

Bridge Decomposition of Restriction Measures

Tom Alberts · Hugo Duminil-Copin

Received: 22 October 2009 / Accepted: 24 May 2010 / Published online: 5 June 2010
© Springer Science+Business Media, LLC 2010

Abstract Motivated by Kesten’s bridge decomposition for two-dimensional self-avoiding walks in the upper half plane, we show that the conjectured scaling limit of the half-plane SAW, the SLE(8/3) process, also has an appropriately defined bridge decomposition. This continuum decomposition turns out to entirely be a consequence of the restriction property of SLE(8/3), and as a result can be generalized to the wider class of *restriction measures*. Specifically we show that the restriction hulls with index less than one can be decomposed into a Poisson Point Process of *irreducible bridges* in a way that is similar to Itô’s excursion decomposition of a Brownian motion according to its zeros.

Keywords Self-avoiding walk · Conformal invariance · SLE · Restriction measures · Poisson process · Excursion theory

1 Introduction

One of the greatest successes of the Schramm-Loewner Evolution (SLE), and the broader study of two-dimensional conformally invariant stochastic processes that it enabled, has been the ability to describe the scaling limits of two-dimensional lattice models that arise in statistical mechanics. There are many known examples: SLE(2) as the scaling limit for loop erased random walk, SLE(3) as the scaling limit of critical Ising interfaces, SLE(6) as the limit of percolation exploration paths, etc. One of the most important open problems in the

Research of Tom Alberts supported in part by NSF Grant OISE 0730136, and a postdoctoral fellowship from the Natural Sciences and Engineering Research Council of Canada. Research of Hugo Duminil-Copin supported in part by project MRTN-CT-2006-035651, Acronym CODY, of the European Commission, and a grant from the Swiss National Science Foundation.

T. Alberts (✉)
Department of Mathematics, University of Toronto, Toronto, ON, Canada
e-mail: alberts@math.toronto.edu

H. Duminil-Copin
Département de Mathématiques et Applications, École Normale Supérieure, Paris, France
e-mail: hugo.duminil@ens.fr

field is to prove that the scaling limit of the infinite *self-avoiding walk* in the upper half plane \mathbb{H} is given by SLE(8/3). It is known that *if* the scaling limit of half-plane SAWs exists *and* is conformally invariant, then the scaling limit must be SLE(8/3). Both the existence and conformal invariance are widely believed to be true, yet proofs remain elusive. For an accessible and relatively recent source on the current status of this problem, we refer the reader to [11]. Even without formally establishing the scaling limit result, it is often still possible to independently check that the various well-studied properties of half-plane SAWs carry over to the SLE(8/3) process. The main results of this paper should be seen in this context. In [6] it is shown that half-plane SAWs admit what is called a *bridge decomposition*, which raised the question of finding a similar decomposition for SLE(8/3). In this paper we will show that an appropriately defined continuum decomposition does exist, and we will describe some of its properties. A somewhat surprising aspect of the existence is that it depends only on the fact that SLE(8/3) satisfies the restriction property, and not on the fine details of the process itself. Specifically, the decomposition has no explicit reliance on the Loewner equation. Using this fact we are able to extend the continuum bridge decomposition beyond SLE(8/3) to a wider class of random sets whose laws are given by the so-called restriction measures. These probability measures were introduced and studied extensively in [8], and they occupy an important position in the hierarchy of two-dimensional conformally invariant processes. We will give a more detailed description of restriction measures in Sect. 2, but we emphasize that the reader who is uninterested in general restriction measures will lose nothing by focusing on SLE(8/3) as the canonical one.

1.1 Motivation: Bridge Decomposition of SAWs

To motivate the continuum bridge decomposition, we first describe the corresponding decomposition for half-plane SAWs. This is thoroughly described in [14], along with many other interesting properties of the self-avoiding walk. In the discrete setting we will work exclusively on the lattice $\mathbb{Z} + i\mathbb{Z}$. An N -step self-avoiding walk ω on $\mathbb{Z} + i\mathbb{Z}$ is a sequence of lattice sites $[\omega(0), \omega(1), \dots, \omega(N)]$ satisfying $|\omega(j + 1) - \omega(j)| = 1$ and $\omega(i) \neq \omega(j)$ for $i \neq j$. We will write $|\omega| = N$ to denote the length of ω . Given walks ω and ω' of length N and M (respectively), the *concatenation* of ω and ω' is defined by

$$\omega \oplus \omega' = [\omega(0), \dots, \omega(N), \omega'(1) + \omega(N), \dots, \omega'(M) + \omega(N)].$$

Letting c_N denote the number of self-avoiding walks of length N , it is easy to see that

$$c_{N+M} \leq c_N c_M$$

since any SAW of length $N + M$ can always be written as the concatenation of two SAWs of length N and M . A standard submultiplicativity argument then proves the existence of a constant $\mu > 0$ such that

$$\lim_{N \rightarrow \infty} \frac{\log c_N}{N} = \log \mu, \tag{1}$$

or $c_N \approx \mu^N$ in the common shorthand. The exact value of μ is not known, nor is it expected to be any special value, but numerically it has been shown that μ is close to 2.638 (see [14, Section 1.2]).

We will mostly deal with half-plane SAWs rooted at the origin, i.e. self-avoiding paths ω such that $\omega(0) = 0$ and $\text{Im } \omega(j) > 0$ for all $j > 0$. Let \mathcal{H} denote the set of all such walks.

The most commonly used probability measure on \mathcal{H} , and the one that we will consider throughout, is the weak limit of the uniform measure on $\{\omega \in \mathcal{H} : |\omega| = N\}$, as $N \rightarrow \infty$. This limit is proven to exist in [14], and again in the appendix of [11]. The key element of both proofs is, in fact, the *bridge decomposition* of the walks in \mathcal{H} , the study of which was initiated by Kesten [6, 7] and goes as follows. A *bridge of length N* is a self-avoiding walk ω such that $|\omega| = N$ and

$$\text{Im } \omega(0) < \text{Im } \omega(j) \leq \text{Im } \omega(N), \quad 1 \leq j \leq N.$$

Note that the concatenation of any two bridges is still a bridge, but that not every bridge is the concatenation of two shorter ones. A bridge with the latter property is said to be *irreducible*, and such bridges are the basic building blocks of walks in \mathcal{H} . Indeed, given any $\omega \in \mathcal{H}$, one performs a bridge decomposition of ω by searching for the smallest time j such that $\text{Im } \omega(k) \leq \text{Im } \omega(j)$ for $k \leq j$ and $\text{Im } \omega(k) > \text{Im } \omega(j)$ for $k > j$. By the minimality of j , the subpath $[w(0), w(1), \dots, w(j)]$ is an irreducible bridge, and the shifted subpath $[0, w(j + 1) - w(j), \dots, w(k) - w(j), \dots]$ for $k \geq j$ is a new element of \mathcal{H} on which we may repeat this procedure. Iterating in this fashion produces the bridge decomposition of ω into a sequence of irreducible bridges, and the decomposition is clearly unique.¹

Much of the study of the infinite self-avoiding walk in the upper half plane therefore reduces to the study of irreducible bridges. Let \mathcal{B} be the set of all irreducible bridges rooted at the origin, and λ_N be the number of length N elements of \mathcal{B} . Using some clever tricks involving generating functions, Kesten was able to prove what is now called **Kesten’s relation**:


$$\sum_{N \geq 1} \lambda_N \mu^{-N} = \sum_{\omega \in \mathcal{B}} \mu^{-|\omega|} = 1, \tag{2}$$

for the same μ as in (1) (for proofs see [6] or [14, Section 4.3]). Kesten’s relation shows that $\mathbf{P}(\omega) := \mu^{-|\omega|}$ is a probability measure on \mathcal{B} , and by concatenating together an independent sequence of irreducible bridges each sampled from \mathbf{P} , a probability measure is induced on \mathcal{H} . In [14] and [11], the latter measure is shown to be the only possible candidate for the weak limit of the uniform measure on $\{\omega \in \mathcal{H} : |\omega| = N\}$, and therefore the question of existence of this weak limit is immediately settled.

The bridge decomposition shows that infinite half-plane SAWs have a renewal structure to them. At the end of each irreducible bridge the future path of the walk lies entirely in the half-plane above the horizontal line where the bridge ended. The future path is again a concatenation of a sequence of irreducible bridges, so that its law is the same as the law of the original path and the future path is independent of the past. In this sense the walk renews itself whenever it is at the end of an irreducible bridge, and it is appropriate to call such times renewal times. Note that the renewal times are functions of the *entire* half-plane SAW, since the algorithm for the bridge decomposition depends upon knowing the entire walk.

