if $k \leq r$ so that the right hand side is less than one, and we use the trivial bound 1 for the probability on the left-hand side otherwise), we see that

$$\mathbb{P}(T_4 - T_3 > t/4) \leq \mathbb{E}(X\mathbf{1}\{X \leq r\})\left(\frac{1}{\sqrt{t}} + \frac{1}{r}\right) + \mathbb{P}(X > r),$$

where X has a tail bounded by Lemma 3.3. By Fubini’s theorem,

$$\mathbb{P}(T_4 - T_3 > t/4) \lesssim \frac{(\log t)^{3+b}}{\sqrt{t}} \tag{3.9}$$

so we get (3.2) with $w = 3 + b$. Since $b > 1$ is arbitrary, $w > 4$ is arbitrary.

It therefore remains to give the proof of Lemma 3.3.

Proof of Lemma 3.3 Let $L = L_1$ be the (horizontal) reflection line. We wish to show that $\mathbb{P}(\sup_{t \leq T_L} |u_t - u_0| \geq k) \lesssim (\log k)^2/k$. Without loss of generality we assume that $v_0 < v'_0$ so \tilde{X} starts below L , and $u_0 = 0$. Let L' be a line parallel to L below L , at distance A from it, where $A = \lfloor k/(\log k)^2 \rfloor$. Let S denote the infinite strip in between these two lines. Let $T = T_L$ denote the hitting time of L and let T' denote the hitting time of L' , and let $\tau = T \wedge T'$ denote the time at which the walk leaves the inside of the strip S . Let T_k denote the first time at which $|u_t| \geq k$. Then

$$\begin{aligned} \mathbb{P}(\sup_{t \leq T} |u_t| \geq k) &\leq \mathbb{P}(T' < T, \sup_{t \leq T} |u_t| \geq k) + \mathbb{P}(T' > T, \sup_{t \leq T} |u_t| \geq k) \\ &\leq \mathbb{P}(T' < T) + \mathbb{P}(T_k \leq \tau). \end{aligned}$$

Now, the event $T' < T$ concerns only the vertical coordinate which (ignoring the times at which it doesn’t move which are irrelevant here) is simple random walk on \mathbb{Z} . Hence $\mathbb{P}(T' < T) = 1/A \lesssim (\log k)^2/k$ by the gambler’s ruin estimate in one dimension for simple random walk.

It remains to show that $\mathbb{P}(T_k \leq \tau) = o((\log k)^2/k)$. We split the event into two events, and show both are overwhelmingly unlikely. We observe that for $T_k \leq \tau$ to occur, one of the following two events must occur: either (i) $T_k \leq n := k^2/(\log k)^2$, or (ii) $\tau > n$. Let E_1 be the first event and let E_2 be the second one. Then by Lemma 3.2,

$$\mathbb{P}(E_1) \lesssim \exp(-c(\log k)^2) \tag{3.10}$$

for some constant $c > 0$. As for the second event E_2 , we note that every A^2 units of time there is a positive chance to leave S (this is a trivial consequence of the fact that

the vertical coordinate is lazy random walk on \mathbb{Z} , with the laziness parameter equal to $1/2 + 1/4 = 3/4$, hence

$$\mathbb{P}(E_2) \leq \exp(-c \frac{n}{A^2}) = \exp(-c(\log k)^2)$$

since $A = k/(\log k)^2$ and $n = k^2/(\log k)^2$. Thus

$$\mathbb{P}(T_k \leq \tau) \leq \mathbb{P}(E_1) + \mathbb{P}(E_2) \lesssim \exp(-c(\log k)^2) = o(k^{-1}),$$

and (3.8) follows. □

3.4 A priori estimate on the gradient of the potential kernel

The purpose of this section is to show the following estimate. This will be useful both for proving that Green’s functions differences converge to differences of the potential kernel in the limit of large box \mathcal{G}_n , $n \rightarrow \infty$, but also as an input to the proof of the scaling limit result for the height function, where such an a priori estimate is needed for the inverse Kasteleyn matrix.

Proposition 3.4 *Let o, x, x' be any vertices of Γ such that $\text{dist}(o, x) = R$ and such that x, x' are of the same class with $\text{dist}(x, x') = 2$. Then for $R \geq 2$,*

$$|\tilde{a}(x, o) - \tilde{a}(x', o)| \lesssim \frac{(\log R)^c}{R}$$

where $c = w + 2 > 2$, and $w > 0$ is as in (3.2).

Proof Let $t = R^2/(\log R)^2$. We note that for $s \leq t$,

$$\begin{aligned} \tilde{p}_s(x, o) &\leq \mathbb{P}_o(d(X_s, o) \geq R) \\ &\leq \mathbb{P}_x(|u_s - u| \geq R/2) + \mathbb{P}(|v_s - v| \geq R/2). \end{aligned}$$

Both terms are easily estimated. The first term is estimated by Lemma 3.2 which shows it is bounded $\exp(-(\log t)^2)$. The same estimate holds (and is of course easier) for the vertical coordinate, since this is simply lazy simple random walk (with laziness parameter $3/4$). Naturally this argument also holds with x' in place of x . Thus

$$\left| \sum_{s=0}^t \tilde{p}_s(x, o) - \tilde{p}_s(x', o) \right| \lesssim \exp(-(\log R)^2). \tag{3.11}$$

On the other hand, for $s \geq t$, we recall that

$$|\tilde{p}_s(x, o) - \tilde{p}_s(x', o)| \leq s^{-3/2}(\log s)^w, a$$

by (3.2). Summing over $s \geq t$,

$$\left| \sum_{s=t}^{\infty} \tilde{p}_s(x, o) - \tilde{p}_s(x', o) \right| \lesssim t^{-1/2} (\log t)^w \lesssim R^{-1} (\log R)^{w+2}. \tag{3.12}$$

Combining with (3.11) this finishes the proof with $c = w + 2$ as desired. □

3.5 Convergence of Green’s function differences to gradient of potential kernel

Let $B_R = B(0, R)$, and x, o be vertices of Γ as before. Define the unnormalised Green’s function

$$\tilde{G}_R(x, o) = \mathbb{E}_x \left(\sum_{n=0}^{\infty} \mathbf{1}_{\{\tilde{X}_n = o, \tau_R > n\}} \right),$$

where τ_R is the first time that the (lazy) walk \tilde{X} leaves B_R . We will prove the following proposition:

Proposition 3.5 *As $R \rightarrow \infty$, for any fixed x, x' of the same class, and any fixed o ,*

$$\tilde{G}_R(x, o) - \tilde{G}_R(x', o) \rightarrow -(\tilde{a}(x, o) - \tilde{a}(x', o)).$$

As a consequence, the same convergence is true also for the nonlazy walk X instead of \tilde{X} .

The proof is based on ideas similar to Proposition 4.6.3 in [30]. We first recall the following lemma which (in the case of finite range irreducible symmetric random walk would be Proposition 4.6.2 in [30]):

Lemma 3.6 *For any x, o vertices of Γ , we have*

$$\tilde{G}_R(x, o) = \mathbb{E}_x (\tilde{a}(\tilde{X}_{\tau_R}, o)) - \tilde{a}(x, o).$$

Proof The proof is simply an application of the optional stopping theorem for the martingale $M_n = \tilde{a}(\tilde{X}_n, o) - L_n^{\tilde{X}}(o)$, where

$$L_n^{\tilde{X}}(o) = \sum_{m=0}^n \mathbf{1}_{\{\tilde{X}_m = o\}} \tag{3.13}$$

denote the local time of \tilde{X} at o by time n . The application of the optional stopping is first done at time $\tau_R \wedge n$ which is bounded. The limit when $n \rightarrow \infty$ can be taken by dominated convergence for the first term and monotone convergence for the second. In fact, for the application of the dominated convergence theorem, one must be a little more careful than with simple random walk, since when leaving B_R , there is an unbounded set of possibilities for \tilde{X}_{τ_R} . However the jump probabilities decay

exponentially and $\tilde{a}(x, o)$ grows at most like $(\log |x - o|)^c$ as $x \rightarrow \infty$ by Proposition 3.4 (note here that x, x' and o is fixed while $R \rightarrow \infty$). This makes the application of the dominated convergence justified. We give full details of this argument for the sake of completeness.

To this end, note that $|\tilde{X}_{n \wedge \tau_R}| \leq 2|\tilde{X}_{\tau_R}|$ almost surely. Let $B'_R \subseteq B_R$ be the set of vertices connected by an edge to the outside of B_R , and let τ'_R be the first hitting time of B'_R . By the strong Markov property we get

$$\mathbb{E}_x(|\tilde{X}_{\tau_R}|) = \sum_{z \in B'_R} \mathbb{P}_x(\tilde{X}_{\tau'_R} = z) \mathbb{E}_z(|\tilde{X}_{\tau_R}|)$$

and

$$\mathbb{E}_z(|\tilde{X}_{\tau_R}|) \leq \mathbb{P}_z(\tau_R = 1) \mathbb{E}_z(|\tilde{X}_{\tau_R}| \mid \tau_R = 1) + \mathbb{P}_z(\tau_R > 1) \max_{w \in B_R} \mathbb{E}_w(|\tilde{X}_{\tau_R}|).$$

Plugging the latter into the former and taking the maximum over $z \in B'_R$, we obtain for all $x \in B_R$,

$$\mathbb{E}_x(|\tilde{X}_{\tau_R}|) \leq \max_{z \in B'_R} \mathbb{E}_z(|\tilde{X}_{\tau_R}| \mid \tau_R = 1) + \max_{z \in B'_R} \mathbb{P}_z(\tau_R > 1) \max_{w \in B_R} \mathbb{E}_w(|\tilde{X}_{\tau_R}|).$$

Finally, taking maximum over $x \in B_R$, we arrive at

$$\mathbb{E}_x(|\tilde{X}_{\tau_R}|) \leq \frac{1}{1 - \max_{z \in B'_R} \mathbb{P}_z(\tau_R > 1)} \max_{z \in B'_R} \mathbb{E}_z(|\tilde{X}_{\tau_R}| \mid \tau_R = 1).$$

The quantities on the right hand side are clearly finite due to exponentially decaying probabilities for the jumps of \tilde{X} and the fact that $z \in B'_R$. Moreover the maximums are taken over a finite set. This together with the fact that $\tilde{a}(x, o) \lesssim (\log |x - o|)^c \lesssim |x| + |o|$ as $x \rightarrow \infty$ completes the proof. \square

Proof of Proposition 3.5 By Lemma 3.6, we have

$$\tilde{G}_R(x, o) - \tilde{G}_R(x', o) = -(\tilde{a}(x, o) - \tilde{a}(x', o)) + \mathbb{E}_x(\tilde{a}(\tilde{X}_{\tau_R}, o)) - \mathbb{E}_{x'}(\tilde{a}(\tilde{X}_{\tau_R}, o)),$$

so it suffices to prove

$$\mathbb{E}_x(\tilde{a}(\tilde{X}_{\tau_R}, o)) - \mathbb{E}_{x'}(\tilde{a}(\tilde{X}_{\tau_R}, o)) \rightarrow 0 \tag{3.14}$$

as $R \rightarrow \infty$. This will follow rather simply from our coupling arguments, where we will choose the parameter r in the stage 3 of the coupling to be $R/(\log R)^2$.

Reasoning as in (3.9), we see that

$$\mathbb{P}_x(\tau_R < T) \lesssim \frac{(\log R)^d}{R} \tag{3.15}$$

for some $d > 0$ as $R \rightarrow \infty$ while x, x' are fixed of the same class. Since we already know from Proposition 3.4 that $\tilde{a}(x, o)$ grows at most like $\log(|x - o|)^c$, (3.15) implies that the difference of expectations in the left hand side of (3.14) is at most $O((\log R)^{c+d}/R)$ and so tends to zero as $R \rightarrow \infty$. \square

We will now consider random walks which are killed on a portion of the boundary of a large box Λ_R but may have different (e.g., reflecting) boundary conditions on other portions of the boundary. We will show that the same result as Proposition 3.5 holds provided that the Dirichlet boundary conditions are, roughly speaking, macroscopic. More precisely, let $\Lambda_R \subset \mathbb{Z}^2$ be such that $B(0, R) \subset \Lambda_R$. Let $\partial\Lambda_R$ denote its (inner) vertex boundary, and let $\partial_D\Lambda_R$ denote a subset of $\partial\Lambda_R$. Suppose that \tilde{X}^{Λ_R} is a (lazy) random walk with transitions given by $\tilde{p}(x, y)$ if $x, y \in \Lambda_R$ and suppose that the walk is absorbed on $\partial_D\Lambda_R$. We suppose that $\partial_D\Lambda_R$ is such that from every vertex in Λ_R , $\partial_D\Lambda_R$ contains a straight line segment of length αR , and at distance at most $\alpha^{-1}R$ from x , where $\alpha > 0$ is a (small) positive constant. Note that these assumptions are satisfied for the domains \mathcal{G}_n we consider in Theorem 1.2 (after gluing together with the reflected copy as in the definition of Γ). Indeed, the (approximate rectangles) \mathcal{G}_n are constructed in such a way that both the odd and even effective bulk random walks are killed on half of the upper side of \mathcal{G}_n .

We do not specify the transition probabilities for \tilde{X}^Λ when it is on $\partial\Lambda_R \setminus \partial_D\Lambda_R$. Let $\tilde{G}^{\Lambda_R}(x, o) = \mathbb{E}_x(\sum_{n=0}^\infty \mathbf{1}_{\{\tilde{X}_n^{\Lambda_R}=o\}})$ denote the corresponding unnormalised Green's function.

Proposition 3.7 *As $R \rightarrow \infty$, for any fixed x, x' of the same class and any fixed o ,*

$$\tilde{G}^{\Lambda_R}(x, o) - \tilde{G}^{\Lambda_R}(x', o) \rightarrow -(\tilde{a}(x, o) - \tilde{a}(x', o)).$$

As a consequence, the same convergence is true also for the nonlazy walk X instead of \tilde{X} .

Proof For this proof we will need the following lemma, which says that from any point there is a good chance to hit the boundary without returning to the point, whence the expected number of visits to that point before hitting the boundary is small.

Lemma 3.8 *There exists a constant such that the following holds for all $k \geq 2$ and vertex o of Γ . Let L be a lattice line at distance k from o and of same class as o . Then*

$$\mathbb{P}_o(T_L < T_o^+) \gtrsim (\log k)^{-1}, \tag{3.16}$$

where T_L is the hitting time of L , T_o^+ is the return time to o .