¹There is a minor technicality to point out here: if the walk oscillates infinitely often in the vertical direction without approaching some limit (including infinity) the decomposition algorithm will terminate after finitely many iterations and the remaining part of the walk will not be a bridge. However, we will see in the next paragraph that this is a probability zero event under the standard measure on \mathcal{H} , and that the vertical component of the SAW always goes to infinity with probability one.

Fig. 1 A sample SLE(8/3) curve in the lighter color, with the bridge points superimposed in black. The bridge heights are plotted on the vertical axis. The SLE(8/3) curve is generated by Tom Kennedy’s algorithm and freely available graphics program; see [5]



1.2 Statement of Results: The Continuum Bridge Decomposition

In the continuum we will show that an analogue of bridge times exists for the so-called *restriction hulls* in \mathbb{H} , and that these times are also renewal times. Using this renewal structure, we proceed to decompose the restriction hulls into countably many continuum irreducible bridges. This continuum decomposition most closely resembles the discrete one in the case of SLE(8/3), but we will see that it also holds for more general restriction hulls with parameter $\alpha < 1$. We will give a more in-depth description of the restriction hulls in Sect. 2, but provide a brief summary here.

Roughly speaking, a restriction hull is a stochastic process taking values in the space of unbounded hulls in \mathbb{H} . An *unbounded hull* is a closed, connected subset $K \subset \overline{\mathbb{H}}$ such that $\mathbb{H} \setminus K$ consists of exactly two connected components. The unbounded hulls that we will consider are closed, connected subsets of $\overline{\mathbb{H}}$ that connect 0 and ∞ , and intersect \mathbb{R} only at zero; moreover it will be possible to time parameterize them into a growing family $(K_t, t \geq 0)$ of *hulls* (closed, connected subsets A of $\overline{\mathbb{H}}$ such that $\mathbb{H} \setminus A$ is simply connected with exactly one connected component) with $K_\infty = K$. This time parameterization is provided by the well-known construction of restriction hulls that was originally laid out in [8] and [13]. Those papers show that attaching the filled-in loops from a realization of the Brownian Loop Soup to an independent SLE curve induces a restriction law on unbounded hulls in \mathbb{H} . By changing the κ parameter for the SLE and the intensity parameter for the loop soup (in a specific way) an entire family \mathbb{P}_α of restriction measures on unbounded hulls is created. Here α is a real parameter with $\alpha \geq 5/8$.

The definition of a continuum bridge is motivated by the algorithm for decomposing half-plane SAWs into irreducible bridges, which essentially searches for horizontal lines that separate the future path from the past.

Definition 1 Let K be a hull (unbounded or not).

- Call $L > 0$ a *bridge height* for K if the horizontal line $y = L$ intersects K at exactly one point, i.e. if $K \cap \{y = L\}$ is a singleton.
- If $z \in \mathbb{H}$ is such a singleton then we call it a *bridge point*. Let C be the set of bridge points of K , and let D be the set of bridge heights (note that $D = \{\text{Im } z : z \in C\}$).

- Let G be the set of *bridge times* at which the hull is at a bridge point, which can be written as $G := \{t \geq 0 : K_t \setminus K_{t-} \cap C \neq \emptyset\}$.
- A *continuum bridge* is a segment of the bridge between two bridge times, i.e. if $s, t \in G$ with $s < t$ then the hull $K_{t-} \setminus K_s$ is a bridge. A continuum bridge is said to be *irreducible* if it contains no bridge points (other than the starting and ending points).

Note that bridge heights, points and times are all functions of the *entire* hull K . A subset of K is, by itself, not enough to determine C, D or G . At any fixed time $t \geq 0$ it is possible to determine what are the bridge points of the hull K_t , but not which of those are bridge points of the entire hull $K_\infty = K$, since some of the bridge points of K_t may ultimately be destroyed by the future hull as it grows.

There are two main steps behind the continuum bridge decomposition. The first is to show that bridge points actually exist for hulls with $\alpha < 1$, which is not a priori clear. We do this by calculating the almost sure Hausdorff dimensions of C and D and showing that they are strictly larger than zero (and in fact the same). Specifically we will show the following:

Theorem 1 *Suppose K has the law of \mathbb{P}_α , then*

1. *the laws of C and D are scale invariant (i.e. $rC \equiv C$ and $rD \equiv D$ for all $r > 0$),*
2. *C and D are almost surely perfect (i.e. closed and without isolated points),*
3. *the Hausdorff dimensions of both C and D are constant, $\mathbb{P}_\alpha - a.s.$,*
4. *$\dim_{\mathbb{H}} C = \dim_{\mathbb{H}} D = \max(2 - 2\alpha, 0)$, $\mathbb{P}_\alpha - a.s.$,*
5. *C and D are empty, $\mathbb{P}_\alpha - a.s.$ if and only if $\alpha \geq 1$.*

The proof of Theorem 1 is taken up in Sect. 3, but we will mention here that the key element is the **restriction formula**:

$$\mathbb{P}_\alpha(K \cap A = \emptyset) = \phi'_A(0)^\alpha, \tag{3}$$

where A is a hull that does *not* contain zero, and ϕ_A is a conformal map from $\mathbb{H} \setminus A$ to \mathbb{H} such that $\phi_A(z) \sim z$ as $z \rightarrow \infty$. Most of the proof of Theorem 1 is based on an analysis of $\phi'_A(0)$ for a specific choice of the hull A . The proof of part (5) builds upon the $\alpha = 1$ case, which is related to Brownian excursions, and uses the fact that the vertical component of a Brownian excursion is a Bessel-3 process.

Given that bridge points exist for $\alpha < 1$, the next step is to prove an analogue of the renewal theory for half-plane SAWs. In Sect. 4 we show that the restriction hulls have an extended Markov property with respect to the information gained by observing the hull as it grows along with the *global* bridge points of K as they appear, and as a corollary we show that the bridge times are actually renewal times for the hull process. In Sect. 5 we will use this Markov property and Theorem 1 to show the existence of a “local time” for the time spent by a restriction hull at its bridge points, and the local time can then be used to prove:

Theorem 2 *There exists a local time λ supported on bridge heights such that $\theta_\lambda(K_\lambda \setminus K_{\lambda-})$ is a Poisson Point Process, where θ_t is an operator that shifts back to the origin the part of the hull that comes after time t . Moreover, the local time is the inverse of a stable subordinator of index $2 - 2\alpha$.*

The general theory of Poisson Point Processes then implies the existence of a sigma-finite measure ν_α on continuum irreducible bridges that is the analogue of the measure \mathbf{P} on irreducible bridges for half-plane SAWs. In Sect. 5 we mention some basic properties of this

measure. We also show that the Poisson Point Process can be used to recover the restriction hull, so that as in the discrete case, the irreducible bridges are the building blocks of the restriction hull processes.

We should mention that most of these ideas are similar in spirit to the excursion decomposition of a one-dimensional Brownian motion according to its zeros, as was first described by Itô. In recent years, similar two-dimensional conformally invariant decompositions of this type have also been considered by Dubédat [4] and Virág [17]. They provide decompositions of unbounded hulls arising from certain variants of $\text{SLE}(\kappa, \rho)$ and Brownian excursions, respectively, although their decompositions are at cutpoints rather than bridge points (i.e. points that, if removed from the set, would disconnect it into two pieces). Clearly bridge points are cutpoints but not vice versa, and there does not appear to be any direct relationship between our decomposition and theirs. In one sense their decompositions are more involved than ours, since their hulls refresh at cutpoints only after conformally mapping away the past, whereas our hulls refresh at bridge points after a simple shift of the future hull back to the origin. This difference is mostly cosmetic, however, and in spirit all these decompositions are quite similar.

The paper is organized as follows: in Sect. 2 we give the necessary background on restriction measures and introduce some notation. Section 3 is devoted to proving the existence of bridge points and Theorem 1, while Sect. 4 proves an extended Markov property and a refreshing property of the restriction hulls with respect to the filtration generated by bridge points as they appear. Section 5 then uses these results to prove the decomposition of Theorem 2. Finally, in Sect. 6 we present a series of open questions that were raised by our work.

2 Restriction Measures

In this section we review the basic construction and properties of restriction measures. We include no proofs but give references to the appropriate sources. For thorough overviews of the subject see [8, 10, 13]. The reader interested only in the bridge decomposition for $\text{SLE}(8/3)$, and not for general restriction measures, can entirely ignore the presence of the loops in this section.

To begin with, consider a simply connected domain D in the complex plane \mathbb{C} (other than the whole plane itself) and two boundary points $z, w \in \partial D$. A *chordal restriction measure* corresponding to the triple (D, z, w) is a probability measure $\mathbb{P}^{(D, z, w)}$ on closed subsets of \overline{D} . The measures are supported on closed, connected subsets of $K \subset \overline{D}$ such that $K \cap \partial D = \{z, w\}$ and $D \setminus K$ has exactly two components (for the triple $(\mathbb{H}, 0, \infty)$ we call these sets unbounded hulls, for obvious reasons). The restriction measures satisfy the following properties, which essentially characterize them uniquely:

- **Restriction property:** for all simply connected subsets D' of D such that $D \setminus D'$ is also simply connected and bounded away from z and w , the law of $\mathbb{P}^{(D, z, w)}$, conditioned on $K \subset D'$, is $\mathbb{P}^{(D', z, w)}$,
- **Conformal invariance:** if $f : D \rightarrow D'$ is conformal and K has $\mathbb{P}^{(D, z, w)}$ as its law, then $f(K)$ is distributed according to $\mathbb{P}^{(f(D), f(z), f(w))}$.

It turns out that for a given triple (D, z, w) there is only a one-parameter family of such laws, indexed by a real number α . We denote the law by $\mathbb{P}_\alpha^{(D, z, w)}$, and due to the conformal invariance property it is enough to define the restriction measure for a single triple (D, z, w) . The canonical choice is $(\mathbb{H}, 0, \infty)$, and for shorthand we will write \mathbb{P}_α for $\mathbb{P}_\alpha^{(\mathbb{H}, 0, \infty)}$. In [8]

it is shown that these restriction measures exist only if the parameter α satisfies $\alpha \geq 5/8$, and that the measure is supported on simple curves only if $\alpha = 5/8$. In the latter case the restriction measure is simply the SLE(8/3) law from z to w in D . For $\alpha = 1$ it turns out that the restriction measure coincides with the law of filled-in Brownian excursions in D from z to w .

For all $\alpha \geq 5/8$, one of the fundamental constructions of [8] is that restriction measures can be realized by adding to an SLE(κ) curve the filled-in loops that it intersects from an independent realization of the Brownian loop soup, for an appropriate choice of κ for the curve and intensity parameter λ for the loop soup. Let

$$\kappa = \frac{6}{2\alpha + 1}, \quad \lambda = (8 - 3\kappa)\alpha,$$

and let γ be a chordal SLE(κ) and \mathcal{L}_λ be an independent realization of the Brownian loop soup (in \mathbb{H}) with intensity parameter λ . The individual loops in \mathcal{L}_λ will be generically denoted by η , they can be thought of as continuous curves $\eta : [0, t_\eta] \rightarrow \mathbb{H}$ such that $\eta(0) = \eta(t_\eta)$. Throughout we will use γ and η to denote the curves as well as their traces, i.e. $\gamma[0, \infty)$ and $\eta[0, t_\eta]$, respectively. It will be clear from the context which we are referring to. Let K be the hull generated by the union of γ and all the (filled-in) $\eta \in \mathcal{L}_\lambda$ such that $\eta \cap \gamma \neq \emptyset$. Then [8] (along with [13]) proves that K is distributed according to \mathbb{P}_α .

This construction allows us to identify restriction hulls with pairs (γ, \mathcal{L}) , where $\gamma : [0, t_\gamma] \rightarrow \mathbb{C}$ is a continuous, simple curve and \mathcal{L} is a set of loops. Furthermore, the curve plus loops structure gives a clean way of time parameterizing the hulls. Letting K be a restriction hull, which we identify with (γ, \mathcal{L}) , we define K_t to be the hull generated by $\gamma[0, t]$ plus the union of all filled-in loops $\eta \in \mathcal{L}_\lambda$ such that $\eta \cap \gamma[0, t] \neq \emptyset$. Then $(K_t)_{t \geq 0}$ is a growing family of hulls that increases to $K_\infty = K$. It is important for us to have such a time parameterization so that we may properly describe the renewal theory for the restriction hulls, but the particular time parameterization is not especially important since we are mostly interested in the restriction hull as a topological object. We remark that this growing family is not continuous with respect to the time parametrization, since loops are added “all at once”, but again it does not really matter for our purposes (nevertheless, notice that the parameterization is right continuous). The only issue to point out is that the bridge points of a restriction hull will always be a subset of the underlying (simple) curve γ , and therefore to each bridge point there is a corresponding unique bridge time. Hence the set of bridge times G is a well defined object.

The curve-plus-loops structure also makes it easy to define various operations on hulls. Given two pairs (γ, \mathcal{L}) and $(\gamma^*, \mathcal{L}^*)$ with $\gamma(0) = \gamma^*(0) = 0$, their *concatenation* is defined by

$$(\gamma, \mathcal{L}) \oplus (\gamma^*, \mathcal{L}^*) = (\gamma \oplus \gamma^*, \mathcal{L} \cup (\gamma(t_\gamma) + \mathcal{L}^*)),$$

where $\gamma \oplus \gamma^*$ is the usual concatenation of curves given by

$$(\gamma \oplus \gamma^*)(t) = \begin{cases} \gamma(t), & 0 \leq t \leq t_\gamma \\ \gamma^*(t - t_\gamma) + \gamma(t_\gamma), & t_\gamma \leq t \leq t_\gamma + t_{\gamma^*}. \end{cases}$$

We also define a time shift for the hulls. For $t \leq s \leq t_\gamma$, define the curve $\gamma^{t,s}$ by $\gamma^{t,s}(t') := \gamma(t + t')$ for $0 \leq t' \leq s - t$, and let

$$\mathcal{L}^{t,s} := \{\eta \in \mathcal{L} : \eta \cap \gamma^{t,s} \neq \emptyset, \eta \cap \gamma[0, t] = \emptyset\}.$$

Then we define $\Lambda_{t,s}K := (\gamma^{t,s}, \mathcal{L}^{t,s})$, which is the future hull between times t and s , and $\theta_{t,s}K := \Lambda_{t,s}K - \gamma(t)$, which shifts the future hull to start at the origin. If $s = t_\gamma$, which usually for us means $s = \infty$, we write Λ_t and θ_t for these operators. In the case that K is an unbounded hull in \mathbb{H} and t is a bridge time for K , it is easy to see that $\theta_t K$ is also an unbounded hull in \mathbb{H} . At non-bridge times $\theta_t K$ does not remain in \mathbb{H} .

Imagine a walker moving along the hull that has discovered K_t at time t . The information that is progressively revealed to the walker is encapsulated by the filtration

$$\mathcal{F}_t := \sigma(K_s; 0 \leq s \leq t).$$

With respect to this filtration, the following Domain Markov property is true:

$$\text{The conditional law of } \Lambda_t K, \text{ given } \mathcal{F}_t, \text{ is } \mathbb{P}_\alpha^{(\mathbb{H} \setminus \gamma[0,t], \gamma(t), \infty)}. \tag{4}$$

This is similar to the Domain Markov property for regular SLE, where the future curve is an independent SLE(κ) curve from $\gamma(t)$ to ∞ in $\mathbb{H} \setminus \gamma[0, t]$, except that in the case of restriction measures one also attaches to the curve the filled-in loops of an independent realization of the Brownian loop soup in the domain $\mathbb{H} \setminus \gamma[0, t]$. Note, however, that both the future curve and loops are sampled from the laws corresponding to the domains $\mathbb{H} \setminus \gamma[0, t]$, *not* the laws corresponding to $\mathbb{H} \setminus K_t$. In short, the future curve and future loops are allowed to intersect the past loops but *not* the past curve $\gamma[0, t]$.

For the domain $(\mathbb{H}, 0, \infty)$ recall that the restriction measures satisfy the restriction formula (3):

$$\mathbb{P}_\alpha(K \cap A = \emptyset) = \phi'_A(0)^\alpha,$$

where A is a hull in \mathbb{H} that is a positive distance from zero, and ϕ_A is a conformal map from $\mathbb{H} \setminus A$ onto \mathbb{H} satisfying $\phi_A(z) \sim z$ as $z \rightarrow \infty$. In fact, specifying the above probabilities for a sufficiently large class of hulls A (so-called *smooth hulls*) uniquely determines \mathbb{P}_α , see [8] for a proof of this fact. For general triples (D, z, w) , the restriction formula is

$$\mathbb{P}_\alpha^{(D,z,w)}(K \cap A = \emptyset) = \phi'_{f(A)}(0)^\alpha, \tag{5}$$

where A is a hull in D not containing z , and f is a conformal map from D onto \mathbb{H} that sends z to 0 and w to ∞ .