Proof of Lemma 3.8 We start by noticing that, up to a factor equal to the total conductance at o , the probability on the left-hand side is equal to the effective conductance (or inverse of the effective resistance $\mathcal{R}_{\text{eff}}(o; L)$) between o and L . Since the total conductance at o is bounded away from 0 and ∞ , it suffices to show that

$$\mathcal{R}_{\text{eff}}(o; L) \lesssim \log k.$$

This can either be proved directly or by comparison with the analogous estimate on \mathbb{Z}^2 through Rayleigh’s monotonicity principle (see Chapter II of [31]). A direct proof is to construct a unit flow θ from o to L and estimating its Dirichlet energy $\mathcal{E}(\theta) = \sum_e \theta(e)^2 \text{res}(e)$, where $\text{res}(e)$ denotes the resistance of e . Such a unit flow can be constructed by the method of random paths, as discussed in (2.17) of [31]: we consider a cone of fixed aperture whose apex is at o and intersects L , then choose a line at random in that cone starting at o and whose angle is uniformly selected among the set of possibilities. We get a directed lattice path π from o to L by selecting a lattice path staying as close as possible to this random line (with ties broken in some arbitrary way), staying on the same sublattice as o and L . Note that this path never uses long range edge along the real line, and in fact jumps only by $\pm 2e_i, i = 1, 2$, at any given steps. A unit flow θ from o to L is obtained by setting $\theta(e) = \mathbb{P}(e \in \pi) - \mathbb{P}(-e \in \pi)$ (where $-e$ denotes the reverse of the edge e). Then if e is at distance j from o ,

$$|\theta(e)| \leq \mathbb{P}(e \in \pi) + \mathbb{P}(-e \in \pi) \lesssim \frac{1}{j}$$

since there are $O(j)$ edges at distance j . Hence

$$\mathcal{E}(\theta) \leq \sum_{j=1}^k O(j) \frac{1}{j^2} \lesssim \log k.$$

Since the effective resistance is smaller than the energy of any flow from o to L , we get the desired bound. □

Now let us return to the proof of Proposition 3.7. We apply the full plane coordinatewise mirror coupling of Proposition 3.5 (that is, with the parameter r chosen to be $R/(\log R)^2$), until the time S_R either of the walks leaves the ball $B_R = B_R(0)$. If they have not coupled before S_R , we consider this a failure and will not try to couple them after: we let them evolve independently.

Then note that (we write \tilde{X} for \tilde{X}^{Λ_R} for simplicity)

$$\tilde{G}^{\Lambda_R}(x, o) - \tilde{G}^{\Lambda_R}(x', o) = \mathbb{E}_x(L_{S_R}^{\tilde{X}}(o)) - \mathbb{E}_{x'}(L_{S_R}^{\tilde{X}'}(o)) \tag{3.17}$$

$$+ \mathbb{E}(L_{(S_R, \infty)}^{\tilde{X}}(o) - L_{(S_R, \infty)}^{\tilde{X}'}(o)), \tag{3.18}$$

where $L_{S_R}^{\tilde{X}}(o)$ is as in (3.13). Note that the term in (3.17) converges to $-(\tilde{a}(x, o) - \tilde{a}(x', o))$ by Proposition 3.5. So it suffices to show that the term in (3.18) converges to zero. However, this is an easy consequence of the following facts:

- If the coupling was successful before S_R , then the random variable in the expectation of (3.18) is zero.
- The probability that the coupling has failed (i.e., that the walks did not meet before leaving B_R) is $\lesssim (\log R)^4/R$, by (3.15).
- Conditionally on not having coupled by time S_R , the expected number of visits to y after that time is $\lesssim \log R$ by Lemma 3.8 and by assumption on the Dirichlet

part $\partial_D \Lambda_R$. (In fact, the lemma is stated for hitting an infinite line, but it is easy checked that the argument shows it is a segment of macroscopic size that is being hit with the stated probability).

This completes the proof. \square

We apply this to the bulk effective random walk of Sect. 2.6. This yields the following corollary which also concludes the proof of Theorem 1.2. Recall Definition 3.1 of \tilde{a} , and let $a = \tilde{a}/2$ be its nonlazy version.

Corollary 3.9 *Let \mathcal{G}_n be an increasing sequence of domains such that $\cup_n \mathcal{G}_n = \mathbb{Z}^2 \cap \mathbb{H}$ as in Theorem 1.2. Consider the free boundary dimer model on \mathcal{G}_n with weights as described in Corollary 2.8. Then, the inverse Kasteleyn matrix converges pointwise as $n \rightarrow \infty$ to a matrix indexed by the vertices of $\mathbb{Z}^2 \cap \mathbb{H}$, called the **coupling function**, and given in matrix notation by*

$$C = -AK^*,$$

where

$$A(u, v) = \frac{1}{2D(u, v)} \tilde{a}(u, v) = \frac{1}{D(u, v)} a(u, v)$$

is the normalised potential kernel associated with the effective (odd and even) bulk (nonlazy) random walks.

In particular, μ_n converges weakly as $n \rightarrow \infty$ to a law μ which describes a.s. a random monomer dimer configuration on $\mathbb{Z}^2 \cap \mathbb{H}$.

Proof The first part of the statement follows from the random walk representation of K^{-1} in finite volume from Corollary 2.8, the interpretation of K^* as a difference operator, and the convergence of differences of Green's functions of the bulk effective walk from Proposition 3.7.

The convergence in law is a standard application of Kasteleyn theory. Indeed, this follows from the fact the local statistics of μ_n are described by local functions of the inverse Kasteleyn matrix (which we will for instance recall in Theorem 5.3). It is also clear that μ is supported on monomer-dimer configurations on $\mathbb{Z}^2 \cap \mathbb{H}$. \square

4 Scaling limit of discrete derivative of potential kernel

Let $x, y \in \bar{\mathcal{G}}_\delta := (\delta\mathbb{Z})^2$. The purpose of this section will be to prove a scaling limit for the discrete derivatives of the potential kernel $\tilde{a}(x, y) - \tilde{a}(x', y)$ associated to the lazy (odd) effective random walk. (Contrary to the previous section, the second variable will more typically be called y than o in this section). As mentioned at the beginning of Sect. 3, the same result holds for both the even and odd walk, but for convenience (and also because this is a slightly more complicated case) we write our proofs in the odd case.

Theorem 4.1 *Let $x, y \in \bar{\mathcal{G}}_\delta$ and let $x' = x \pm 2\delta e_i \in \bar{\mathcal{G}}_\delta, i = 1, 2$. Suppose $\Im(x)\Im(y) \geq 0$ and $\min(|\Im(x)|, |\Im(y)|) \geq \rho$ for some fixed arbitrary $\rho > 0$. Then there exists $\varepsilon > 0$ depending only on ρ such that as the mesh size $\delta \rightarrow 0$, uniformly over such points x, y ,*

$$\begin{aligned} &\tilde{a}(x', y) - \tilde{a}(x, y) \\ &= \begin{cases} \frac{2}{\pi} \Re\left(\frac{x' - x}{x - \bar{y}}\right) + o(\delta^{1+\varepsilon}) & \text{if } x, y \text{ are of different class} \\ \frac{4}{\pi} \Re\left(\frac{x' - x}{x - y}\right) - \frac{2}{\pi} \Re\left(\frac{x' - x}{x - \bar{y}}\right) + o(\delta^{1+\varepsilon}) + O\left(\frac{\delta}{|x-y|}\right)^2 & \text{if } x, y \text{ are of the same class} \end{cases} \end{aligned} \tag{4.1}$$

To prove this theorem, we will first show that the potential kernel can be compared to that of a **coloured random walk** on the lattice. The coloured random walk is a lazy simple random walk on the lattice $(2\delta\mathbb{Z})^2$ which carries a black or white colour (in addition to its position). Its position moves like simple random walk on the lattice. It changes colour with some fixed probability $p \in (0, 1)$ each time it touches the real line independently of the rest, and otherwise remains constant. If X is a coloured random walk, we will use $\sigma(X_s)$ to denote the colour of the coloured walk X at time s (and again, this is different from the colour of the vertex X_s): thus, we will write $\sigma(X_s) = \bullet$ if X is black at time s , and $\sigma(X_s) = \circ$ if X is white at time s . Although X_s consists both of a position $x \in (2\delta\mathbb{Z})^2$ and a colour, we will sometimes with an abuse of notation refer to X_s as only a position.

Remark 4.2 We warn the reader that this should *not* be confused with the black/white colouring (which we call class precisely to avoid confusion) of the vertices of our graph $\bar{\mathcal{G}}_\delta$: indeed, the position of the coloured walk is in $(2\delta\mathbb{Z})^2$ and so its ‘‘class’’ in $\bar{\mathcal{G}}_\delta$ remains constant.

Note that x and x' are necessarily of the same class (hence the same colour). However, y may be of a different colour. We will choose p to correspond to the probability that the odd effective walk makes a jump of odd length when it touches the real line: thus,

$$p = \frac{1}{4} \sum_{k \in \mathbb{Z}} q_{0, (2k+1)e_1}^\infty \tag{4.2}$$

where q^∞ is as in (2.9).

We will prove the following two results. Let $y, x \in (2\delta\mathbb{Z})^2$ and choose a colour among $\{\circ, \bullet\}$, say \bullet . Let $\tilde{a}^\bullet(x, y)$ denote the potential kernel of the coloured random walk, constructed as in Definition 2.8 but only counting visits to y with the predetermined colour \bullet : that is,

$$\tilde{a}^\bullet(x, y) = \sum_{s=1}^\infty \mathbb{P}(X_s = y; \sigma(X_s) = \bullet)$$

where X is a coloured walk starting from x with initial colour \bullet . The fact that the series defining \tilde{a}^\bullet converges is an immediate consequence of the arguments in Sect. 3.3, which apply much more directly here.

The first result below shows that the potential kernel of the lazy effective walk and of the coloured walk are quite close to one another, in the sense that the difference in their discrete derivatives are of lower order than δ , our target for Theorem 4.1. In the next statement we write $y \approx x$ to denote that x and y are of different class.

Proposition 4.3 Fix $\rho > 0$. Let $x, y \in \mathcal{G}_\delta$, and let $z = x + \delta \mathbf{1}_{y \approx x}$ (resp $z' = x' + \delta \mathbf{1}_{y \approx x}$), so that z and z' are of the same class as y . Let us write $\nabla_x f(x)$ for $f(x') - f(x)$ (resp. $\nabla_z f(z) = f(z') - f(z)$). Then there exists $\varepsilon > 0$ such that as $\delta \rightarrow 0$,

$$|\nabla_x \tilde{a}(x, \bar{y}) - \nabla_z \tilde{a}^\bullet(z, \bar{y})| \lesssim \delta^{1+\varepsilon},$$

uniformly over x, y , with $\min(\Im(x), \Im(y)) \geq \rho$.

The next proposition says that the potential kernel of the coloured walk is close to $1/2$ that of the regular lazy simple random walk: this is because when the walk touches the real line, it does so many times in a row typically, and so is roughly equally likely to end up with the colour \bullet or \circ . Moreover in the above setting the walk is forced to touch the real line in order to go from x to \bar{y} . Let $\tilde{b}(x, y) = \tilde{b}(x - y)$ denote the potential kernel of lazy simple random walk on $(2\delta\mathbb{Z})^2$.

Proposition 4.4 In the same setting as Proposition 4.3,

$$|\nabla_x \tilde{a}^\bullet(z, \bar{y}) - \frac{1}{2} \nabla_z \tilde{b}(z, \bar{y})| \lesssim \delta^{1+\varepsilon},$$

for some $\varepsilon > 0$.

Proof of Theorem 4.1 given Proposition 4.3 and Proposition 4.4 It is enough to combine Propositions 4.3 and 4.4 as well as known estimates on the two-dimensional simple random walk potential kernel.

Let us give a few details. Suppose we are in the first case where x, y are of different class. This means only walks going through the boundary have the possibility to contribute to the potential kernel. By the reflection symmetry, the walks from x to y going through the boundary have the same weight as the walks from x to \bar{y} . In the full plane for simple random walk, (see e.g. Theorem 4.4.4. in [30]), the potential kernel has the form

$$b(z, 0) = \frac{2}{\pi} \log |z| + C + o(|z|^{-1})$$

for some constant $C > 0$, as $z \rightarrow \infty$. Let us rescale the lattice so that it becomes $\delta\mathbb{Z}^2$, and let us adopt complex notation, so $\log |x| = \Re(\log x)$, and let $h = x' - x = \pm 2\delta e_i$. Then

$$\begin{aligned} b(x, \bar{y}) - b(x', \bar{y}) &= \frac{2}{\pi} \Re(\log(x - \bar{y} + h) - \log(x - \bar{y})) + o(\delta) \\ &= \frac{2}{\pi} \Re\left(\frac{h}{x - \bar{y}}\right) + o(\delta). \end{aligned} \tag{4.3}$$

Now, multiplying by 2 to account for laziness, and by 1/2 to account for the loss at the boundary (the real line) (Proposition 4.4) and we get the first line in (4.1).

To get the second line, we observe that if x and y are of the same class, there are two types of effective random walks to consider: the effective random walks going from x to y in the full plane without touching the boundary (type I), and those which do touch the boundary (type II). The effective random walks of type I can be written as all simple random walks going from x to y in the plane (type III) minus simple random walks going from x to y through the boundary (type IV). By Propositions 4.3 and 4.4, the walks of type IV contribute roughly twice as much as those of type II. So we have to count walks of type III minus those of type II. Those of type III contribute $\frac{4}{\pi} \Re\left(\frac{x'-x}{x-y}\right) + O\left(\frac{\delta}{|x-y|}\right)^2$ to the gradient of the potential kernel (the factor in front is twice that of (4.3) due to laziness, the error term comes from Corollary 4.4.5 in [30]). The contribution of type II on the other hand is exactly counted by the first line of (4.1). This proves Theorem 4.1. \square

Now we derive the version which is useful for later, which includes folding the plane onto itself so that the walk is reflected on the real line, and is not lazy. The corresponding potential kernel satisfies $a = \tilde{a}/2$.

Corollary 4.5 *Let us assume that $x' = x \pm 2\delta e_i \in \delta\mathbb{Z}^2 \cap \mathbb{H}$, $i = 1, 2$. Let $y \in \delta\mathbb{Z}^2 \cap \mathbb{H}$. Then there exists $\varepsilon > 0$ such that as the mesh size $\delta \rightarrow 0$, uniformly over points x, y such that $\min(\Im(x), \Im(y)) \geq \rho > 0$,*

$$a(x', y) - a(x, y) = \begin{cases} \frac{2}{\pi} \Re\left(\frac{x' - x}{x - \bar{y}}\right) + o(\delta^{1+\varepsilon}) & \text{if } x, y \text{ are of different class,} \\ \frac{2}{\pi} \Re\left(\frac{x' - x}{x - y}\right) + o(\delta^{1+\varepsilon}) + O\left(\frac{\delta}{|x-y|}\right)^2 & \text{if } x, y \text{ are of the same class.} \end{cases} \tag{4.4}$$

Proof of Corollary 4.5 given Theorem 4.1 As before the first case (when x, y are of different classes) is easiest to compute. Since the walk is now nonlazy, we need to multiply the values of the potential kernel by 1/2, but also add the walks from x to \bar{y} ; both are counted by the same formula in the first line of (4.1), and so the factor remains $2/\pi$ overall.

In the second case when x, y are of the same class, we note that the number of lazy walks from x to y that don't touch the boundary are, as observed above, given by $\frac{4}{\pi} \Re\left(\frac{x'-x}{x-y}\right)$ (type I). On the other hand, when we do the folding, we must add the walks that touch boundary and go from x to y , to those going from x to \bar{y} . This gives us one extra group of walks of type II and so these cancel. Multiplying by 1/2 to account for non-laziness gives us the second line of (4.4).

Thus it remains to prove the two propositions 4.3 and 4.4. We do so in the following two subsections.

4.1 Proof of Proposition 4.3

We will prove this by coupling. We will need to compare $\nabla_x \tilde{p}_t(x, o)$ and $\nabla_z \tilde{p}_t^\bullet(z, o)$, where $\tilde{p}_t^\bullet(z, o) = \mathbb{P}_z(X_t = o, \sigma(X_t) = \bullet)$ for the coloured walk, where we take

$o = \bar{y}$, and z is a vertex chosen as in Proposition 4.3. We will see that by coupling our effective walks with coloured walks we will gain an order of magnitude compared with (3.2): that is, we will show that

$$|\nabla_x \tilde{p}_t(x, o) - \nabla_z \tilde{p}_t^\bullet(z, o)| \leq t^{-3/2-\varepsilon}; \quad t \leq \delta^{-2-\varepsilon}, \tag{4.5}$$

for some $\varepsilon > 0$. Given (4.5), reasoning as in the proof of Proposition 3.4 (with $R = \delta^{-1}$, and using the improved (4.5) instead of (3.12) in the range up to $t = \delta^{-2-\varepsilon}$), we immediately deduce Proposition 4.3.