The restriction formula will be heavily used throughout this paper. For a given hull A there are various techniques from both complex analysis and probability theory that can be used to compute $\phi'_A(0)$. We will exclusively use probabilistic techniques involving Brownian motion; these are described in the next section.

3 Bridge Lines and Bridge Points

The main focus of this section is proving Theorem 1. Specifically, we establish the existence of bridge points and lines for restriction hulls with $\alpha < 1$, and also prove the non-existence for $\alpha \geq 1$.

First observe that part (1) of Theorem 1 is trivial. The scale invariance of C and D follows immediately from the scale invariance of the restriction hulls (which itself follows from the scale invariance of SLE and of the loop soup). To prove part (2), first recall that bridge points of a restriction hull are always on the SLE curve itself and never on a loop, and that there is

always a unique bridge time corresponding to every bridge point. We refer to the end of the section for the proof.

The most involved proofs are for calculating the Hausdorff dimensions of C and D . The computation of the Hausdorff dimensions in Theorem 1 follows standard “one-point” and “two-point” arguments, as in, for example, [1, 2, 9, 16]. The idea behind this argument is to approximate C and D by “thickened” sets C_ϵ and D_ϵ , and then obtain estimates on the probability that a given set of points belongs to the thickened sets. A specific bound on the probability that one point belongs to the thickened set gives an upper bound on the Hausdorff dimension, and a similar bound on the probability that two points are in the thickened sets, together with the order of magnitude of the one-point estimate, gives a lower bound on the dimension. We recall the result that we will use in the remainder; throughout this paper we use the notation $f(\epsilon) \asymp g(\epsilon)$ to indicate that there exists constants C_1 and C_2 independent of ϵ such that $C_1g(\epsilon) \leq f(\epsilon) \leq C_2g(\epsilon)$, for all ϵ sufficiently small.

Proposition 1 *Let H be a random subset of \mathbb{C} and H_ϵ be the set of points at distance less than ϵ from H . Suppose that the two following conditions are fulfilled for some $s \geq 0$ and constant $c > 0$:*

- for all $z \in \mathbb{H}$, $\mathbb{P}(z \in H_\epsilon) \asymp \epsilon^s$,
- for all distinct $w, z \in \mathbb{H}$, $\mathbb{P}(w, z \in H_\epsilon) \leq c\epsilon^s \wedge c(\epsilon^{2s}/|w - z|^s)$.

Then $\dim_{\mathbb{H}} H \leq 2 - s$ with probability one, and with some strictly positive probability we also have $\dim_{\mathbb{H}} H \geq 2 - s$. If H is a random subset of \mathbb{R} then the same conclusion holds with $2 - s$ replaced by $1 - s$.

Note that Proposition 1 by itself is not enough to conclude that the Hausdorff dimension of H is a constant, since the lower bound only holds on some event of positive probability. In our situation we are able to conclude that the Hausdorff dimension of C and D is constant by using a 0-1 law. The argument that follows uses the Blumenthal 0-1 Law and is modified from [9].

Proof (Proof of Theorem 1, part (3)) We will prove the result for C , a similar argument holds for D . For $0 \leq t \leq s$, define $C_t(s) := \{\text{bridge points of } K_s\} \cap K_t$. For a fixed $d > 0$, let $W_t(s) := \{\dim_{\mathbb{H}} C_t(s) \geq d\}$. It is enough to show that $\mathbb{P}_\alpha(W_\infty(\infty)) = 0$ or 1.

First note that for fixed s , both the sets $C_t(s)$ and $W_t(s)$ are increasing in t , while for fixed t they are decreasing in s . Defining

$$V_s := \bigcap_{n=1}^\infty W_{\frac{1}{n}}(s) = \{\dim_{\mathbb{H}} C_t(s) \geq d \forall 0 < t \leq s\},$$

it follows that V_s is also decreasing in s . For each element of the event $V_s \setminus V_\infty$, there exists a t_0 such that $0 < t_0 \leq s$ and for all $0 < t \leq t_0$,

$$\dim_{\mathbb{H}} C_t(\infty) < d \leq \dim_{\mathbb{H}} C_t(s).$$

But this can only happen if for every $0 < t \leq t_0$, the future hull $A_s K$ destroys bridge points of K_s that are in K_t , and since this happens for every $0 < t \leq t_0$ and $K_t \rightarrow \{0\}$ as $t \rightarrow 0$, this forces that the future hull comes arbitrarily close to the real axis. But this is clearly an event of measure zero. Hence for every $s > 0$, $\mathbb{P}_\alpha(V_s \setminus V_\infty) = 0$, from which it immediately

follows that

$$\mathbb{P}_\alpha \left(\bigcap_{n=1}^\infty V_{\frac{1}{n}} \right) = \mathbb{P}_\alpha(V_\infty).$$

However, the intersection of the $V_{1/n}$ is \mathcal{F}_{0+} -measurable, and in the case of SLE(8/3) it follows that $\mathbb{P}_{5/8}(V_\infty) = 0$ or 1 by the Blumenthal 0-1 Law, since the corresponding measure $\mathbb{P}_{5/8}$ is a pushforward of Wiener measure through the Loewner equation. For general $\alpha > 5/8$, the same type of Blumenthal 0-1 Law holds via the usual argument. Indeed, the Domain Markov property implies that $\phi_{K_t}(A_t K)$ is a restriction hull that is independent of \mathcal{F}_t , hence for $A \in \mathcal{F}_{0+}$ and $t > 0$ and any bounded, continuous function f on hulls we have

$$\mathbf{E}[f(\phi_{K_t}(A_t K))\mathbf{1}_A] = \mathbf{E}[f(\phi_{K_t}(A_t K))]\mathbb{P}_\alpha(A).$$

Taking a limit of both sides as $t \downarrow 0$ and using the fact that f is continuous and ϕ_{K_t} goes continuously to the identity we get that

$$\mathbf{E}[f(K)\mathbf{1}_A] = \mathbf{E}[f(K)]\mathbb{P}_\alpha(A),$$

which shows that A is independent of all elements of \mathcal{F}_∞ , and therefore of itself. □

We now use Proposition 1 to prove part (4) of Theorem 1. We use the following events to define our thickened sets.

Definition 2 For $z \in \mathbb{H}$ and $\epsilon > 0$, let $I(z, \epsilon)$ be the horizontal line $y = \text{Im } z$ with the gap of width 2ϵ centered around z removed. That is

$$I(z, \epsilon) := \{w \in \mathbb{H} : \text{Im } w = \text{Im } z, |\text{Re}(w - z)| \geq \epsilon\}.$$

Define the sets C_ϵ and D_ϵ by

$$C_\epsilon := \{z \in \mathbb{H} : I(z, \epsilon) \cap K = \emptyset\},$$

$$D_\epsilon := \{L > 0 : I(n\epsilon + iL, \epsilon) \cap K = \emptyset \text{ for some } n \in \mathbb{Z}\}.$$

Lemma 1 *With the definitions above, the following is true \mathbb{P}_α -a.s.:*


$$C = \bigcap_{\epsilon > 0} C_\epsilon, \quad D = \bigcap_{\epsilon > 0} D_\epsilon.$$

Proof Recall that C consists of $z \in \mathbb{H}$ for which $K \cap \{y = \text{Im } z\} = \{z\}$. Hence if $z \in C$ then $z \in C_\epsilon$ for all $\epsilon > 0$. To prove the converse, note that if $z \in C_\epsilon$ for every $\epsilon > 0$ then z is the only possible element in the set $K \cap \{y = \text{Im } z\}$. But the latter set is always non-empty, since restriction hulls are connected and their vertical component goes from zero to infinity (\mathbb{P}_α -a.s.), and therefore with \mathbb{P}_α -probability 1 the set $K \cap \{y = L\}$ is non-empty for all $L > 0$. The proof for D is exactly the same. □

The restriction formula makes it easy to compute the probability that a point $z \in \mathbb{H}$ is in C_ϵ . Indeed, by formula (3) we have

$$\mathbb{P}_\alpha(z \in C_\epsilon) = \mathbb{P}_\alpha(I(z, \epsilon) \cap K = \emptyset) = \phi'_{I(z, \epsilon)}(0)^\alpha,$$

Fig. 2 The dotted point is z and the two horizontal lines on either side form the set $I(z, \epsilon)$. This figure depicts the event that an SLE(8/3) avoids the hull $I(z, \epsilon)$



where $\phi_{I(z, \epsilon)}$ is a conformal map from $\mathbb{H} \setminus I(z, \epsilon)$ onto \mathbb{H} such that $\phi_{I(z, \epsilon)}(w) \sim w$ as $w \rightarrow \infty$. Similarly,

$$\mathbb{P}_\alpha(w, z \in C_\epsilon) = \phi'_{I(w, \epsilon) \cup I(z, \epsilon)}(0)^\alpha.$$

By Proposition 1, the Hausdorff computation for C and D therefore comes down to an estimate of the derivative of these conformal maps at zero. We list three possible methods for these estimates. One deals only with conformal maps and is entirely analytic. The others use probabilistic techniques. We recall the analytic method but do not enter into details.