We will couple the effective walk X and a coloured walk Z as follows; as in the previous section we work with lazy versions. The coupling will be similar to the one in Sect. 3.3, but it is simpler since we are allowed to choose the starting point of Z . We have chosen z in the statement of Proposition 4.3 so that X and Z start immediately from the same horizontal line. This simplifies the coupling analysis; as in the previous coupling this property will be preserved forever under the coupling, (so essentially only the last stage, stage 4, needs to be described). More precisely, recall that we set $z = x$ if x and \bar{y} are of the same class, and $z = x + \delta$ otherwise. In any case Z will always be of the same class as o . Until hitting the real line, we take X and Z to evolve in parallel, with equal jumps. After hitting the real line, we may arrange the coupling so that they are always on the same horizontal line by always first tossing the Coordinate coin, so that any movement in the vertical coordinate is replicated for both walks no matter what. Beyond the Coordinate and Laziness coins, we will need a third coin which we use to indicate changes in the sublattice (for X) and in colour (for Z). This coin is only used when the walks are on the real line and a horizontal movement is to take place. We call this coin Parity. Unlike the other two coins, Parity comes up heads with the fixed probability $p \in (0, 1)$ from (4.2) which in general is not $1/2$.

It remains to specify what to do if the Coordinate coin indicates a horizontal movement. To describe this, we need to introduce the following stopping times. Let $\sigma_0 = \inf\{t \geq 0 : X_t \in \mathbb{R}\}$ denote the hitting of \mathbb{R} by X (or equivalently by Z), and let $\tau_0 = \inf\{t \geq \sigma_0 : \Im(X_t) \leq \Im(\bar{y})/3\}$ be the hitting time of the line

$$\Delta = \{z \in \mathbb{C} : \Im(z) = \lfloor \Im(\bar{y})/3 \rfloor\}$$

by X (or equivalently Z , since X and Z are always on the same horizontal line). Then define σ_n, τ_n inductively as follows:

$$\sigma_n = \inf\{t \geq \tau_{n-1} : X_t \in \mathbb{R}\}; \quad \tau_n = \inf\{t \geq \sigma_n : X_t \in \Delta\}.$$

Write $X_t = (u_t, v_t)$ and $Z_t = (u'_t, v'_t)$ with $v_t = v'_t$ as explained above.

- If $X_t, Z_t \in \mathbb{R}$. Toss the Parity coin: if it comes heads, let X_t take a jump from its conditional distribution given that it is odd, and let Z_t change colour and make an independent jump. If it is tails, let X_t take a jump from its conditional distribution given that it is even, and let Z_t keep its current colour and make an independent jump.

- Now suppose $X_t, Z_t \notin \mathbb{R}$. If X_t, \bar{y} are of a different class, then let X_t and Z_t evolve in parallel (with equal jumps). This will remain so until hitting again the real line, where there will be a chance to change class again.
- X_t, \bar{y} are of the same class, and thus also of the same class as Z_t . In that case, the evolution depends on whether $t \in [\sigma_n, \tau_n]$ for some $n \geq 0$ or $t \in (\tau_n, \sigma_{n+1})$ for some n : If $t \in [\sigma_n, \tau_n]$ then the walks evolve in parallel. Otherwise, we use Laziness to first ensure that $u_t - u'_t = 0 \pmod{4\delta}$ after a number of steps which has geometric tail. Once that is the case, we let u_t and u'_t evolve in mirror from one another, so $(u_{t+1} - u_t) = -(u'_{t+1} - u'_t)$.

In general the walks get further from each other during a phase of the form $[\sigma_n, \tau_n]$ but get closer together again during the phase $[\tau_n, \sigma_{n+1}]$. Note that a visit to o necessarily occurs during such a phase. In fact we will see that typically the walks agree (if they are on the same sublattice) by the time they reach 2Δ or return to \mathbb{R} . Furthermore, only a small number of phases need to be considered if $t \leq \delta^{-2-\varepsilon}$ (of order at most $\delta^{-\varepsilon}$). Let us say that a non coupled visit to o occurs at time t if $\{X_t = o\} \Delta \{Z_t = o, \sigma(Z_t) = \bullet\}$ occurs (where Δ denotes symmetric difference).

The coordinatewise mirror coupling between X and X' on the one hand, and the above described coupling between X and Z on the other hand, induce a coupling between four processes: X, X' (effective walks starting from x, x') and Z, Z' (coloured walks started from z, z'). Here we take $z' - z = x' - x = \delta$, as in the statement of Proposition 4.3. Note that if X and X' merge before hitting the real line, then so do Z and Z' . The difference between the gradient of the transition probabilities can now be written as an expectation

$$\nabla_x \tilde{p}_t(x, o) - \nabla_z \tilde{p}_t(z, o) = \mathbb{E}(\mathbf{1}_{\{X_t=o\}} - \mathbf{1}_{\{X'_t=o\}} - \mathbf{1}_{\{Z_t=o; \sigma(Z_t)=\bullet\}} + \mathbf{1}_{\{Z'_t=o; \sigma(Z'_t)=\bullet\}}) \tag{4.6}$$

To get a nonzero contribution it is necessary that X did not couple with X' by time $(T_{\mathbb{R}} \wedge t/2)$ or that Z did not couple with Z' by time $t/2$. Both have a probability which is given by $(\log t)^a/t^{1/2}$ by a slight modification of (3.2) (in fact, since the walks start far from the real line, the proof is much simpler than what is given in Sect. 3.3, and follows directly from gambler’s ruin). Furthermore, given this, it is also necessary that a non coupled visit to o occurs at time t by (X, Z) or by (X', Z') .

To estimate the latter conditional probability, we may condition on everything which happened until time $T_{\mathbb{R}} \wedge t/2$, and we will call s the remaining amount of time until time t , i.e., $s = t - (T_{\mathbb{R}} \wedge t/2) \in [t/2, t]$ so $s \asymp t$. Since until that time the walks have not yet touched the real line, the discrepancy between X and Z is therefore equal to the initial discrepancy $z - x \in \{0, \delta e_1\}$.

Lemma 4.6 *Suppose $s \leq \delta^{-2-\varepsilon}$. Let $N_s = \max\{k : \tau_k \leq s\}$. Then there exists some $c_1, c_2 > 0$ such that $\mathbb{P}(N_s \geq c_1 \delta^{-\varepsilon}) \leq \exp(-c_2 \delta^{-\varepsilon})$.*

Proof Each journey between \mathbb{R} and Δ and back may take more than δ^{-2} with fixed positive probability p , independently of one another. Hence the probability in the lemma is bounded by the probability that a Binomial random variable with parameters $c_1 \delta^{-\varepsilon}$ and p , is less than $\delta^{-\varepsilon}$. Choosing c_1 such that $c_1 p > 1$, the result follows from straightforward large deviations of binomial random variables.

We will need to control the discrepancy between X and Z at the beginning of a coupling stage, of the form τ_k (for $0 \leq k \leq \delta^{-\varepsilon}$), assuming that $\sigma(Z_{\tau_k}) = \bullet$ or equivalently that $X_{\tau_k} \sim \bar{y}$. Let us say that this coupling phase **succeeds** if by the time the walks next hit \mathbb{R} or 2Δ , the discrepancy has been reduced to zero.

We note that the discrepancy between X and Z is typically accumulated when the two walks hit the real line; on the other hand they tend to be reduced to zero during a coupling phase, meaning a coupling phase is likely to be successful. However, we will not aim to control the discrepancy if at any point the coupling phase does not succeed.

The key argument will be to say that so long as there has been no unsuccessful coupling phase, the discrepancy at the beginning of any coupling phase is small. To this end, we introduce ρ_n the first time that the real line has been visited more than n times by either (both) walks. We let Δ_n the (horizontal) discrepancy accumulated by the walks at this n th visit: that is,

$$\Delta_n = \langle (X_{\rho_{n+1}} - X_{\rho_n}) - (Z_{\rho_{n+1}} - Z_{\rho_n}); e_1 \rangle$$

Note that by construction of the coupling, Δ_n are i.i.d. and centered random variables with exponential moments (each of them of order the mesh size δ). We then introduce the martingale

$$M_n = \sum_{i=0}^n \Delta_i$$

which counts the accumulated discrepancy at the n th visit to the real line. If $0 \leq u \leq s$ is a time, let us call $n(u)$ the number of visits to \mathbb{R} by time u . At the end of a successful coupling phase σ_k , the discrepancy is reduced to zero, so in fact in the future (until the beginning of the next coupling phase at time τ_k), the discrepancy will be of the form $M_{n(u)} - M_{n(\sigma_k)}$.

Lemma 4.7 *With probability at least $1 - s^{-2\varepsilon}$, we have*

$$\max_{0 \leq k \leq N_s} |X_{\tau_k} - Z_{\tau_k}| \mathbf{1}_{\mathcal{G}_k} \leq \delta s^{1/4+\varepsilon},$$

where \mathcal{G}_k is the good event that there was no unsuccessful coupling by time σ_k .

Proof Fix $0 \leq k \leq N_s$. Let $j = j(k) = \max\{j \leq k : \text{the coupling starting at } \tau_j \text{ was successful}\}$. Suppose that the event \mathcal{G}_k holds otherwise there is nothing to prove. Then as observed above, the discrepancy at time τ_k is given by

$$|X_{\tau_k} - Z_{\tau_k}| = |M_{n(\tau_k)} - M_{n(\tau_j)}| \leq 2 \max_{n \leq n(\tau_k)} |M_n|.$$

By Chebyshev’s inequality and Doob’s maximal inequality,

$$\mathbb{P}\left(\max_{0 \leq k \leq \delta^{-\varepsilon}} |X_{\tau_k} - Z_{\tau_k}| \mathbf{1}_{\mathcal{G}_k} \geq \delta s^{1/4+\varepsilon}\right) \lesssim \frac{1}{\delta^2 s^{1/2+2\varepsilon}} \mathbb{E}\left(\max_{n \leq n(\tau_{N_s})} |M_n|^2\right)$$

$$\lesssim \frac{1}{\delta^{2s^{1/2+2\varepsilon}}} \mathbb{E} \left(M_{n(\tau_{N_s})}^2 \right).$$

Now, $M_n^2 - c\delta^2 n$ is a martingale for some constant $c > 0$ corresponding to the (rescaled) variance of the increments of the martingale M , so (since $n(\tau_{N_s})$ is trivially bounded by s),

$$\mathbb{E} \left(M_{n(\tau_{N_s})}^2 \right) = c\delta^2 \mathbb{E}(n(\tau_{N_s})) = c\delta^2 \mathbb{E}(L_{\mathbb{R}}(s)),$$

where $L_{\mathbb{R}}(s)$ denote the number of visits to \mathbb{R} by both (either) walks by time s . Since the vertical coordinate performs a delayed simple random walk on the integers, this is less than the expected number of visits to 0 by time s of a one-dimensional walk starting from zero, which is at most $\lesssim \sqrt{s}$. Hence

$$\mathbb{P} \left(\max_{0 \leq k \leq \delta^{-\varepsilon}} |X_{\tau_k} - Z_{\tau_k}| \mathbf{1}_{\mathcal{G}_k} \leq \delta s^{1/4+\varepsilon} \right) \lesssim \frac{1}{s^{2\varepsilon}}$$

as desired. □

We now deduce that all coupling phases are successful with high probability.

Lemma 4.8 *We have that for ε small enough (fixed),*

$$\mathbb{P} \left(\bigcup_{k=0}^{N_s} \mathcal{G}_k^c \right) \lesssim s^{-2\varepsilon}.$$

Proof We may work on the event $\mathcal{N} = \{N_s \lesssim \delta^{-\varepsilon}\}$ and the event \mathcal{D} of Lemma 4.7. On $\mathcal{N} \cap \mathcal{D}$ the probability of an unsuccessful coupling starting from time τ_k may be bounded as follows. Supposing that $\sigma(Z_{\tau_k}) = \bullet$ (or equivalently $X_{\tau_k} \sim \bar{y}$), the walks X and Z start a mirror coupling at time τ_k and they are initially spaced by no more than $\delta s^{1/4+\varepsilon}$, if \mathcal{G}_{k-1} holds. By the gambler’s ruin estimate, the probability for X to avoid the reflection line until hitting either \mathbb{R} or 2Δ is then at most $\delta s^{1/4+\varepsilon}$. Hence

$$\mathbb{P}(\mathcal{G}_k^c; \mathcal{G}_{k-1} \cap \mathcal{N} \cap \mathcal{D}) \leq \delta s^{1/4+\varepsilon} \lesssim \delta^{1/2-3\varepsilon}.$$

Summing over $k \leq \delta^{-\varepsilon}$, we get

$$\mathbb{P} \left(\bigcup_{k=0}^{N_s} \mathcal{G}_k^c; \mathcal{N} \cap \mathcal{D} \right) \lesssim \delta^{1/2-4\varepsilon}.$$

We conclude by Lemma 4.7 and Lemma 4.6. □

Proof of Proposition 4.3 We estimate the right hand side of (4.6). For the random variable in the right hand side to be nonzero, it is necessary that:

- X and X' did not couple prior to time $T_{\mathbb{R}} \wedge t/2$;
- one of the \mathcal{G}_k^c occurs for some $k \leq N_s$;
- and still one of the four walks must visit \bar{y} at exactly time t .

The first event has probability bounded by $\lesssim 1/\sqrt{t}$ by straightforward gambler’s ruin. The second has probability at most $1/t^{2\varepsilon}$ by Lemma 4.8 (since $s \asymp t$). To bound the probability of the third event, we observe the following: if $w \in 2\Delta$, the maximum over all times u of the probability to visit \bar{y} at the specific time u is small:

Lemma 4.9 *We have*

$$\sup_{w \in 2\Delta} \sup_{u \geq 0} \tilde{p}_u(w, \bar{y}) \leq \delta^2 (\log 1/\delta)^c,$$

for some $c > 0$.

Proof This follows from the facts (already used before, so we will be brief) that if $u \leq \delta^{-2}/(\log 1/\delta)^c$ then the probability to be at \bar{y} at time u is at most $\exp(-(\log 1/\delta)^2)$ by subdiffusivity, while for $u \geq \delta^{-2}/(\log 1/\delta)^c$ we have a bound of the form $1/u$ thanks to (3.3). \square

All in all, putting these three events together we find

$$|\nabla_x \tilde{p}_t(x, o) - \nabla_z \tilde{p}_t^\bullet(z, o)| \lesssim t^{-1/2} \times t^{-2\varepsilon} \times \delta^2 (\log 1/\delta)^c$$

Summing over $t \in [\delta^{-2}/\log(1/\delta)^c, \delta^{-2-\varepsilon}]$ we see that this is at most $(\log 1/\delta)^c \delta^{1+7\varepsilon/2}$, which is sufficient. \square

4.2 Proof of Proposition 4.4

At this point we may work exclusively with the simple random walk on $(2\mathbb{Z}) \times (2\mathbb{Z})$ or the coloured simple random walk on the same lattice. Let us write $\mathbb{P}_{x \rightarrow \bar{y}; t}$ for the law of a random walk bridge, i.e., the law of a (lazy) simple random walk on $(2\mathbb{Z})^2$ conditioned to go from x to \bar{y} in time t .

Let $\tilde{q}_t(x, y)$ denote the transition probability for (lazy) simple random walk on $(2\mathbb{Z})^2$. Then note that

$$\tilde{p}_t^\bullet(x, y) = \tilde{q}_t(x, y) \mathbb{P}_{x \rightarrow \bar{y}; t}(\sigma(X_t) = \bullet),$$

where $\sigma(X_t)$ is the colour of the process which changes with probability p every time this process touches the real line. Now, let N denote the number of visits to \mathbb{R} and observe that by conditioning on N ,

$$\mathbb{P}_{x \rightarrow \bar{y}; t}(\sigma(X_t) = \bullet | N = n) = \frac{1}{2} \pm \frac{1}{2} \lambda^n$$

where $\lambda = 1 - 2p$ is the eigenvalue of the 2-state Markov chain which switches state with probability p at each step, and the \pm sign depends on the initial colour $\sigma(X_0)$. Therefore,

$$\nabla_x \tilde{p}_t^\bullet(x, \bar{y}) = \frac{1}{2} \nabla_x \tilde{q}_t(x, \bar{y}) \pm \frac{1}{2} \nabla_x \left(\tilde{q}_t(x, \bar{y}) \mathbb{E}_{x \rightarrow \bar{y}; t}(\lambda^N) \right).$$

Since $\sum_{t=0}^{\infty} \frac{1}{2} \nabla_x \tilde{q}_t(x, \bar{y})$ is by definition the potential kernel of the (lazy) simple random walk $\frac{1}{2} \nabla_x \tilde{b}(x, \bar{y})$, to prove Proposition 4.4, as we already observed before, it suffices to show that there exists $\epsilon' > 0$ such that

$$\left| \tilde{q}_t(x, \bar{y}) \mathbb{E}_{x \rightarrow \bar{y}; t}[\lambda^N] - \tilde{q}_t(x', \bar{y}) \mathbb{E}_{x' \rightarrow \bar{y}; t}[\lambda^N] \right| \lesssim \frac{1}{t^{3/2+\epsilon'}}, \tag{4.7}$$

for $t \in [\delta^{-2}/(\log \delta)^2, \delta^{-2-\epsilon}]$. We recall first that if $0 \leq u \leq t$ and $E \in \mathcal{F}_u = \sigma(X_0, \dots, X_u)$, then by the Markov property:

$$\mathbb{P}_{x \rightarrow \bar{y}; t}(E) = \mathbb{E}_x \left(\mathbf{1}_E \frac{\tilde{q}_{t-u}(X_u, \bar{y})}{\tilde{q}_t(x, \bar{y})} \right). \tag{4.8}$$

Let $T_{\mathbb{L}}$ denote the hitting time of the reflection line bisecting x and x' ; and let $T_{\mathbb{R}}$ denote the hitting time of \mathbb{R} . We introduce the following bad events:

- $B_1 = \{T_{\mathbb{R}} > t - s\}$, where $s = \lceil t/(\log t)^2 \rceil \wedge \lceil \delta^{-2}/(\log \delta)^2 \rceil$.
- $B_2 = \{T_{\mathbb{R}} \leq t - s\} \cap \{T_{\mathbb{L}} > T_{\mathbb{R}}\} \cap \{N_{t-s/2} \leq (\log t)^2\}$, where N_u is the number of visits to \mathbb{R} by time u .

We will first show that both events are highly unlikely. In words, B_1 is unlikely because it requires going to \bar{y} in the remaining s units of time starting from above \mathbb{R} , which means \bar{y} is too far away compared to the time remaining. B_2 is unlikely because it requires avoiding the reflection line for a long time (until touching \mathbb{R}) and thereafter making very few visits to \mathbb{R} .

Lemma 4.10 *For $t \in [\delta^{-2}/(\log \delta)^2, \delta^{-2-\epsilon}]$, we have*

$$\mathbb{P}_{x \rightarrow \bar{y}; t}(B_1) \lesssim \exp(-(\log t)^2) \tilde{q}_t(x, \bar{y})^{-1}.$$

Proof Note that by (4.8),

$$\mathbb{P}_{x \rightarrow \bar{y}; t}(B_1) \leq \mathbb{E}_{x \rightarrow \bar{y}; t} \left(\mathbf{1}_{T_{\mathbb{R}} > t-s} \frac{\tilde{q}_s(X_{t-s}, \bar{y})}{\tilde{q}_t(x, \bar{y})} \right).$$

Now, $\tilde{q}_t(x, \bar{y})$ satisfies the Gaussian behaviour $\tilde{q}_t(x, \bar{y}) \asymp (1/t) \exp(-\frac{|x-\bar{y}|^2}{2t})$ in the range $t \geq \delta^{-2}/(\log \delta)^2$ (see Theorem 2.3.11 in [30]). Since $|X_{t-s} - \bar{y}| \gtrsim \delta^{-1}$ when $T_{\mathbb{R}} > t - s$, and since $|x - \bar{y}| \lesssim \delta^{-1}$, we deduce that for some constant $c > 0$,

$$\tilde{q}_s(X_{t-s}, \bar{y}) \leq \exp(-c\delta^{-2}/s) \leq \exp(-c(\log 1/\delta)^2)$$

on the event $T_{\mathbb{R}} > t - s$, where we used that $s \leq \delta^{-2}/(\log \delta)^2$. The desired inequality follows since $t \leq \delta^{-2-\epsilon}$.

Lemma 4.11 *For $t \in [\delta^{-2}/(\log \delta)^2, \delta^{-2-\epsilon}]$, we have*

$$\mathbb{P}_{x \rightarrow \bar{y};t}(B_2) \lesssim \frac{1}{t^{3/2+\varepsilon'}} (\log t)^6 \tilde{q}_t(x, \bar{y})^{-1},$$

where $\varepsilon' = \frac{1-\varepsilon}{2+\varepsilon}$.

Proof Using (4.8),

$$\mathbb{P}_{x \rightarrow \bar{y};t}(B_2) \leq \mathbb{E}_x \left(\mathbf{1}_{\{T_{\mathbb{R}} < t-s, T_{\mathbb{R}} < T_L\}} \mathbf{1}_{\{N_{t-s/2} \leq (\log t)^2\}} \frac{\tilde{q}_{s/2}(X_{t-s/2}, \bar{y})}{\tilde{q}_t(x, \bar{y})} \right).$$

We estimate the off-diagonal heat kernel term $\tilde{q}_{s/2}(X_{t-s/2}, \bar{y})$ by its diagonal behaviour which is at most $\lesssim 1/s = [(\log t)^2/t] \vee [(\log \delta)^2 \delta^2] \lesssim (\log t)^2 t^{-\frac{2}{2+\varepsilon}}$, where in the last bound we used that $t \leq \delta^{-2-\varepsilon}$ and hence $\delta \leq t^{-\frac{1}{2+\varepsilon}}$. Therefore

$$\mathbb{P}_{x \rightarrow \bar{y};t}(B_2) \lesssim t^{-\frac{2}{2+\varepsilon}} \frac{(\log t)^2}{\tilde{q}_t(x, \bar{y})} \mathbb{P}_x(T_{\mathbb{R}} < t-s, T_{\mathbb{R}} < T_L, N_{t-s/2} \leq (\log t)^2). \tag{4.9}$$

We already know by gambler’s ruin estimates that, since x is at distance $O(\delta)$ from L and at distance $\gtrsim 1$ from \mathbb{R} that $\mathbb{P}_x(T_{\mathbb{R}} < T_L) \leq O(\delta)$. Conditioning on everything up to time $T_{\mathbb{R}}$, and applying the strong Markov property at this time,

$$\mathbb{P}_x(T_{\mathbb{R}} < t-s, T_{\mathbb{R}} < T_L, N_{t-s/2} \leq (\log t)^2) \lesssim \delta \sup_{z \in \mathbb{R}} \mathbb{P}_z(N_{s/2} \leq (\log t)^2).$$

Let T_i denote the length of the intervals between successive visits to the real line. Thus T_i are i.i.d. and $\mathbb{P}(T_i \geq r) \asymp 1/\sqrt{r}$ when $r \rightarrow \infty$ by elementary one-dimensional random walk arguments. Fix $z \in \mathbb{R}$. Then by a union bound,

$$\begin{aligned} \mathbb{P}_z(N_{s/2} \leq (\log t)^2) &= \mathbb{P}_z \left(\sum_{i=1}^{(\log t)^2} T_i \geq s/2 \right) \\ &\leq \mathbb{P}(T_i \geq s/(2(\log t)^2) \text{ for some } 1 \leq i \leq (\log t)^2) \\ &\lesssim (\log t)^3/\sqrt{s} = (\log t)^4/\sqrt{t}. \end{aligned}$$

Therefore, plugging this into (4.9), we find

$$\mathbb{P}_{x \rightarrow \bar{y};t}(B_2) \lesssim t^{-1/2-\frac{2}{2+\varepsilon}} \frac{\delta (\log t)^6}{\tilde{q}_t(x, \bar{y})} \leq t^{-1/2-\frac{3}{2+\varepsilon}} \frac{\delta (\log t)^6}{\tilde{q}_t(x, \bar{y})} = t^{-3/2-\varepsilon'} \frac{(\log t)^6}{\tilde{q}_t(x, \bar{y})},$$

where we again used that $\delta \leq t^{-\frac{1}{2+\varepsilon}}$. □

Finally, we turn to the remaining contribution. Together with Lemma 4.10 and Lemma 4.11, this shows that (4.7) holds with any $0 < \varepsilon < 1/2$.

Lemma 4.12

$$|\nabla_x (\tilde{q}_t(x, \bar{y}) \mathbb{E}_{x \rightarrow \bar{y}; t}(\lambda^N; (B_1 \cup B_2)^c))| \leq |\lambda|^{(\log t)^2}. \tag{4.10}$$

where $\lambda = (1 - 2p) < 1$.

Proof On $(B_1 \cup B_2)^c$, we see that $T_{\mathbb{R}} < t - s$, and either $T_L < T_{\mathbb{R}}$ or $N_{t-s/2} \geq (\log t)^2$. In the latter case, $|\lambda^N| = |\lambda|^{N_t} \leq |\lambda|^{(\log t)^2}$, so this event contributes at most the right hand side of (4.10) to the expectation. To conclude, it therefore suffices to show

$$\tilde{q}_t(x, \bar{y}) \mathbb{E}_{x \rightarrow \bar{y}; t}(\lambda^N; T_L < T_{\mathbb{R}}) = \tilde{q}_t(x', \bar{y}) \mathbb{E}_{x' \rightarrow \bar{y}; t}(\lambda^N; T_L < T_{\mathbb{R}}) \tag{4.11}$$

so that the contribution of this event to the left hand side of (4.10) vanishes exactly. To see this, let us rewrite the left hand side of (4.11) as an expectation involving random walk rather than bridge, and observe that when $T_L < T_{\mathbb{R}}$ the walks from x and x' are coupled *before* hitting the real line, so that the overall number of visits to the real line is the same for both walks. Hence

$$\begin{aligned} \tilde{q}_t(x, \bar{y}) \mathbb{E}_{x \rightarrow \bar{y}; t}(\lambda^N; T_L < T_{\mathbb{R}}) &= \mathbb{E}_x(\lambda^N \mathbf{1}_{\{T_L < T_{\mathbb{R}}\}} \mathbf{1}_{\{X_t = \bar{y}\}}) \\ &= \mathbb{E}_{x'}(\lambda^N \mathbf{1}_{\{T_L < T_{\mathbb{R}}\}} \mathbf{1}_{\{X_t = \bar{y}\}}) \\ &= \tilde{q}_t(x', \bar{y}) \mathbb{E}_{x' \rightarrow \bar{y}; t}(\lambda^N; T_L < T_{\mathbb{R}}), \end{aligned}$$

as desired. □

As explained above, this concludes the proof of Proposition 4.4, and thus also of Theorem 4.1.

5 Convergence to the Neumann Gaussian Free field

From now on we work in the upper-half plane \mathbb{H} with the local (infinite volume) limit μ (depending on z) of the free boundary dimer model from Theorem 1.2. We will write μ to denote both the probability and expectation with respect to μ .

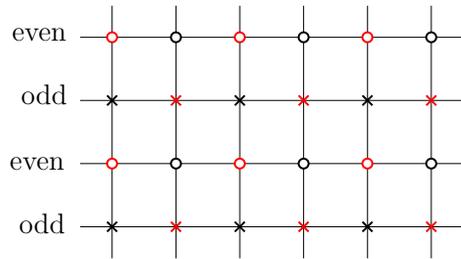
5.1 Infinite volume coupling function and its scaling limit

Let C be the **coupling function**, as defined in Corollary 3.9, i.e., the pointwise limit, as $n \rightarrow \infty$, of the inverse Kasteleyn matrix on \mathcal{G}_n given in matrix notation by

$$C = -AK^*, \tag{5.1}$$

where $A(x, y) = \frac{1}{D(y, y)} a(x, y)$ is the normalised potential kernel of the infinite volume bulk effective (nonlazy) walk. We write A_{even} and A_{odd} for the restriction of A to the even and odd rows respectively. When we unpack (5.1) we find that its meaning is different depending on the respective type of the pair of vertices. We denote the black

Fig. 6 Different types of vertices. The black vertices are drawn in red, and the white vertices are drawn in black (color figure online)



and white vertices in the even and odd rows by the symbols $\circ, \circ, \times, \times$ respectively as illustrated in Fig. 6. In the following, $\frac{\delta}{\delta x_2}$ (resp. $\frac{\delta}{\delta y_2}$) will denote the discrete derivative in the x (resp. y) direction of the second coordinate of the Green’s function. Fix v_1, v_2 two vertices in $\mathbb{H} \cap \mathbb{Z}^2$. Suppose for instance that $v_1 \in \circ$ and $v_2 \in \circ$. Then (5.1) says

$$C(v_1, v_2) = a(v_1, v_2 + 1) - a(v_1, v_2 - 1) \\ := \frac{\delta}{\delta x_2} A_{\text{even}}(v_1, v_2).$$

Here we only see a horizontal derivative since A assigns a nonzero transition weight only between vertices in rows of the same parity, and since v_1 and v_2 belong to rows of the same parity in this example. Note that

$$\frac{\delta}{\delta x_2} \sim 2\delta \frac{\partial}{\partial x_2} \text{ as } \delta \rightarrow 0,$$

where δ is the mesh size. Likewise, if instead we have $v_1 \in \circ$ and $v_2 \in \times$, then

$$C(v_1, v_2) = i \frac{\delta}{\delta y_2} A_{\text{even}}(v_1, v_2).$$

Here the derivative is vertical since v_1 and v_2 belong to rows of different parity. We summarise these computations in a table:

$v_1 \setminus v_2$	$v_2 \in \times$	$v_2 \in \circ$	$v_2 \in \circ$	$v_2 \in \times$
$v_1 \in \times$	$\frac{\delta}{\delta x_2} A_{\text{odd}}$	$i \frac{\delta}{\delta y_2} A_{\text{odd}}$	$i \frac{\delta}{\delta y_2} A_{\text{odd}}$	$\frac{\delta}{\delta x_2} A_{\text{odd}}$
$v_1 \in \circ$	$i \frac{\delta}{\delta y_2} A_{\text{even}}$	$\frac{\delta}{\delta x_2} A_{\text{even}}$	$\frac{\delta}{\delta x_2} A_{\text{even}}$	$i \frac{\delta}{\delta y_2} A_{\text{even}}$

Furthermore, when v_1 is in the black lattice ($v_1 \in \circ$ or $v_1 \in \times$) we obtain the corresponding table simply by translation invariance:

$v_1 \setminus v_2$	$v_2 \in \times$	$v_2 \in \circ$	$v_2 \in \circ$	$v_2 \in \times$
$v_1 \in \times$	$\frac{\delta}{\delta x_2} A_{\text{odd}}$	$i \frac{\delta}{\delta y_2} A_{\text{odd}}$	$i \frac{\delta}{\delta y_2} A_{\text{odd}}$	$\frac{\delta}{\delta x_2} A_{\text{odd}}$
$v_1 \in \circ$	$i \frac{\delta}{\delta y_2} A_{\text{even}}$	$\frac{\delta}{\delta x_2} A_{\text{even}}$	$\frac{\delta}{\delta x_2} A_{\text{even}}$	$i \frac{\delta}{\delta y_2} A_{\text{even}}$

Remark 5.1 It is useful to point out that the terms involving mixed colours and those involving matching colours behave very differently: indeed, if both v_1 and v_2 are of the same colour, then the arguments of the corresponding potential kernel are of different colours. This corresponds to only considering walks that go through the boundary in the definition of the potential kernel (see Corollary 4.5).