Analytic method: While it is not possible to write down $\phi_{I(z, \epsilon)}$ explicitly, one can write down the general form of its inverse. Let

$$f_{z, \epsilon}(w) := \lambda w + \frac{\text{Im} z}{\pi} (\log(w - a) - \log(w - b) + \pi i),$$

where the imaginary part of the logarithm is zero along the positive real axis and π on the negative real axis. For appropriate choices of real constants $\lambda, a,$ and b (with $a < b, \lambda > 0$), $f_{z, \epsilon}$ maps \mathbb{H} onto $\mathbb{H} \setminus I(z, \epsilon)$. These constants implicitly depend on z and ϵ , although it is difficult to give closed-form expressions for them. Close analysis of the asymptotic behavior of $\lambda, a,$ and b could be used to get estimates on $\phi'_{I(z, \epsilon)}(0)$ as $\epsilon \downarrow 0$, but we will mostly avoid this strategy. We will, however, mention that a and b are determined mostly by z , while λ is proportional to ϵ^{-2} .

Brownian excursion method: The first probabilistic method uses a well-known formula, due to Bálint Virág [17], for *Brownian excursions* in the upper half plane. Recall that a Brownian excursion in \mathbb{H} can be thought of as a Brownian motion that is started at zero and conditioned to have a positive imaginary part at all later times. Such excursions can be realized by a random path whose horizontal component is a one-dimensional Brownian motion and whose vertical component is an independent Bessel-3 process.

Lemma 2 ([17]) *Let A be a compact hull in the upper half plane such that $\mathbb{H} \setminus A$ is simply connected and $\text{dist}(0, A) > 0$, and ϕ_A be a conformal map from $\mathbb{H} \setminus A$ into \mathbb{H} such that*

$\phi_A(0) = 0$ and $\phi_A(z) \sim z$ as $z \rightarrow \infty$. If BE denotes the path of a Brownian excursion in \mathbb{H} from 0 to ∞ , then

$$\phi'_A(0) = \mathbb{P}(BE \text{ does not intersect } A).$$

In particular, this lemma shows that the filling in of a Brownian excursion has the law of a restriction measure with index 1. It can also be used to get the estimates of Proposition 1, but we prefer the following method that produces asymptotic results (even if they are not necessary in our setting).

Brownian motion method: Instead of using Brownian excursions to compute $\phi'_A(0)$, one can use Brownian motion directly. Oftentimes this is easier as it doesn't require dealing with the conditioning. In an appropriate sense, $\phi'_A(0)$ is the exit density at zero (with respect to Lebesgue measure) of a Brownian motion in $\mathbb{H} \setminus A$, starting from ∞ . This is also called the *excursion Poisson kernel* as seen from ∞ . In what follows we let B be a complex Brownian motion.

Definition 3 Given a simply connected domain D with $z \in D$, $w \in \partial D$, let $H_D(z, w)$ denote the Poisson kernel. In the case $D = \mathbb{H} \setminus A$, we will often be interested in the ‘‘Poisson kernel as seen from infinity’’, for which we introduce the notation

$$H_{\mathbb{H} \setminus A}(\infty, w) := \lim_{L \uparrow \infty} LH_{\mathbb{H} \setminus A}(iL, w).$$

The following estimates will be useful when using Lemma 3 to estimate $\phi'_A(0)$. For $x > 0$, $H_{\mathbb{H}}(z, x) = \frac{1}{\pi} \text{Im}(z)/|z - x|^2$ and consequently $H_{\mathbb{H}}(\infty, x) = \frac{1}{\pi}$. Recall that under a conformal map $f : D \rightarrow D'$, $H_D(z, w)$ changes according to the scaling rule $H_D(z, w) = |f'(w)|H_{f(D)}(f(z), f(w))$. In particular, we have the scaling rule $H_{\mathbb{H} \setminus A}(\infty, w) = H_{\mathbb{H} \setminus rA}(\infty, rw)$.

The next lemma outlines how to use Brownian motion directly to estimate $\phi'_A(0)$. The method of proof is virtually identical to the one for Lemma 2, so we refer the reader to [17] for details.

Lemma 3 For a complex Brownian motion and a compact hull A in the upper half-plane such that $\mathbb{H} \setminus A$ is simply connected and $\text{dist}(0, A) > 0$,

$$\phi'_A(0) = H_{\mathbb{H} \setminus A}(\infty, 0).$$

The computation of $\phi'_{I(z, \epsilon)}(0)$ is thus reduced to some estimates on the exit density of a Brownian motion in the domain $\mathbb{H} \setminus I(z, \epsilon)$. In order to simplify the computations, we first estimate exit densities for an intermediate set S_ϵ .

Lemma 4 Let $S_\epsilon = \mathbb{R} \times [0, 2i] \setminus I(i, \epsilon)$. Then for $x \in \mathbb{R}$ and $\lambda \in [-1, 1]$,

$$H_{S_\epsilon}(\lambda\epsilon + i, x) \sim \frac{\pi\sqrt{1 - \lambda^2}}{8 \cosh^2(\pi x/2)} \epsilon \tag{6}$$

as $\epsilon \downarrow 0$, where ‘‘ \sim ’’ means that the ratio of the two terms converges to 1 uniformly with respect to x and λ . In particular, the probability that the Brownian motion started at i exits S_ϵ on \mathbb{R} is of order ϵ .

Proof Let $z_\epsilon = \lambda\epsilon + i$. In this case, it is easy to find an explicit conformal map from S_ϵ onto \mathbb{H} . A simple one is given by

$$f_\epsilon(z) = \left(\frac{e^{\pi z} + e^{\pi\epsilon}}{e^{\pi z} + e^{-\pi\epsilon}} \right)^{1/2}.$$

By the scaling rule for the Poisson kernel

$$H_{S_\epsilon}(z_\epsilon, x) = |f'_\epsilon(x)| H_{\mathbb{H}}(f_\epsilon(z_\epsilon), f_\epsilon(x)) = \frac{|f'_\epsilon(x)|}{\pi} \frac{\text{Im}(f_\epsilon(z_\epsilon))}{|f_\epsilon(z_\epsilon) - f_\epsilon(x)|^2}.$$

It is straightforward to verify that

$$f_\epsilon(x) \sim 1,$$

as $\epsilon \downarrow 0$, and

$$\begin{aligned} |f'_\epsilon(x)| &= \frac{1}{2f_\epsilon(x)} \frac{2\pi e^{\pi x} \sinh(\pi\epsilon)}{(e^{\pi x} + e^{-\pi\epsilon})^2} \\ &\sim \frac{\pi^2\epsilon}{4 \cosh^2(\pi x/2)}. \end{aligned}$$

Similarly

$$\begin{aligned} f_\epsilon(z_\epsilon) &= \left(\frac{e^{\pi\epsilon} - e^{\pi\lambda\epsilon}}{e^{-\pi\epsilon} - e^{\pi\lambda\epsilon}} \right)^{1/2} \\ &\sim \left(\frac{1 - \lambda}{-1 - \lambda} \right)^{1/2} \\ &= i \left(\frac{1 - \lambda}{1 + \lambda} \right)^{1/2}. \end{aligned}$$

Assembling the pieces proves (6), and then integrating (6) over x proves the last statement. \square

Lemma 5 *Let $x \in \mathbb{R}$ and $\lambda \in [-1, 1]$. Then*

$$H_{\mathbb{H} \setminus I(i, \epsilon)}(\lambda\epsilon + i, x) \sim H_{S_\epsilon}(\lambda\epsilon + i, x)$$

as $\epsilon \downarrow 0$.