There is a convenient algebraic rewriting of these different values. Suppose v_1 and v_2 are two arbitrary vertices (of any colour), and let

$$s(v) = (-1)^{\text{row \# of } v}$$

be the signed parity of the row of v . Then we have

$$C(v_1, v_2) = \frac{1}{4} \left[(1 + s(v_1)s(v_2)) \frac{\delta}{\delta x_2} + (1 - s(v_1)s(v_2)) i \frac{\delta}{\delta y_2} \right] \times [(1 - s(v_1)) A_{\text{odd}}(v_1, v_2) + (1 + s(v_1)) A_{\text{even}}(v_1, v_2)]. \tag{5.2}$$

We will now combine (5.2) with Corollary 4.5 to obtain the scaling limit of the inverse Kasteleyn matrix in the upper half-plane.

Theorem 5.2 *Let z and w be two vertices on $\delta\mathbb{Z}^2 \cap \mathbb{H}$, and fix $\rho > 0$. Then there exists $\varepsilon > 0$ such that uniformly over $z \neq w$ with $\min(\Im(z), \Im(w)) \geq \rho$, as the mesh size $\delta \rightarrow 0$,*

$$C(z, w) = \begin{cases} -\frac{\delta}{2\pi} \left(s(z)s(w) \frac{1}{z-w} + \frac{1}{\bar{z}-\bar{w}} \right) + o(\delta^{1+\varepsilon}) + O\left(\frac{\delta}{|z-w|}\right)^2 & \text{if } z, w \text{ are of different class} \\ \frac{\delta}{2\pi} \left(s(z) \frac{1}{z-w} + s(w) \frac{1}{\bar{z}-w} \right) + o(\delta^{1+\varepsilon}) & \text{if } z, w \text{ are of the same class} \end{cases}$$

Proof We could use the master formula (5.2) but in order to avoid making mistakes it is perhaps easier to consider all the possible cases for the types of vertices z and w using the tables above. We start with the case when z and w are of different colour. We will use the symmetry of the potential kernel $a(w', z) = a(z, w')$ and Corollary 4.5 (applied to the case when the arguments z, w' are of the same colour since by (5.1), multiplying A by K^* changes the colour of the second argument).

- For z, w with $s(z) = s(w) = -1$, we have

$$\begin{aligned} C(z, w) &= \frac{\delta}{\delta x_2} A_{\text{odd}}(z, w) \\ &= \frac{1}{4} (a(w + \delta, z) - a(w - \delta, z)) \\ &= \frac{1}{4} \times \frac{2}{\pi} \Re \left(\frac{2\delta}{w-z} \right) + o(\delta^{1+\varepsilon}) + O \left(\frac{\delta}{|z-w|} \right)^2 \\ &= -\frac{\delta}{\pi} \Re \left(\frac{1}{z-w} \right) + o(\delta^{1+\varepsilon}) + O \left(\frac{\delta}{|z-w|} \right)^2. \end{aligned}$$

The factor $1/4$ comes from the fact that A_{odd} is normalised by the degree of w which is equal to 4 (see (2.29)).

- For z, w with $s(z) = -1$ and $s(w) = 1$, we have

$$\begin{aligned} C(z, w) &= i \frac{\delta}{\delta y_2} A_{\text{odd}}(z, w) \\ &= -i \frac{\delta}{\pi} \Re \left(\frac{i}{z-w} \right) + o(\delta^{1+\varepsilon}) + O \left(\frac{\delta}{|z-w|} \right)^2 \\ &= i \frac{\delta}{\pi} \Im \left(\frac{1}{z-w} \right) + o(\delta^{1+\varepsilon}) + O \left(\frac{\delta}{|z-w|} \right)^2. \end{aligned}$$

- For z, w with $s(z) = s(w) = 1$, since z, w are of different colors, z and $w \pm \delta$ are of the same colour. Note that G_{even} is a signed function such that $G_{\text{even}}(z, w) < 0$ for z, w of different colors, and $G_{\text{even}}(z, w) > 0$ for z, w of the same colour. Therefore we have

$$C(z, w) = \frac{\delta}{\delta x_2} A_{\text{even}}(z, w) = -\frac{\delta}{\pi} \Re \left(\frac{1}{z-w} \right) + o(\delta^{1+\varepsilon}) + O \left(\frac{\delta}{|z-w|} \right)^2.$$

- For z, w with $s(z) = 1$ and $s(w) = -1$, z and $w \pm \delta i$ are again of the same colour. We have

$$\begin{aligned} C(z, w) &= i \frac{\delta}{\delta y_2} A_{\text{even}}(z, w) \\ &= -i \frac{\delta}{\pi} \Re \left(\frac{i}{z-w} \right) + o(\delta^{1+\varepsilon}) + O \left(\frac{\delta}{|z-w|} \right)^2 \\ &= i \frac{\delta}{\pi} \Im \left(\frac{1}{z-w} \right) + o(\delta^{1+\varepsilon}) + O \left(\frac{\delta}{|z-w|} \right)^2. \end{aligned}$$

Let us now consider the case where z and w are of different colors, by applying Corollary 4.5 (when the arguments are of different colors).

- For z, w with $s(z) = s(w) = -1$, we have

$$C(z, w) = \frac{\delta}{\delta x_2} A_{\text{odd}}(z, w) = -\frac{\delta}{\pi} \Re \left(\frac{1}{\bar{z}-w} \right) + o(\delta^{1+\varepsilon}),$$

- For z, w with $s(z) = -1$ and $s(w) = 1$, we have

$$C(z, w) = \frac{\delta}{\delta y_2} A_{\text{odd}}(z, w) = \frac{\delta}{\pi} \Im \left(\frac{1}{\bar{z}-w} \right) + o(\delta^{1+\varepsilon}).$$

- For z, w with $s(z) = s(w) = 1$, since z, w are of the same colour, z and $w \pm \delta$ are of different colors. We have

$$C(z, w) = \frac{\delta}{\delta x_2} A_{\text{even}}(z, w) = \frac{\delta}{\pi} \Re \left(\frac{1}{\bar{z} - w} \right) + o(\delta^{1+\varepsilon}).$$

- For z, w with $s(z) = 1$ and $s(w) = -1$, z and $w \pm \delta i$ are again of different colors. We have

$$C(z, w) = i \frac{\delta}{\delta y_2} A_{\text{odd}}(z, w) = -i \frac{\delta}{\pi} \Im \left(\frac{1}{\bar{z} - w} \right) + o(\delta^{1+\varepsilon}).$$

Combined, we have proved the theorem. □

5.2 Pfaffians and Kasteleyn theory

In this section we recall basics of Kasteleyn theory. In particular we will express local statistics of μ in terms of the coupling function C .

Let A be $2k \times 2k$ skew-symmetric matrix indexed by vertices $w_1, b_1, \dots, w_k, b_k$ of k edges $(w_1, b_1), \dots, (w_k, b_k)$. Then a Pfaffian can be expressed as a sum over matchings of these $2k$ vertices (identified with numbers from 1 to $2k$), in a similar way as the determinant can be expressed as a sum over permutations. Let M be such a perfect matching. We can write it as $(i_1, j_1), \dots, (i_k, j_k)$ where in each pair (i, j) we have $i < j$, and moreover $i_1 < \dots < i_k$. This defines a permutation

$$\pi_M = \begin{bmatrix} 1 & 2 & 3 & 4 & \dots & 2k - 1 & 2k \\ i_1 & j_1 & i_2 & j_2 & \dots & i_k & j_k \end{bmatrix}.$$

Then we have

$$\text{Pf}(A) = \sum_{M \text{ matching}} \text{sgn}(\pi_M) a_{i_1, j_1} \dots a_{i_k, j_k}. \tag{5.3}$$

Based on Kasteleyn’s theorem Kenyon derived the following description of local statistics for the dimer model [21]. Recall that μ denotes the probability measure of the infinite volume free boundary dimer model on $\mathbb{H} \cap \mathbb{Z}^2$.

Theorem 5.3 *Let E be a set of pairwise distinct edges $e_1 = (w_1, b_1), \dots, e_k = (w_k, b_k)$ with the convention that the white vertex comes first. Then*

$$\mu(e_1, \dots, e_k \in \mathcal{M}) = a_E \text{Pf}(C),$$

where \mathcal{M} is the random monomer-dimer configuration under μ , and where $C = C(v_1, v_2)$ is the coupling function restricted to the vertices $v_1, v_2 \in \{w_1, \dots, w_k\} \cup \{b_1, \dots, b_k\}$ (implicit here is the fact that the vertices are ordered from black to white,

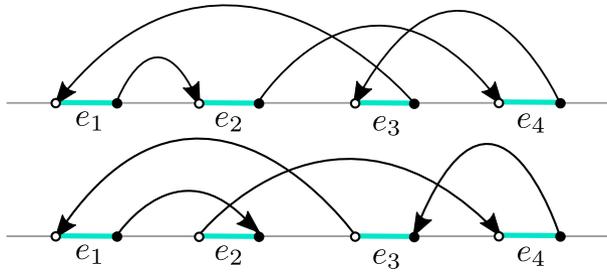


Fig. 7 Graphical representation of the directed matchings $M = \{(b_1, w_2), (b_2, w_4), (b_3, w_1), (b_4, w_3)\}$ (top) and $M' = \{(b_1, b_2), (w_2, w_4), (w_3, w_1), (b_4, b_3)\}$ (bottom) corresponding a single cyclic permutation $\sigma = 1 \rightarrow 2 \rightarrow 4 \rightarrow 3 \rightarrow 1$ with signs $v = (1, 1, 1, 1)$ and $v' = (1, -1, -1, 1)$ respectively. We have $\text{sgn}(\pi_M) = \text{sgn}(\pi_{M'}) = -1$ since the number of arc crossings is odd in both cases

and from 1 to k), and where

$$a_E = \prod_{i=1}^k K(w_i, b_i)$$

is the product of the Kasteleyn weight of each edge, oriented from white to black.

To compute the scaling limit of the height function, we will need to study the centered dimer-dimer correlations. When expanding the Pfaffian into matchings, this leads to a simplification which is the analogue of Lemma 21 in [22]:

Lemma 5.4 *In the setting as above, we have*

$$\begin{aligned} & \mu[\mathbf{1}_{\{e_1 \in \mathcal{M}\}} - \mu(e_1 \in \mathcal{M})] \dots [\mathbf{1}_{\{e_k \in \mathcal{M}\}} - \mu(e_k \in \mathcal{M})] \\ &= a_E \sum_{M \text{ restricted matching}} \text{sgn}(\pi_M) \prod_{\substack{\{u,v\} \in M \\ u < v}} C(u, v) \end{aligned}$$

where a **restricted matching** is a matching M such that w_i cannot be matched to b_i for any $1 \leq i \leq k$, and where $u < v$ means that u comes before v in the fixed order on vertices.

Proof Let M be a matching of the vertices of e_1, \dots, e_k . We will call a pair of matched vertices an M -edge to distinguish it from the edges of the underlying graph. We mark the vertices $w_1, b_1, \dots, w_k, b_k$ (see Fig. 7) in the order from left to right on the real line. For each M -edge, draw an arc (a simple continuous curve) in the upper half-plane connecting the vertices matched by this M -edge (see Fig. 7). Moreover, draw the arcs in such a way that any two arcs cross at most once.

A standard result that can be checked by induction says that

$$\text{sgn}(\pi_M) = (-1)^{\# \text{ arc crossings}}. \tag{5.4}$$

Note that an arc connecting b_i to w_i does not cross any other arc. Using this and Theorem 5.3, we can write

$$\begin{aligned}
 & \mu[(\mathbf{1}_{\{e_1 \in \mathcal{M}\}} - \mu(e_1 \in \mathcal{M})) \dots (\mathbf{1}_{\{e_k \in \mathcal{M}\}} - \mu(e_k \in \mathcal{M}))] \\
 &= \sum_{E' \subseteq E} (-1)^{|E \setminus E'|} \mu(e \in \mathcal{M} \text{ for all } e \in E') \prod_{e \in E \setminus E'} \mu(e \in \mathcal{M}) \\
 &= a_E \sum_{E' \subseteq E} (-1)^{|E \setminus E'|} \sum_{M \in \Pi(E')} \operatorname{sgn}(\pi_M) \prod_{\substack{\{u,v\} \in M \\ u < v}} C(u, v) \prod_{\substack{\{u,v\} \in E \setminus E' \\ u < v}} C(u, v) \\
 &= a_E \sum_{E' \subseteq E} (-1)^{|E \setminus E'|} \sum_{\substack{M \in \Pi(E) \\ E \setminus E' \subset M}} \operatorname{sgn}(\pi_M) \prod_{\substack{\{u,v\} \in M \\ u < v}} C(u, v) \\
 &= a_E \sum_{M \in \Pi(E)} \operatorname{sgn}(\pi_M) \prod_{\substack{\{u,v\} \in M \\ u < v}} C(u, v) \left(\sum_{E' \subseteq E \cap M} (-1)^{|E'|} \right),
 \end{aligned}$$

where $\Pi(E')$ is the set of matchings of the vertices of E' . To finish the proof it is enough to notice that the sum of signs in the last expression is equal to one if $M \cap E = \emptyset$ and it vanishes otherwise. □

5.3 Matchings and permutations

In this section we discuss the combinatorics of matchings and permutations which will be used in the computation of moments of the height function in Sect. 1.2.

Let M be a restricted matching of the vertices of e_1, \dots, e_k (recall that restricted means that the endpoints of an edge cannot be matched with each other). We stress the fact that the objects paired by the matching are the vertices of the edges and not the edges themselves. This will be important in the following combinatorial considerations. We will call a pair of matched vertices an M -edge to distinguish it from the edges of the underlying graph.

We can turn a matching M into a **directed matching** \mathbf{m} by assigning to each M -edge a direction in such a way that each edge e_i has exactly one outgoing and one incoming M -edge. Let \mathfrak{S}_k^* be the set of permutations on k elements with no fixed points. Observe that a directed restricted matching \mathbf{m} defines a permutation $\sigma \in \mathfrak{S}_k^*$ of the k edges: indeed, simply define $\sigma(e_i) = e_j$ where e_j is the edge pointed to by the unique outgoing M -edge emanating from e_i . We will say that the directed matching is **compatible** with the permutation σ . Note that since the matching is restricted, σ does not have fixed points. Let DM_σ be the class of restricted directed matchings compatible with $\sigma \in \mathfrak{S}_k^*$. Note that if σ has $n = n(\sigma)$ cycles and $\mathbf{m} \in \operatorname{DM}_\sigma$, then there are 2^n oriented matchings that correspond to the same unoriented matching as \mathbf{m} (one can choose the orientation of each cycle independently of the choice for other cycles).

Conversely, fix $\sigma \in \mathfrak{S}_k^*$. We now describe how to encode any directed matching \mathbf{m} compatible with σ by a sequence of signs $v \in \{-1, 1\}^k$. The sign v_i denotes the choice of the vertex of e_i from which the outgoing edge of \mathbf{m} will emanate, i.e., if $v_i = +1$

(resp. -1) then the outgoing edge of e_i emanates from the black (resp. white) vertex of e_i . This choice implies that the directed M -edge corresponding to the pair (i, j) such that $\sigma(i) = j$ points to the white (resp. black) vertex of e_j if $v_j = +1$ (resp. $v_j = -1$). The resulting map

$$\{-1, 1\}^k \rightarrow \text{DM}_\sigma \tag{5.5}$$

is clearly a bijection.

We can now rewrite the truncated correlation function from Lemma 5.4 as follows:

$$\begin{aligned} & \mu \left[(\mathbf{1}_{\{e_1 \in \mathcal{M}\}} - \mu(e_1 \in \mathcal{M})) \dots (\mathbf{1}_{\{e_k \in \mathcal{M}\}} - \mu(e_k \in \mathcal{M})) \right] \\ &= a_E \sum_{\sigma \in \mathfrak{S}_k^*} \sum_{M \in \text{DM}_\sigma} \text{sgn}(\pi_M) \frac{1}{2^n} \prod_{(u,v) \in M} C(u, v) (-1)^{\mathbf{1}_{u>v}} \end{aligned} \tag{5.6}$$

where $n = n(\sigma)$ is, as above, the number of cycles of σ . To explain (5.6), we simply recall that each undirected matching that can be oriented as to be compatible with σ corresponds to 2^n directed matchings by choosing the orientation of each cycle arbitrarily. The factor $(-1)^{\mathbf{1}_{u>v}}$ comes from the fact that $C(u, v)$ is antisymmetric and that we always have $i_l < j_l$ in the expansion of the Pfaffian as a sum over matchings (5.3). Here $u > v$ means that u comes later than v in the order defined by $w_1, b_1, \dots, w_k, b_k$.

We will later need the following lemma. What is specifically interesting to us in the expression below is that the right hand side depends very little on the permutation σ , given the signs v .

Lemma 5.5 *Let \mathbf{m} be the restricted directed matching compatible with $\sigma \in \mathfrak{S}_k^*$ and encoded by $v \in \{-1, 1\}^k$ by the map (5.5). We have*

$$\prod_{(u,v) \in \mathbf{m}} (-1)^{\mathbf{1}_{u>v}} \text{sgn}(\pi_M) = (-1)^n \prod_{i=1}^k v_i, \tag{5.7}$$

where $n = n(\sigma)$ is the (total) number of cycles in σ .

Proof Mark the vertices $w_1, b_1, \dots, w_k, b_k$ in the order from left to right on the real line as in Lemma 5.4 and recall formula (5.4). Note that if we flip exactly one sign v_i , then both sides of (5.7) change sign since the parity of the number of crossings between arcs changes (we either cross or uncross the arcs ending at e_i and we do not change the number of crossings for other pairs of arcs), and since the number of decreasing edges (u, v) of \mathbf{m} , i.e., satisfying $u > v$, does not change. We can hence assume that $v_i = +1$ for all i .

Equipped with the graphical representation as in Lemma 5.4 we proceed by induction on k . One can check that the statement is true for $k = 2$. We therefore assume that $k > 2$. Let \mathbf{m} be a directed restricted matching on e_1, \dots, e_k , and let $\sigma \in \mathfrak{S}_k^*$ be the permutation associated with M . Let i, j be such that $\sigma(i) = k$ and $\sigma(k) = j$. Consider a graphical representation of \mathbf{m} . Imagine infinitesimally deforming the path composed of the arcs connecting e_i to e_k and e_k to e_j together with the line segment representing

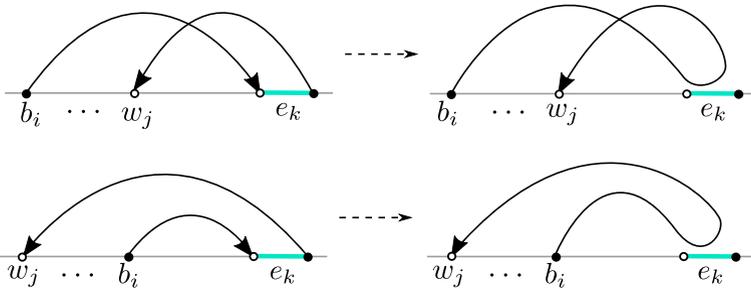


Fig. 8 The induction step from the proof of Lemma 5.5

e_k in such a way that the path is fully contained in \mathbb{H} . This path hence becomes an arc (modulo a possible self-crossing) representing an \mathbf{m}' -edge (b_i, w_j) , where \mathbf{m}' is a directed matching on e_1, \dots, e_{k-1} . Note that \mathbf{m}' is not necessarily restricted (this happens when $i = j$). Let $\sigma' \in \mathfrak{S}_{k-1}$ be the permutation associated to \mathbf{m}' . Note that σ' has the same number of cycles as σ .

In this transformation we replaced an increasing edge (b_i, w_k) and a decreasing edge (b_k, w_j) by the edge (b_i, w_j) . For topological reasons, the deformed path representing the \mathbf{m}' -edge (b_i, w_j) has a self-crossing if and only if (b_i, w_j) is an increasing edge (see Fig. 8). To finish the proof we use (5.4) to evaluate and compare (5.7) for \mathbf{m} and \mathbf{m}' (or \mathbf{m}' with the decreasing edge (b_i, w_i) (and hence one cycle) removed in the case that \mathbf{m}' is not restricted), and we use the induction assumption. \square

5.4 Moments of the height function

In this section we compute the scaling limit of the pointwise moments of the height function on $\delta\mathbb{Z}^2 \cap \mathbb{H}$, which is the penultimate step in establishing its convergence as a random distribution.

We fix $k \geq 1$, and $2k$ faces $a_1, b_1, \dots, a_k, b_k$ of $\delta\mathbb{Z}^2 \cap \mathbb{H}$. We consider disjoint paths γ_i in the dual lattice $(\delta\mathbb{Z}^2 \cap \mathbb{H})^*$ connecting a_i to b_i for $1 \leq i \leq k$. The following is the analogue of Proposition 20 in [22]. Let D denote the minimal distance in the complex plane between any pair of points within $\{a_i, b_i\}_{1 \leq i \leq k}$.

Proposition 5.6 *Let $k \geq 1$. Let $\rho > 0$ be fixed and let $\beta > 0$ be sufficiently small (possibly depending on k). As $\delta \rightarrow 0$,*

$$\left| \mu \left[(h^\delta(a_1) - h^\delta(b_1)) \cdots (h^\delta(a_k) - h^\delta(b_k)) \right] - \right. \tag{5.8}$$

$$\left. \sum_{\mathbf{m} \in \mathcal{M}(1, \dots, k)} \prod_{(i, j) \in \mathbf{m}} -\frac{1}{2\pi^2} \Re \log \frac{(a_i - a_j)(b_i - b_j)(\bar{a}_i - a_j)(\bar{b}_i - b_j)}{(a_i - b_j)(b_i - a_j)(\bar{a}_i - b_j)(\bar{b}_i - a_j)} \right| \rightarrow 0, \tag{5.9}$$

uniformly over the choice of $a_1, b_1, \dots, a_k, b_k$ such that $D \geq \delta^\beta$ and $\min_{1 \leq i \leq k} (\Im(a_i), \Im(b_i)) \geq \rho$.

Proof As in Kenyon [22] we can assume without loss of generality that the paths γ_i are piecewise parallel to the axes and that each straight portion is of even length. In this way, we can pair the edges of a straight portion of the path in groups of two consecutive edges. In order to distinguish between the two edges in a given pair it will be useful to have a notation which emphasises this difference, and following the notations of Kenyon we will call a generic pair of edges α and β respectively; an α -edge will have a black vertex on the right while a β -edge will have a black vertex on its left. The point is that considering their contributions together will lead to cancellations that are crucial in the computation. Also, in this way the contribution from a pair of edges does not depend anymore on the microscopic types of its vertices and has a scaling limit which depends only the macroscopic position.

Let α_t^i (resp. β_t^i) be the indicator that the t -th α -edge (resp. β -edge) in the path γ_i is present in the dimer cover, minus its expectation. In this way due to the definition of the height function and the choice of reference flow,

$$h(a_i) - h(b_i) = \sum_t \alpha_t^i - \beta_t^i.$$

(Note we do not have a factor 4 as in Kenyon because our choice of reference is slightly different in order to deal directly with a centered height function: more precisely, the total flow out of a vertex is one instead of four in Kenyon’s work [22]). We are ignoring here possibly one term on the boundary if the faces a_i and b_i do not have the correct parity; but in any case it is clear that the contribution of a single term in such a sum is of order $O(\delta)$ and so can be ignored in what follows.

We therefore have

$$\mu[(h^\delta(a_1) - h^\delta(b_1)) \cdots (h^\delta(a_k) - h^\delta(b_k))] = \sum_{t_1, \dots, t_k} \mu[(\alpha_{t_1}^1 - \beta_{t_1}^1) \cdots (\alpha_{t_k}^k - \beta_{t_k}^k)]. \tag{5.10}$$

We fix a choice of t_i s and analyse this product. We first expand this product into a sum of 2^k terms containing for each i a term which is either $\alpha_{t_i}^i$ or $-\beta_{t_i}^i$. Consider for simplicity the term containing all of the $\alpha_{t_i}^i$. Write w_i, b_i for the white and black vertices of the edge corresponding to $\alpha_{t_i}^i$. Let E be the set of edges $(w_1, b_2), \dots, (w_k, b_k)$ and let $a_E = \prod_{e \in E} K(e)$. Then by (5.6) we have

$$\mu(\alpha_{t_1}^1 \cdots \alpha_{t_k}^k) = a_E \sum_{\sigma \in \mathfrak{S}_k^*} \sum_{\mathbf{m} \in \text{DM}_\sigma} \text{sgn}(\mathbf{m}) \frac{1}{2^n} \prod_{(u,v) \in \mathbf{m}} C(u, v) (-1)^{1_{u>v}}.$$

We rewrite the sum over directed matchings $\mathbf{m} \in \text{DM}_\sigma$ as a sum over $(v_i)_{1 \leq i \leq k}$ using (5.5), and get (writing \mathbf{m} for the unique directed matching determined by σ and $v = (v_i)_{1 \leq i \leq k} \in \{-1, 1\}^k$),

$$a_E \sum_v \sum_{\sigma \in \mathfrak{S}_k^*} \text{sgn}(\mathbf{m}) \frac{1}{2^n} \prod_{(u,v) \in \mathbf{m}} C(u, v) (-1)^{1_{u>v}}$$

$$= a_E \sum_v \left(\prod_{i=1}^k v_i \right) \sum_{\sigma \in \mathfrak{S}_k^*} (-1)^n \frac{1}{2^n} \prod_{(u,v) \in \mathbf{m}} C(u, v)$$

using Lemma 5.5.

Fix ν and σ (i.e., we fix a directed matching \mathbf{m}) and use Theorem 5.2 to approximate $C(u, v)$. Let $(u, v) \in \mathbf{m}$ and let ν and ν' be the respective values of the variables ν_i associated with the two edges containing the vertices u and v . Note that if $\nu\nu' = 1$ then u and v must be of different colours and so we fall in case 2 of the approximation given by Theorem 5.2, while if $\nu\nu' = -1$ then u and v are of the same colour and so we fall in the first case of this approximation. Hence we get

$$C(u, v) = \begin{cases} -\frac{\delta}{2\pi} \left[s(u)s(v) \frac{1}{u-v} + \frac{1}{\bar{u}-\bar{v}} \right] + o(\delta^{1+\varepsilon}) + O\left(\frac{\delta}{|u-v|}\right)^2 & \text{if } \nu\nu' = 1 \\ \frac{\delta}{2\pi} \left[s(u) \frac{1}{u-\bar{v}} + s(v) \frac{1}{\bar{u}-v} \right] + o(\delta^{1+\varepsilon}) & \text{if } \nu\nu' = -1. \end{cases}$$

The terms $o(\delta^{1+\varepsilon})$ here are uniform on u, v (subject only to the imaginary parts being $\geq \rho$). Since $|u - v| \geq D \geq \delta^\beta$, we see that $O\left(\frac{\delta}{|u-v|}\right)^2 \leq O(\delta^{2-2\beta}) = o(\delta^{1+\varepsilon})$ if β is sufficiently small. We can thus absorb the term $O\left(\frac{\delta}{|u-v|}\right)^2$ into the term $o(\delta^{1+\varepsilon})$ under our assumptions on D .

We expand $\prod_{(u,v) \in \mathbf{m}} C(u, v)$ using the above formula. This gives us another sum of 2^k terms, which we view as a polynomial in the variables $s(w_i), s(b_i)$. We group the terms by their monomials; and since $s(z)^2 = 1$ for any z , these monomials can only be of degree at most one in each variable. We now claim that any monomial such that $s(b_i)$ appears but not $s(w_i)$ for some $1 \leq i \leq k$, or vice-versa, will contribute $o(\delta^{1+\varepsilon})$ when we take into account the equivalent term coming from the same expansion where $\alpha_{t_i}^i$ has been replaced by $-\beta_{t_i}^i$. Indeed, since σ and ν have been fixed, consider what happens when $\alpha_{t_i}^i$ is replaced by $-\beta_{t_i}^i$:

- There is a $-$ sign coming from the change $\alpha_{t_i}^i \rightarrow -\beta_{t_i}^i$.
- The sign of a_E changes by -1 always (consider separately the cases of a horizontal or vertical edge to see this).
- Crucially both $s(b_i)$ **and** $s(w_i)$ change.
- Yet the coefficients accompanying $s(b_i)$ and $s(w_i)$ (both of which are terms of the form $\frac{1}{z-w} + o(\delta^{1+\varepsilon}), \dots$ or $\frac{1}{\bar{z}-\bar{w}} + o(\delta^{1+\varepsilon})$) do not change in the scaling limit, since this term is determined only by the choice of ν , which is fixed.

As a consequence, as we sum over all choices of α and β in the 2^k terms of (5.10), and we expand in terms of monomials as described above, we only keep terms that contain for each $1 \leq i \leq k$ either, simultaneously $s(b_i)$ and $s(w_i)$, or neither of them.