Proof If a Brownian motion started at $\lambda\epsilon + i$ exits S_ϵ at x , then it also exits $\mathbb{H} \setminus I(i, \epsilon)$ at x . Consequently, the Poisson kernel on the left hand side is bigger than the one on the right. They are not the same because the Brownian motion in $\mathbb{H} \setminus I(i, \epsilon)$ can hit the line $y = 2i$ before hitting zero, which the Brownian motion in S_ϵ is not allowed to do. Asymptotically this event contributes nothing; indeed there is only an $O(\epsilon)$ chance that the Brownian motion even makes it up to $y = 2i$, and then another $O(\epsilon)$ chance that it passes back through the gap. Overall this makes the event of order ϵ^2 (uniformly in x and λ), which, by Lemma 4, is negligible compared to $H_{S_\epsilon}(\lambda\epsilon + i, x)$. \square

Proposition 2 For $z = y(x + i) \in \mathbb{H}$,

$$\phi'_{I(z,\epsilon)}(0) \sim U(z)\epsilon^2$$

as $\epsilon \downarrow 0$, where

$$U(y(x + i)) = \frac{\pi}{16y^2 \cosh^2(\pi x/2)}.$$

Proof It suffices to prove the result in the case $z = x + i$, for the general form use the scaling rule. We use Brownian motion coming down from infinity as in Lemma 3. In order to reach 0, the Brownian motion coming down from infinity must first pass through the gap of width 2ϵ centered at z , and then from the gap it must transition to zero while avoiding $I(z, \epsilon)$. The two events are independent by the Strong Markov property, and each one is $O(\epsilon)$. More precisely, by Lemmas 4 and 5,

$$\begin{aligned} \phi'_{I(z,\epsilon)}(0) &= H_{\mathbb{H} \setminus I(x+i,\epsilon)}(\infty, 0) \\ &= \int_{[-\epsilon,\epsilon]} H_{\mathbb{H}}(\infty, x + y) H_{\mathbb{H} \setminus I(x+i,\epsilon)}(x + y + i, 0) dy \\ &= \int_{-\epsilon}^{\epsilon} \frac{1}{\pi} H_{\mathbb{H} \setminus I(i,\epsilon)}(y + i, -x) dy \\ &= \frac{\epsilon}{\pi} \int_{-1}^1 H_{\mathbb{H} \setminus I(i,\epsilon)}(\lambda\epsilon + i, -x) d\lambda \\ &\sim \frac{\epsilon^2}{8 \cosh^2(\pi x/2)} \int_{-1}^1 \sqrt{1 - \lambda^2} d\lambda. \end{aligned} \quad \square$$

From Proposition 2 and the restriction formula, it is easy to derive the probability that a bridge point is within distance ϵ of a given point z decays like $\epsilon^{2\alpha}$. From this the first part of Proposition 1 follows easily, but we need a last proposition in order to derive the two point estimate.

Proposition 3 Let $z, w \in \mathbb{H}$, with $\text{Im}(z) > \text{Im}(w)$, and $\epsilon_z, \epsilon_w > 0$. Let $A = I(z, \epsilon_z) \cup I(w, \epsilon_w)$. Then

$$\phi'_A(0) \asymp U(z - w)U(w)\epsilon_z^2\epsilon_w^2,$$

as $\epsilon_z, \epsilon_w \downarrow 0$.

Proof The argument is virtually the same as for the one-point estimate in Proposition 2, the only difference being that the Brownian motion, after passing through the first gap at z then has to pass through a second gap at w . The probability of the latter event can be estimated using Proposition 2; indeed, after temporarily shifting w to zero, there is a $U(z - w)\epsilon_z^2\epsilon_w$ chance that the Brownian motion hits in an ϵ_w neighborhood of w (and therefore also the second gap). With some positive probability it hits in the middle of the second gap, where the probability of moving to zero is, up to a constant, given by $U(w)\epsilon_w$. These two probabilities multiply since, by the Strong Markov property, the path before the second gap is independent of the path after the second gap. □

Remark 1 By carefully decomposing the path according to the points it passes through in the gaps and then integrating, the statement of Proposition 3 could be strengthened to an asymptotic result rather than just up to constants. For our purposes, however, this is not required.

Proof (Proof of Theorem 1, part (4)) Propositions 2 and 3 combine with Proposition 1 to prove the result for C .

For D , the key observation is that if two gaps on a horizontal line do not overlap, then the curve can only avoid the line by going through one of them. Consequently, for $n \neq m$, the events $I(n\epsilon + iL, \epsilon/2) \cap K = \emptyset$ and $I(m\epsilon + iL, \epsilon/2) \cap K = \emptyset$ are disjoint, and therefore

$$\begin{aligned} \mathbb{P}_\alpha(L \in D_\epsilon) &= \mathbb{P}_\alpha\left(\bigcup_{n \in \mathbb{Z}} \{I(n\epsilon + iL, \epsilon/2) \cap K = \emptyset\}\right) \\ &= \sum_{n \in \mathbb{Z}} \mathbb{P}_\alpha(I(n\epsilon + iL, \epsilon/2) \cap K = \emptyset) \\ &\sim \frac{1}{L^{2\alpha}} \sum_{n \in \mathbb{Z}} U\left(\frac{n\epsilon}{L} + i\right)^\alpha \left(\frac{\epsilon^{2\alpha}}{4^\alpha}\right) \\ &\sim \frac{\epsilon^{2\alpha-1}}{4^\alpha L^{2\alpha-1}} \int_{\mathbb{R}} U(x + iL)^\alpha dx \\ &\sim \frac{\pi^\alpha \epsilon^{2\alpha-1}}{32^\alpha L^{2\alpha-1}} \int_{\mathbb{R}} \cosh^{-2\alpha}(\pi x/2) dx. \end{aligned}$$

The transition from sum to integral is a Riemann sum approximation. By $2\alpha > 1$, the integral is a finite constant depending only on α . This gives the one-point estimate for D .

Similarly, for $0 < L < L'$,

$$\begin{aligned} \mathbb{P}_\alpha(L, L' \in D_\epsilon) &= \mathbb{P}_\alpha\left(\bigcup_{m, n \in \mathbb{Z}} \{n\epsilon + iL, m\epsilon + iL' \in C_{\epsilon/2}\}\right) \\ &= \sum_{m, n \in \mathbb{Z}} \mathbb{P}_\alpha(n\epsilon + iL, m\epsilon + iL' \in C_{\epsilon/2}) \\ &\asymp \sum_{m, n \in \mathbb{Z}} \epsilon^{4\alpha} U[(m - n)\epsilon + i(L' - L)]^\alpha U(n\epsilon + iL)^\alpha \\ &\asymp \epsilon^{4\alpha-2} \int_{\mathbb{R}} U(x + i(L' - L))^\alpha dx \int_{\mathbb{R}} U(x + iL)^\alpha dx \\ &\asymp \frac{\epsilon^{4\alpha-2}}{L^{2\alpha-1}(L' - L)^{2\alpha-1}}. \end{aligned}$$

We use the same transition from sum to integral as in the one-point bound. Proposition 1 now completes the proof. □

We show that C and D are almost surely empty for $\alpha \geq 1$. For $\alpha < 1$, the Hausdorff dimension is strictly positive and the set is non empty.

Proof (Proof of Theorem 1, part (5)) For $\alpha = 1$, recall that the imaginary part of a Brownian excursion is a Bessel(3) process, and a bridge height for the hull necessarily corresponds to

a point of increase for the Bessel(3) process. However, it is well known that Bessel(3) has no point of increase since, for example, a Bessel(3) process reversed from its last passage time of a level has the same law as a Brownian motion up to its first hitting time of zero, and Brownian motion is known to have no points of increase (see [15] for details of both facts).