As it turns out, given σ and ν , only very few terms do not cancel out. In fact, for each cycle of σ there will be only two terms. For example, consider the case $k = 4$, $\sigma = (1234)$ a four-cycle, and $\nu = + - + +$. This means we are expanding

$$C(b_1, b_2)C(w_2, w_3)C(b_3, w_4)C(b_4, w_1).$$

Letting z_i be the point in the middle of the edge (b_i, w_i) , the expansion looks like

$$\begin{aligned} & \frac{\delta}{2\pi} \left[s(b_1) \frac{1}{z_1 - \bar{z}_2} + s(b_2) \frac{1}{\bar{z}_1 - z_2} \right] \times \frac{\delta}{2\pi} \left[s(w_2) \frac{1}{z_2 - \bar{z}_3} + s(w_3) \frac{1}{\bar{z}_2 - z_3} \right] \\ & \times \left(-\frac{\delta}{2\pi} \right) \left[s(b_3)s(w_4) \frac{1}{z_3 - z_4} + \frac{1}{\bar{z}_3 - \bar{z}_4} \right] \\ & \times \left(-\frac{\delta}{2\pi} \right) \left[s(b_4)s(w_1) \frac{1}{z_4 - z_1} + \frac{1}{\bar{z}_4 - \bar{z}_1} \right] + o(\delta^{4+\varepsilon}/D^4). \end{aligned}$$

The only terms that survive this expansion with the above requirements are the monomials corresponding to $s(b_1)s(w_1)s(b_3)s(w_3)s(b_4)s(w_4)$ and $s(b_2)s(w_2)$: indeed, choosing or not the term containing $s(b_1)$ in the first line imposes a choice on every other line, which is why just two terms survive this expansion.

Furthermore, crucially, in the corresponding coefficients of the surviving monomials, the variables z_i or \bar{z}_i occurs exactly twice, either twice in the type z_i or twice in the type \bar{z}_i (but never in a mixed fashion). For instance, in the above example, the coefficient will involve either z_1, \bar{z}_2, z_3, z_4 or the other way around: $\bar{z}_1, z_2, \bar{z}_3, \bar{z}_4$. Note that the dependence on z_i or \bar{z}_i is consistent with the choice of signs coming from ν : more precisely, for $z \in \mathbb{C}$ and $\varepsilon = \pm 1$, define z^ε to be z if $\varepsilon = +1$ and \bar{z} if $\varepsilon = -1$. Then for a cyclic permutation σ , the two monomials which survive the expansion have a coefficient proportional to

$$\prod_{i=1}^k \frac{1}{z_i^{\nu_i} - z_{\sigma(i)}^{\nu_{\sigma(i)}}} \quad \text{and} \quad \prod_{i=1}^k \frac{1}{\bar{z}_i^{\nu_i} - \bar{z}_{\sigma(i)}^{\nu_{\sigma(i)}}}$$

and a similar property holds for a general permutation σ by considering each of its cycles separately.

Note furthermore that each such coefficient comes with a factor $\pm(\delta/2\pi)^k \times 2^k$: indeed, when a monomial survives it arises exactly once in each of the 2^k terms from the $\alpha - \beta$ expansion (5.10). The sign itself is determined purely by the parity of the cycle of the permutation: indeed, for an even length cycle the number of times the colour changes as we follow the directed matching must be even; while it must be odd for an odd length cycle.

Suppose $C = \{c_1, \dots, c_n\}$ is the cycle structure of σ . We will use variables $(\varepsilon_c)_{c \in C} \in \{-1, 1\}^n$ to denote which type of monomials we consider. Thus the right hand side of (5.10) (still for a fixed choice of t_i 's) becomes

$$\begin{aligned} & = a\delta^k \sum_{\nu} \left(\prod_{i=1}^k \nu_i \right) \\ & \sum_{\sigma \in \mathfrak{S}_k^*} (-1)^n \frac{1}{2^n} \sum_{\varepsilon} \prod_{i=1}^k \frac{1}{\pi} \frac{1}{z_i^{\nu_i \varepsilon_{c(i)}} - z_{\sigma(i)}^{\nu_{\sigma(i)} \varepsilon_{c(\sigma(i))}}} [s(b_i)s(w_i)]^{(1+\varepsilon_{c(i)}\nu_i)/2} + o(\delta^{k+\varepsilon}/D^k) \end{aligned} \tag{5.11}$$

where $c(i)$ is the cycle containing i .

We now claim that if $\sigma \in \mathfrak{S}_k^*$ has any cycle c of length $|c| > 2$ then it contributes zero to the sum. We start by considering odd cycles. Indeed consider the case where k is odd and σ is a cyclic permutation of length k . Then apply the bijection $v \rightarrow -v$ and $\varepsilon \rightarrow -\varepsilon$ to find that all the terms are unchanged except for a negative sign coming from $\prod_{i=1}^k v_i$. Hence this contribution must be equal to zero, and a similar argument can easily be made when σ contains a cycle of odd length.

In particular, k itself must be even for the contribution to be nonzero. To get rid of permutations containing cycles of even length > 2 , we will rely on the following lemma.

Lemma 5.7 *Let $k > 2$ be even and let $(x_i)_{1 \leq i \leq k}$ be pairwise distinct complex numbers. Let \mathfrak{C}_k be the set of cyclic permutation of length k . Then*

$$\sum_{\sigma \in \mathfrak{C}_k} \prod_{i=1}^k \frac{1}{x_i - x_{\sigma(i)}} = 0.$$

Note in particular that it follows from Lemma 5.7 that if A is the matrix $A_{ij} = 1_{i \neq j} / (x_i - x_j)$ then $\det(A)$ can be written as a sum over matchings (which can be thought of as permutations with no fixed points and where each cycle has length 2):

$$\det(A) = \sum_{\mathbf{m}} \prod_{(u,v) \in \mathbf{m}} \frac{1}{(x_u - x_v)^2} \tag{5.12}$$

which is Lemma 3.1 of Kenyon [23].

Proof of Lemma 5.7. First of all, the case $k = 4$ must be true because of (5.12) (note that the odd cycles clearly give a zero contribution to the determinant by an argument similar to the above).

Using again (5.12) but for $k = 6$ gives the desired identity for $k = 6$, since the terms corresponding to \mathfrak{C}_4 in the expansion of the determinant into permutations contribute zero by the case $k = 4$. Proceeding by induction, we deduce the result for every even $k \geq 4$. □

By Lemma 5.7, the number of cycles n is necessarily $k/2$. Note also that in a two-cycle we get a term of the form $C(z, w)$ and another one of the form $C(w, z) = -C(z, w)$, which results in a term of the form $-C(z, w)^2$. Hence the moment (5.11) becomes

$$\begin{aligned} & a\delta^k \sum_v \left(\prod v_i \right) \sum_{\mathbf{m} \in \mathcal{M}(1, \dots, k)} \left(-\frac{1}{2}\right)^{k/2} \times \\ & \times \prod_{(i,j) \in \mathbf{m}} \left[\frac{(s(b_i)s(w_i))^{(1+v_i)/2} (s(b_j)s(w_j))^{(1+v_j)/2}}{\pi^2 (z_i^{v_i} - z_j^{v_j})^2} \right. \\ & \left. + \frac{(s(b_i)s(w_i))^{(1-v_i)/2} (s(b_j)s(w_j))^{(1-v_j)/2}}{\pi^2 (\bar{z}_i^{v_i} - \bar{z}_j^{v_j})^2} \right] + o(\delta^{k+\varepsilon}/D^k) \end{aligned} \tag{5.13}$$

where $\mathcal{M}(1, \dots, k)$ are the matchings of $1, \dots, k$.

Recall that in the above expression $s(b_i)$ and $s(b_j)$ refer to the sign (parity) of the white and black vertex respectively of the α -edge in position t_i of the path γ_i (we have already accounted for the corresponding β -edge). We will now sum over i and interpret the corresponding sums as discrete Riemann sums converging to integrals. For a horizontal edge (w_i, b_i) , we have $s(b_i)s(w_i) = 1$ whereas it is -1 for a vertical edge: this is simply because s measures the parity of the row. We claim that (as in Kenyon’s proof of Proposition 20 in [22], see the equation between (20) and (21)),

$$2\delta(s(b_i)s(w_i))^{(1+v_i)/2}v_iK(w_i, b_i) = -i\delta z_i^{v_i}. \tag{5.14}$$

Indeed, suppose for instance that γ_i moves horizontally from left to right in step t_i . Then the corresponding α -edge is vertical, and has a black vertex at the bottom so $K(w_i, b_i) = +i$. Furthermore, $\delta z_i = \delta \bar{z}_i = 2\delta$ (since one step of the path corresponds to two faces of length δ each). The vertical cases can be checked similarly (keeping in mind the corresponding values of δz_i and $\delta \bar{z}_i$).

From (5.14), we can multiply by v_i both sides of the equation and take the product over i . Then, recalling that $a = \prod_i K(w_i, b_i)$, and observing also that the second term in each bracket of the right hand side of (5.13) is the same as the first term but with v_i replaced by $-v_i$ and v_j replaced by $-v_j$, (5.13) becomes

$$\sum_v \sum_{\mathbf{m} \in \mathcal{M}(1, \dots, k)} (-1)^k \frac{1}{2^{3k/2}} \times \prod_{(i,j) \in \mathbf{m}} - \left[\frac{\delta z_i^{v_i} \delta z_j^{v_j}}{\pi^2(z_i^{v_i} - z_j^{v_j})^2} + \frac{\delta z_i^{-v_i} \delta z_j^{-v_j}}{\pi^2(z_i^{-v_i} - z_j^{-v_j})^2} \right] + o(\delta^{k+\varepsilon}/D^k) \tag{5.15}$$

(We have kept a term $(-1)^k$ even though k is even to indicate that this comes from $(-1)^{k/2}$ at the top of (5.13) and a factor -1 in each of the $k/2$ terms of the product in the bottom of the same equation. On the other hand, the coefficient -1 in each of the $k/2$ of the product in (5.15) above comes from the coefficient $-i$ (squared) in the right hand side of (5.14).) Fixing the matching and summing over v (so exchanging order of summation) we get

$$(-1)^{3k/2} \frac{1}{2^{3k/2}} \sum_{\mathbf{m} \in \mathcal{M}(1, \dots, k)} \prod_{(i,j) \in \mathbf{m}} \times 2 \left[\frac{\delta z_i \delta z_j}{\pi^2(z_i - z_j)^2} + \frac{\delta \bar{z}_i \delta \bar{z}_j}{\pi^2(\bar{z}_i - \bar{z}_j)^2} + \frac{\delta \bar{z}_i \delta z_j}{\pi^2(\bar{z}_i - z_j)^2} + \frac{\delta z_i \delta \bar{z}_j}{\pi^2(z_i - \bar{z}_j)^2} \right] + o(\delta^{k+\varepsilon}/D^k) \tag{5.16}$$

(The term $(-1)^{3k/2}$ in front comes from the previous $(-1)^k$ in (5.15) and another factor (-1) in each of the $k/2$ terms of the product of the same equation.) Summing over the choice of t_i in (5.10), and since k is even (so $(-1)^{3k/2} = (-1)^{k/2}$), we obtain

$$\begin{aligned} \mu[(h^\delta(a_1) - h^\delta(b_1)) \cdots (h^\delta(a_k) - h^\delta(b_k))] = \\ \frac{(-1)^{k/2}}{2^k} \sum_{\mathbf{m} \in \mathcal{M}(1, \dots, k)} \prod_{(i,j) \in \mathbf{m}} \int_{\gamma_i} \int_{\gamma_j} \left[\frac{dz_i dz_j}{\pi^2(z_i - z_j)^2} + \frac{d\bar{z}_i d\bar{z}_j}{\pi^2(\bar{z}_i - \bar{z}_j)^2} \right. \\ \left. + \frac{d\bar{z}_i dz_j}{\pi^2(\bar{z}_i - z_j)^2} + \frac{dz_i d\bar{z}_j}{\pi^2(z_i - \bar{z}_j)^2} + O\left(\frac{\delta}{D^6}\right) \right] + o\left(\frac{\delta^\epsilon}{D^k}\right). \end{aligned} \tag{5.17}$$

To understand the bound on the error above, the term outside of the brackets corresponds to summing the error in (5.15) over k paths (each of length at most $O(\delta^{-1})$); the term inside corresponds to approximating a Riemann sum by an integral. When we do so, for each sum/integral, we make an error of size at most $O(\delta/D)|\sup f'|$, since each path is at least of length D/δ , and f is the function being integrated, so that here $\sup |f'| = O(D^{-3})$. Furthermore, as these are double integrals, we need to multiply this error for a single integral by the overall value the other integral which we bound crudely by $O(1/D^2)$.

Now observe that

$$\int_{\gamma_i} \int_{\gamma_j} \frac{dz_i dz_j}{(z_i - z_j)^2} = \log \frac{(a_i - a_j)(b_i - b_j)}{(a_i - b_j)(b_i - a_j)}$$

Noting that the four integrals give two pairs of conjugate complex numbers, and recalling that $x + \bar{x} = 2\Re(x)$, we obtain

$$\begin{aligned} \mu[(h^\delta(a_1) - h^\delta(b_1)) \cdots (h^\delta(a_k) - h^\delta(b_k))] \\ = \sum_{\mathbf{m} \in \mathcal{M}(1, \dots, k)} \prod_{(i,j) \in \mathbf{m}} -\frac{1}{2\pi^2} \Re \log \frac{(a_i - a_j)(b_i - b_j)(\bar{a}_i - a_j)(\bar{b}_i - b_j)}{(a_i - b_j)(b_i - a_j)(\bar{a}_i - b_j)(\bar{b}_i - a_j)} + \text{err.} \end{aligned} \tag{5.18}$$

where

$$\text{err.} = o\left(\frac{\delta^\epsilon}{D^k}\right) + O\left(\frac{\delta}{D^6}\right) O(\log D)^{k/2-1} = o(\delta^{\epsilon-k\beta})$$

for β sufficiently small, since $D \geq \delta^\beta$. In particular, if β is sufficiently small (depending on k but not on anything else) then this error goes to zero as $\delta \rightarrow 0$. This concludes the proof of Proposition 5.6. □

5.5 Convergence of the height function

In this section we finish the proof of the scaling limit result from Theorem 1.3.

One can think of the height function on $\delta\mathbb{Z}^2 \cap \mathbb{H}$ (which is defined up to a constant) as a random distribution (generalised function) acting on bounded test functions f

with compact support and mean zero. We follow [6] and write the action as

$$(h^\delta, f) = \int_{\mathbb{H}} \int_{\mathbb{H}} (h^\delta(a) - h^\delta(b)) \frac{f^+(a)f^-(b)}{Z_f} da db, \tag{5.19}$$

where $f^\pm = \max\{\pm f, 0\}$ and $Z_f = \int_{\mathbb{H}} f^+(a) da = \int_{\mathbb{H}} f^-(a) da$. Here we think of h^δ as extended in a piecewise constant way to the interiors of faces. Note that this is well defined as the additive indeterminate constant in h^δ cancels out in this expression. One can also check that this gives the same result as just integrating the height function against f . Note that by Fubini's theorem (h^δ, f) is centered as $\mu(h^\delta(a) - h^\delta(b)) = 0$ for all $a, b \in \mathbb{H}$ by our choice of the reference flow from Sect. 1.1.

The result in Proposition 5.6 is the key step to prove the main result of the paper, which we rephrase below for convenience. In fact, in [22], no further justification beyond the analogue of Proposition 5.6 is provided (this is also the case in [36]). The fact that an argument is missing was already pointed out by de Tilière in [10] (see Lemma 20 in that paper). Here, we follow an approach similar to the one used in [29] and in Toninelli's lecture notes [40] (see in particular Theorem 5.4 and the following discussion), but tailored to our setup since our a priori error estimates are somewhat different.

Theorem 5.8 *Let $\Phi_{\mathbb{H}}^{\text{Neu}}$ be the Neumann Gaussian free field in \mathbb{H} , and let $f_1, \dots, f_k \in \mathcal{D}_0(\mathbb{H})$ (smooth test functions of compact support and mean zero). Then for $l_1, \dots, l_k \in \mathbb{N}$,*

$$\mu \left[\prod_{i=1}^k (h^\delta, f_i)^{l_i} \right] \rightarrow \mathbf{E} \left[\prod_{i=1}^k \left(\frac{1}{\sqrt{2\pi}} \Phi_{\mathbb{H}}^{\text{Neu}}, f_i \right)^{l_i} \right], \quad \text{as } \delta \rightarrow 0,$$

where \mathbf{E} is the expectation associated with $\Phi_{\mathbb{H}}^{\text{Neu}}$.