For $\alpha > 1$ consider the rectangle $R = [-1, 1] \times [1/2, 1]$. Cover it with 2^{2n} squares each of side length 2^{-n} , and let $\{S_i\}_{1 \leq i \leq 2^{2n}}$ be the boxes and z_i be their centers. Then, by Proposition 2, the expected number of squares containing a bridge point decays exponentially fast since

$$\begin{aligned} \mathbb{E}_\alpha \left[\sum_{i=1}^{2^{2n}} \mathbf{1}\{C \cap S_i \neq \emptyset\} \right] &\asymp \sum_{i=1}^{2^{2n}} \mathbb{P}_\alpha(I(z_i, 2^{-n}) \cap K \neq \emptyset) \\ &\asymp \sum_{i=1}^{2^{2n}} U(z_i)(2^{-n})^{2\alpha} \\ &= 2^{(2-2\alpha)n} 2^{-2n} \sum_{i=1}^{2^{2n}} U(z_i) \\ &\leq C 2^{(2-2\alpha)n}, \end{aligned}$$

for some constant $C > 0$. The last inequality is a simple consequence of the fact that U is Riemann integrable and hence

$$2^{-2n} \sum_{i=1}^{2^{2n}} U(z_i) \rightarrow \int_R U(z) dA(z) < \infty,$$

where $dA(z)$ is two-dimensional Lebesgue measure. The Borel–Cantelli lemma then proves that R almost surely contains no bridge points. By scale invariance any scaled version of R also contains no bridge points. Translates of R in the horizontal direction also contain no bridge points, since clearly the expected number of bridge points in translates of R decreases as the rectangle is moved away from the imaginary axis. Finally, since the entire half-plane can be covered with countably many scaled and translated versions of R , the entire plane must almost surely be free of bridge points. □

We end this section with the proof of part (2) of Theorem 1. The lack of isolated points in C and D is also a consequence of the renewal property of restriction hulls at bridge points, so we defer the proof of this fact until the end of Sect. 4.

Proof (Proof of Theorem 1, part (2)) We prove the result for D ; the proof for C is similar. To prove that D is closed, suppose that L is a limit point of D . Without loss of generality we may assume that the limiting sequence of bridge heights L_n that converges to L is strictly increasing. If t is the bridge time corresponding to L , then the restriction hull after time t must reside in the domain $\text{Im } z \geq L$ (since each L_n is a bridge height). Then L is not in D if and only if the future hull touches the line $\text{Im } z = L$ but does not cross it, which is clearly an event of probability zero. Indeed, for two points z and w on the same horizontal line let us define $A(z, \epsilon_z, w, \epsilon_w)$ to be the event that the hull goes through the balls $B(z, \epsilon_z)$ and $B(w, \epsilon_w)$ while avoiding $I(z, \epsilon_z) \cap I(w, \epsilon_w)$. The estimates of Proposition 3 can be used to show that the probability of $A(z, \epsilon_z, w, \epsilon_w)$ is of order $\epsilon_z^{2\alpha} \epsilon_w^{2\alpha}$, which easily implies the result since $\alpha > 1/2$. □

4 Renewal at Bridge Lines

In this section we show that the restriction hulls renew themselves at bridge heights. Most of the section is technical, so first we would like to give the intuition behind the renewal property. It is almost entirely a consequence of restriction. Suppose that K is a restriction hull with the law \mathbb{P}_α . Given \mathcal{F}_t , the Domain Markov property (4) says that the future hull has the restriction law corresponding to the domain $(\mathbb{H} \setminus \gamma[0, t], \gamma(t), \infty)$. But if we also know that t is a bridge time, then the future hull is separated from the past by the bridge line that the hull is currently at. The future hull is therefore conditioned not to go below this bridge line, and this conditioning is, by the restriction property, “equivalent” to sampling the future hull from the restriction measure corresponding to the half plane above the bridge line. Shifting the bridge point back to the origin, this means that the shifted future hull $\theta_t K$ also obeys the law \mathbb{P}_α and is independent of \mathcal{F}_t .

There are two main technical obstacles to this intuition. The first is that the event that t is a bridge time for K is not measurable with respect to \mathcal{F}_t , since the set of bridge times is a function of the entire hull. To address this problem and still have a meaningful notion of renewal, we simply expand our filtration to a larger one \mathcal{G}_t that tells us which bridge heights of K_t are also bridge heights of K . The second and more problematic technicality is that t being a bridge time is an event of measure zero, and so conditioning on it requires some care. Theorem 3 deals with this latter problem by showing that the restriction hulls obey a certain Domain Markov property with respect to \mathcal{G}_t , and from this concludes that they refresh themselves at \mathcal{G}_t -stopping times τ such that $\mathbb{P}_\alpha(\tau \in G) = 1$ (recall that G is the set of bridge times).

We make the following definitions:

Definition 4 For $t \geq 0$, let D_t be the set of bridge heights of K_t . Note that D_t is \mathcal{F}_t -measurable and $D_\infty = D$. Observe that $D_t \cap D$ is the set of bridge heights of K_t that are also bridge heights of K , and $D_t \setminus D$ is the set of bridge heights of K_t that are *not* bridge heights of K . We also define

$$L_t := \sup D_t \cap D, \quad L'_t := \inf D_t \setminus D.$$

Note that neither of these quantities, nor $D_t \cap D$ or $D_t \setminus D$, are \mathcal{F}_t -measurable. However, they are measurable with respect to the enlarged filtration

$$\mathcal{G}_t := \sigma(K_s, D_s \cap D; 0 \leq s \leq t).$$

Clearly $\mathcal{F}_t \subset \mathcal{G}_t$, and in this larger filtration the bridge lines (and points, and times) of K that belong to K_t are measurable objects.

Notice that $D_t \cap D$ is almost surely closed, and therefore L_t is actually a maximum rather than a supremum (i.e. $L_t \in D_t \cap D$). Hence L_t is the largest bridge height of K_t that is also a bridge height of K . Clearly $L_t \leq L'_t$. The next result follows easily from these definitions.

Proposition 4 *The σ -algebra \mathcal{G}_t is generated by K_t and L_t , i.e.*

$$\mathcal{G}_t = \sigma(\mathcal{F}_t, L_t).$$

Proof Clearly $\sigma(\mathcal{F}_t, L_t) \subset \mathcal{G}_t$, since L_t is determined by $D_t \setminus D$. For the other direction, it is clear that $D_t \cap D = \{L \in D_t : L \leq L_t\}$. Hence $D_t \cap D$ is determined by both D_t (which

is itself determined by K_t and L_t . This is sufficient because for $s < t$ we have $D_s \cap D \subset D_t \cap D$, and hence $D_s \cap D$ is the intersection of D_s , which is \mathcal{F}_s -measurable, and $D_t \cap D$, which we have just shown is $\sigma(\mathcal{F}_t, L_t)$ -measurable. \square

Proposition 5 *For a fixed $t > 0$, $L_t < L'_t$ with probability one.*

Proof First observe that t is almost surely not a bridge time. It is easy to see that the distance between $\gamma[t, \infty)$ and the last bridge line $\text{Im}(z) = L_t$ is strictly positive (for instance, there must exist another bridge height higher than L_t , and between, it is a continuous compact curve). But a bridge height for $\gamma[0, t)$ that is not a bridge height for the whole curve must be greater than $\inf \text{Im}(\gamma[t, \infty))$. We deduce that L'_t is strictly greater than L_t . \square

Definition 5 Given a subset K of \mathbb{C} , define $J(K) := \inf\{\text{Im } z : z \in K\}$.

With this definition in hand we state the paper’s main technical theorem.

Theorem 3 *Suppose $K = (\gamma, \mathcal{L})$ obeys the law \mathbb{P}_α , and let τ be a \mathcal{G}_t -stopping time. On the event that τ is a bridge time the \mathcal{G}_τ -conditional law of $\theta_\tau K$ is simply the law of a restriction hull in \mathbb{H} . If τ is not a bridge time then the conditional law of $\Lambda_\tau K$, given \mathcal{G}_τ , is the same as the law of a restriction hull K' in $\mathbb{H} \setminus \gamma[0, \tau]$ whose distribution is the restriction measure corresponding to the triple $(\mathbb{H} \setminus \gamma[0, \tau], \gamma(\tau), \infty)$, but further conditioned on the event $L_\tau < J(K') \leq L'_\tau$.*

Remark 2 Note that if τ is a bridge time then $L_\tau = \text{Im } \gamma(\tau)$ and $L_{\tau'} = \infty$. In this situation the notation $L_\tau < J(K') < L_{\tau'}$ can be interpreted as meaning that the future hull lies strictly above the bridge line, which is an event of measure zero. To fully emphasize this very important point we have handled this case with a separate statement at the beginning of the theorem.

Theorem 3 should be seen as the extension of the Domain Markov property (4) to the enlarged filtration \mathcal{G}_t . In words, it simply says that the extra information in \mathcal{G}_τ forces the future restriction hull to go below the horizontal line $y = L'_\tau$ but stay above the horizontal line $y = L_\tau$. This extra conditioning stops L'_τ from being a bridge height for K but preserves L_τ as a bridge height. A detailed proof of the theorem follows. It uses a standard procedure, which we modified from [17], to bootstrap from the easy case of τ being a deterministic time to the general case that τ is a stopping time.