Proof For a function g , we write $g(a; b) = g(a) - g(b)$. To simplify the exposition, we only treat the case of the second moment; the other cases are similar but with heavier notation. To start with, note that

$$\begin{aligned} \mu[(h^\delta, f_1)(h^\delta, f_2)] &= \int_{\mathbb{H}^4} \mu[h^\delta(a_1; b_1)h^\delta(a_2; b_2)] \frac{f_1^+(a_1)f_1^-(b_1)}{Z_{f_1}} \\ &\quad \frac{f_2^+(a_2)f_2^-(b_2)}{Z_{f_2}} da_1 db_1 da_2 db_2. \end{aligned} \tag{5.20}$$

Let $\rho > 0$ be such that $\Im(z) \geq \rho$ whenever $z \in \text{Supp}(f_1) \cup \text{Supp}(f_2)$, and let

$$H(a_1, b_1, a_2, b_2) = -\frac{1}{2\pi^2} \Re \log \frac{(a_1 - a_2)(b_1 - b_2)(\bar{a}_1 - a_2)(\bar{b}_1 - b_2)}{(a_1 - b_2)(b_1 - a_2)(\bar{a}_1 - b_2)(\bar{b}_1 - a_2)}.$$

By Proposition 5.6, since all relevant points have imaginary parts greater than ρ , we have

$$\mu[h^\delta(a_1; b_1)h^\delta(a_2; b_2)] = H(a_1, b_1, a_2, b_2) + o(1), \tag{5.21}$$

where the error $o(1)$ is **uniform** over

$$\mathcal{D}_\delta := \{(a_1, b_1, a_2, b_2) \in \mathbb{H}^4 : D \geq \delta^\beta\},$$

where, as before, $D = D(a_1, b_1, a_2, b_2)$ denotes the minimal distance in the complex plane between any pair of points within $\{a_1, b_1, a_2, b_2\}$. We now split the integral in (5.20) into the integral over \mathcal{D}_δ and over \mathcal{D}_δ^c . Now the important observation is that since the the error is uniform, the limit of the integral over \mathcal{D}_δ is given by

$$\begin{aligned} & \int_{\mathbb{H}^4} H(a_1, \dots, b_2) \frac{f_1^+(a_1)f_1^-(b_1)}{Z_{f_1}} \frac{f_2^+(a_2)f_2^-(b_2)}{Z_{f_2}} da_1 db_1 da_2 db_2 \\ &= \frac{1}{2\pi^2} \mathbf{E}[(\Phi, f_1)(\Phi, f_2)]. \end{aligned}$$

Therefore, we are left with proving that the contribution to the integral (5.20) coming from \mathcal{D}_δ^c is negligible. To do that, we proceed somewhat crudely, noting that by Cauchy–Schwarz,

$$\mu[h^\delta(a_1; b_1)h^\delta(a_2; b_2)] \leq (\text{Var}_\mu(h^\delta(a_1; b_1)) \text{Var}_\mu(h^\delta(a_2; b_2)))^{1/2}. \tag{5.22}$$

Since the volume of \mathcal{D}_δ^c is polynomially small in δ , it will therefore suffice to show that

$$\text{Var}_\mu(h^\delta(a; b)) = O(\log \delta)^C \tag{5.23}$$

for some $C > 0$ and arbitrary points a, b within some fixed compact of \mathbb{H} . To prove this, we will go back to the definition of the height function as a sum of increments over a path, and we will use our *a priori* bound on K^{-1} coming from Proposition 3.4, which gives

$$K^{-1}(u, v) = O\left((\log \text{dist}(u, v))^C \frac{1}{\text{dist}(u, v)}\right).$$

Fix two paths γ_1, γ_2 from a to b . It will be advantageous to take these paths at positive macroscopic distance from one another except near the endpoints, where they must necessarily come together. We will explain more precisely below how we construct them. By the triangle inequality and Theorem 5.3, we have

$$\begin{aligned} \text{Var}_\mu(h^\delta(a; b)) &\leq \sum_{e_1 \in \gamma_1, e_2 \in \gamma_2} |\text{Cov}_\mu(\mathbf{1}_{\{e_1 \in \mathcal{M}\}}, \mathbf{1}_{\{e_2 \in \mathcal{M}\}})| \\ &\lesssim \sum_{e_1 \in \gamma_1, e_2 \in \gamma_2} \left((\log \text{dist}(e_1, e_2))^C \frac{1}{\text{dist}(e_1, e_2)} \right)^2 \end{aligned}$$

$$\lesssim (\log \delta)^{2C} \sum_{e_1 \in \gamma_1, e_2 \in \gamma_2} \left(\frac{1}{\text{dist}(e_1, e_2)} \right)^2.$$

We may assume that near a and b , the paths γ_1 and γ_2 form straight segments with different directions (say opposite directions), until they reach a fixed positive distance α , taken to be small enough that these paths remain at positive distance from the real line. The segments near a and b are then joined by portions of paths staying at distance (in the plane) at least $\alpha/2$ from one another to form γ_1 and γ_2 . Then

$$\begin{aligned} \sum_{e_1 \in \gamma_1, e_2 \in \gamma_2} \left(\frac{1}{\text{dist}(e_1, e_2)} \right)^2 &\leq \sum_{r=1}^{O(1/\delta)} \frac{1}{r^2} \#\{(e_1, e_2) : \text{dist}(e_1, e_2) = r\} \\ &\lesssim \sum_{r=R}^{\alpha/(2\delta)} \frac{1}{r^2} r + O(1) \leq \log(\delta^{-1}). \end{aligned} \tag{5.24}$$

We provide a brief explanation for the crucial point above, which is the bound on $\#\{e_1, e_2 : \text{dist}(e_1, e_2) = r\}$, and which comes from the choice of paths γ_1 and γ_2 . Indeed, if $1 \leq r \leq \alpha/2$ fixed, then elementary geometric considerations imply that any choice of e_1 in the segment of γ_1 at distance at most r from a , will give at most one corresponding point e_2 on γ_2 such that $\text{dist}(e_1, e_2) = r$. On the other hand, for $r \geq \alpha/2$, there are at most $O(\delta^{-2})$ pairs of edges on the whole path, so $\#\{(e_1, e_2) : \text{dist}(e_1, e_2) \geq \alpha/(2\delta)\} = O(\delta^{-2})$, so the contribution of such edges to the sum is indeed $O(1)$ as claimed above.

This proves (5.23) and therefore completes the proof of Theorem 5.8 in the case of second moments. In the general case of a moment of order $k \geq 2$, the same proof works, where we replace the use of Cauchy–Schwarz in (5.22) by a Hölder inequality, so that it suffices to show that $\mu[(h^\delta(a; b))^{2k}] \lesssim (\log 1/\delta)^{Ck}$ (we wrote here a moment of order $2k$ rather than k to account for the possibility that k is odd). We therefore need to choose $2k$ paths leading from a to b . As above, these paths may be chosen as being straight line segments up to a small distance α (in the plane) away from a or b , and with distinct directions; we simply choose the angles between these segments to be π/k , and otherwise require that these paths stay at positive distance from one another. It is easy to check that the analogue of (5.24) holds also in this case.

A standard argument says that since all moments of h^δ converge to the corresponding moments of $\frac{1}{\sqrt{2\pi}} \Phi_{\mathbb{H}}^{\text{Neu}}$, and since $\Phi_{\mathbb{H}}^{\text{Neu}}$ is a Gaussian process, we can conclude that $h^\delta \rightarrow \frac{1}{\sqrt{2\pi}} \Phi_{\mathbb{H}}^{\text{Neu}}$ in distribution as $\delta \rightarrow 0$ (in the sense of finite dimensional distributions, where Φ is viewed as a stochastic process indexed by smooth test functions with compact support and mean zero). This is straightforward to see: the convergence of the second moments, say, implies tightness. On the other hand, since the limiting moments are Gaussian, they identify the limiting distribution uniquely. As a result we have proved Theorem 1.3.

Acknowledgements This project was initiated when all three authors were at the University of Cambridge. During the project, M.L. was mainly at the University of Vienna. W.Q. was at the University of Cambridge.

and later at the CNRS and Laboratoire de Mathématiques d'Orsay, Université Paris-Saclay. N.B.'s work was supported by: EPSRC grant EP/L018896/1, University of Vienna start-up grant, and FWF grant P33083 on "Scaling limits in random conformal geometry". The authors are also grateful for an invitation to visit the Département de Mathématiques et Applications at Ecole Normale Supérieure in April 2018, where part of this work took place, and during which N.B. was an invited professor. The hospitality of that department, and the stimulating atmosphere of the discussions (including with D. Chelkak, who helped us with the computation (2.10)) are gratefully acknowledged. We also thank the anonymous referee who brought to our attention the alternative strategy outlined in Remark 1.6, and the connection to isoradial walks in (1.8). Finally, we also thank Mingchang Liu for pointing out an error in an earlier version of the paper, and to another anonymous referee for their careful reading and insightful comments.

Funding Open access funding provided by Austrian Science Fund (FWF).

Data Availability Data sharing is not applicable to this article as data was generated nor analyzed during its preparation.

Declarations

Conflict of interest The authors state that there is no conflict of interest.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Aru, J., Sepúlveda, A., Werner, W.: On bounded-type thin local sets of the two-dimensional gaussian free field. *J. Inst. Math. Jussieu* **18**(3), 591–618 (2019)
2. Bai, T., Wan, Y.: On the crossing estimates of simple conformal loop ensembles. *Int. Math. Res. Notices*, 2022, rnac173
3. Barlow, M. T.: Random walks and heat kernels on graphs, volume 438. Cambridge University Press, Cambridge (2017)
4. Basok, M., Chelkak, D.: Tau-functions à la dubédat and probabilities of cylindrical events for double-dimers and CLE(4). arXiv preprint [arXiv:1809.00690](https://arxiv.org/abs/1809.00690), 2018
5. Berestycki, N., Laslier, B., Ray, G.: Dimers on Riemann surfaces, I: Temperleyan forests. arXiv (2019)
6. Berestycki, N., Laslier, B., Ray, G.: Dimers and imaginary geometry. *Ann. Probab.* **48**(1), 1–52 (2020)
7. Berestycki, N., Liu, M.: Piecewise Temperleyan dimers and a multiple SLE_g. In preparation, 2022+
8. Berestycki, N., Powell, E.: Gaussian free field, multiplicative chaos and Liouville quantum gravity. to appear (2020)
9. Boutillier, C., De Tilière, B.: Height representation of XOR-Ising loops via bipartite dimers. *Electron. J. Probab.* **19** (2014)
10. de Tilière, B.: Scaling limit of isoradial dimer models and the case of triangular quadri-tilings. *Annales de l'IHP Probabilités et statistiques* **43**(6), 729–750 (2007)
11. Dubédat, J.: Exact bosonization of the Ising model. arXiv preprint [arXiv:1112.4399](https://arxiv.org/abs/1112.4399) (2011)
12. Dubédat, J.: Double dimers, conformal loop ensembles and isomonodromic deformations. *J. Eur. Math. Soc.* **21**(1), 1–54 (2018)
13. Duminil-Copin, H., Lis, M.: On the double random current nesting field. *Probab. Theory Relat. Fields* **175**(3–4), 937–955 (2019)
14. Folz, M.: Gaussian upper bounds for heat kernels of continuous time simple random walks. *Electron. J. Probab.* **16**, 1693–1722 (2011)

15. Freedman, D. A.: On tail probabilities for martingales. *Ann. Probab.* 100–118 (1975)
16. Giuliani, A., Jauslin, I., Lieb, E.H.: A pfaffian formula for monomer-dimer partition functions. *J. Stat. Phys.* **163**(2), 211–238 (2016)
17. Grigor'yan, A.: Gaussian upper bounds for the heat kernel on arbitrary manifolds. *J. Diff. Geom.* **45**, 33–52 (1997)
18. Heilmann, O.J., Lieb, E.H.: Monomers and dimers. *Phys. Rev. Lett.* **24**(25), 1412 (1970)
19. Heilmann, O.J., Lieb, E.H.: Theory of monomer-dimer systems. *Commun. Math. Phys.* **25**(3), 190–232 (1972)
20. Jerrum, M.: Two-dimensional monomer-dimer systems are computationally intractable. *J. Stat. Phys.* **48**(1–2), 121–134 (1987)
21. Kenyon, R.: Local statistics of lattice dimers. *Annales de l'Institut Henri Poincaré (B) Probability and Statistics*, 33(5):591–618 (1997)
22. Kenyon, R.: Conformal invariance of domino tiling. *Ann. Probab.* 759–795 (2000)
23. Kenyon, R.: Dominos and the Gaussian free field. *Ann. Probab.* **29**(3), 1128–1137 (2001)
24. Kenyon, R.: Laplacian and Dirac operators on critical planar graphs. *Invent. Math.* **150**(2), 409–439 (2002)
25. Kenyon, R.: Conformal invariance of loops in the double-dimer model. *Commun. Math. Phys.* **326**(2), 477–497 (2014)
26. Kenyon, R., Wilson, D.: Boundary partitions in trees and dimers. *Trans. Am. Math. Soc.* **363**(3), 1325–1364 (2011)
27. Kenyon, R. W., Propp, J. G., Wilson, D. B.: Trees and matchings. *Electron. J. Combin.*, 7, (2000)
28. Kenyon, R.W., Sheffield, S.: Dimers, tilings and trees. *J. Combin. Theory Ser. B* **92**(2), 295–317 (2004)
29. Laslier, B., Toninelli, F.L.: Lozenge tilings, glauher dynamics and macroscopic shape. *Commun. Math. Phys.* **338**(3), 1287–1326 (2015)
30. Lawler, G.F., Limic, V.: *Random walk: a modern introduction*. Cambridge Studies in Advanced Mathematics, vol. 123. Cambridge University Press, Cambridge (2010)
31. Lyons, R., Peres, Y.: *Probability on trees and networks*, volume 42. Cambridge University Press (2017)
32. Miller, J., Sheffield, S., Werner, W.: CLE percolations. *Forum Math. Pi*, 5:e4, 102, 2017
33. Percus, J.K.: One more technique for the dimer problem. *J. Math. Phys.* **10**(10), 1881–1884 (1969)
34. Popov, S.: *Two-Dimensional Random Walk: From Path Counting to Random Interlacements*. Institute of Mathematical Statistics Textbooks, Cambridge University Press (2021)
35. Qian, W., Werner, W.: Coupling the Gaussian free fields with free and with zero boundary conditions via common level lines. *Commun. Math. Phys.* **361**(1), 53–80 (2018)
36. Russkikh, M.: Dimers in piecewise temperleyan domains. *Commun. Math. Phys.* **359**(1), 189–222 (2018)
37. Sheffield, S.: Exploration trees and conformal loop ensembles. *Duke Math. J.* **147**(1), 79–129 (2009)
38. Sheffield, S., Werner, W.: Conformal loop ensembles: the Markovian characterization and the loop-soup construction. *Ann. Math.*, 1827–1917 (2012)
39. Thurston, W.P.: Conway's tiling groups. *Am. Math. Mon.* **97**(8), 757–773 (1990)
40. Toninelli, F.: Lecture notes on the dimer model, 2019. <http://math.univ-lyon1.fr/homes-www/toninelli/noteDimeri.pdf>

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.