Proof To simplify notation, we will write

$$\mathbb{P}'_\alpha := \mathbb{P}_\alpha^{(\mathbb{H} \setminus \gamma[0, t], \gamma(t), \infty)}(\cdot | L_t < J(K') \leq L'_t)$$

throughout this proof. The goal of the proof is to show that the \mathcal{G}_τ -conditional law of $\Lambda_\tau K$ is \mathbb{P}'_α .

Consider first the case that τ is a deterministic time t . Recall that conditioning on \mathcal{G}_t is the same as conditioning on \mathcal{F}_t and L_t , by Proposition 4. Conditional on \mathcal{F}_t , the Domain Markov property (4) says that $\Lambda_t K$ has the restriction law for the triple $(\mathbb{H} \setminus \gamma[0, t], \gamma(t), \infty)$. Conditioning again on L_t forces the future hull to stay above $y = L_t$ but to go below $y = L'_t$, and since $L_t < L'_t$ with positive probability this conditioning is well-defined. Hence the law conditioned on \mathcal{G}_t is exactly \mathbb{P}'_α .

Another way of stating the above is as follows: let X be a bounded, continuous² function on hulls. Then

$$\mathbf{E}_\alpha[X(\Lambda_t K)|\mathcal{G}_t] = \mathbf{E}'_\alpha[X], \tag{7}$$

where \mathbf{E}_α and \mathbf{E}'_α denote expectations with respect to \mathbb{P}_α and \mathbb{P}'_α , respectively. To finish the proof we need to extend (7) to \mathcal{G}_t -stopping times instead of just fixed times. First suppose that τ only takes values in some countable set \mathcal{T} . Then

$$\begin{aligned} \mathbf{E}_\alpha[X(\Lambda_\tau K)|\mathcal{G}_\tau] &= \sum_{t \in \mathcal{T}} \mathbf{E}_\alpha[X(\Lambda_t K)\mathbf{1}\{\tau = t\}|\mathcal{G}_\tau] \\ &= \sum_{t \in \mathcal{T}} \mathbf{E}_\alpha[X(\Lambda_t K)\mathbf{1}\{\tau = t\}|\mathcal{G}_t] \\ &= \sum_{t \in \mathcal{T}} \mathbf{1}\{\tau = t\} \mathbf{E}_\alpha[X(\Lambda_t K)|\mathcal{G}_t] \\ &= \sum_{t \in \mathcal{T}} \mathbf{1}\{\tau = t\} \mathbf{E}'_\alpha[X] \\ &= \mathbf{E}^\tau_\alpha[X]. \end{aligned}$$

From this we can bootstrap up to the case of general τ . Let τ_n be the smallest element of $2^{-n}\mathbb{N}$ that is greater than or equal to τ . Then the last argument applies to τ_n , so that

$$\mathbf{E}_\alpha[X(\Lambda_{\tau_n} K)|\mathcal{G}_{\tau_n}] = \mathbf{E}^{\tau_n}_\alpha[X]. \tag{8}$$

However, since τ_n is determined at time τ (i.e. τ_n is \mathcal{G}_τ -measurable),

$$\mathbf{E}_\alpha[X(\Lambda_{\tau_n} K)|\mathcal{G}_{\tau_n}] = \mathbf{E}_\alpha[X(\Lambda_{\tau_n} K)|\mathcal{G}_\tau].$$

Since $\Lambda_{\tau_n} K \rightarrow \Lambda_\tau K$ as $n \rightarrow \infty$, and X is bounded and continuous, it follows that the left hand side of (8) converges to

$$\mathbf{E}_\alpha[X(\Lambda_\tau K)|\mathcal{G}_\tau].$$

Hence, if we can show that $\mathbf{E}^{\tau_n}_\alpha[X]$ converges to $\mathbf{E}^\tau_\alpha[X]$ then we are done. Since X is bounded and continuous, this is equivalent to showing that almost surely the law $\mathbb{P}^{\tau_n}_\alpha$ converges weakly to \mathbb{P}^τ_α , which we prove in the next lemma. □

Lemma 6 *Let τ be a \mathcal{G}_t -stopping time and τ_n be the smallest element of $2^{-n}\mathbb{N}$ that is greater than or equal to τ . Then $\mathbb{P}^{\tau_n}_\alpha$ converges weakly to \mathbb{P}^τ_α with probability one, where we define $\mathbb{P}^\tau_\alpha(\cdot) := \mathbb{P}_\alpha(\theta_\tau \cdot)$ in the case that τ is a bridge time.*

Proof Throughout this proof we will let $H_t := (\mathbb{H} + iL_t) \setminus \gamma[0, t]$.

As shown in [8, Lemma 3.2], a probability measure on unbounded hulls in the plane is uniquely determined by the collection of probabilities

$$\mathbb{P}(K \cap A = \emptyset)$$

²The topology we consider is close to the Caratheodory topology and has been defined in [8, Lemma 3.5].

that is indexed by a sufficiently large class of hulls A . Hence it is enough to show that

$$\mathbb{P}_\alpha^{\tau_n}(K' \cap A = \emptyset) \rightarrow \mathbb{P}_\alpha^\tau(K' \cap A = \emptyset) \tag{9}$$

for all hulls A in this class, with probability one. In our case, it is sufficient to prove that for each fixed restriction hull in \mathbb{H} , the convergence (9) holds for all hulls A in H_τ that are a positive distance from $\gamma(\tau)$. Note that since $\tau_n \downarrow \tau$ and γ is continuous, for sufficiently large n one must have that A is at positive distance from $\gamma(\tau_n)$ also. Hence the probabilities on both sides are well defined. We prove (9) in the two distinct cases that τ is and is not a bridge time.

CASE 1: τ IS NOT A BRIDGE TIME

First observe that in the definition of \mathbb{P}_α^t , the conditioning $J(K') > L_t$ forces the hull K' to avoid the region $\{\text{Im } z \leq L_t\}$, and by the restriction property this can equally be achieved by sampling K' from the restriction measure corresponding to the triple $(H_t, \gamma(t), \infty)$. Thus we have the relation

$$\mathbb{P}_\alpha^{(\mathbb{H} \setminus \gamma[0,t], \gamma(t), \infty)}(\cdot | L_t < J(K') \leq L'_t) = \mathbb{P}_\alpha^{(H_t, \gamma(t), \infty)}(\cdot | J(K') \leq L'_t).$$

Let g_t be the conformal map from H_t onto \mathbb{H} such that $g_t(\gamma(t)) = 0$ and $g_t(z) \sim z$ as $z \rightarrow \infty$. Let $R_t := \{z \in H_t : \text{Im } z \leq L'_t\}$. Then

$$\mathbb{P}_\alpha^t(\cdot) = \mathbb{P}_\alpha^{(H_t, \gamma(t), \infty)}(\cdot | K' \cap R_t \neq \emptyset).$$

The first key observation is that for all n sufficiently large we have that $L_{\tau_n} = L_\tau$. This equality is clear since G is closed, and hence τ_n must belong to the same connected component of G^c that τ belongs to, for n sufficiently large. For these n we have $L_{\tau_n} = L_\tau$. For L'_τ there are two distinct possibilities, which we now treat separately.

First note that necessarily $L'_\tau < \infty$. Indeed, the maximum of the imaginary part of $\text{Im } K_\tau$ is always an element of D_τ , and since τ is not a bridge time this maximum cannot be in D . So first consider the case that $L'_\tau < \text{Im } \gamma(\tau)$. By formula (5), we have that

$$\begin{aligned} \mathbb{P}_\alpha^t(K' \cap A = \emptyset) &= \frac{\mathbb{P}_\alpha^{(H_t, \gamma(t), \infty)}(K' \cap A = \emptyset, K' \cap R_t \neq \emptyset)}{\mathbb{P}_\alpha^{(H_t, \gamma(t), \infty)}(K' \cap R_t \neq \emptyset)} \\ &= \frac{\phi'_{A_t}(0)^\alpha - \phi'_{A_t \cup S_t}(0)^\alpha}{1 - \phi'_{S_t}(0)^\alpha}, \end{aligned} \tag{10}$$

where $A_t = g_t(A)$ and $S_t = g_t(R_t)$ (this is justified since neither A nor R_t contains $\gamma(\tau)$). Equation (10) shows that it is sufficient to prove

$$\phi'_{A_{\tau_n}}(0) \rightarrow \phi'_{A_\tau}(0), \quad \phi'_{A_{\tau_n} \cup S_{\tau_n}}(0) \rightarrow \phi'_{A_\tau \cup S_\tau}(0), \quad \phi'_{S_{\tau_n}}(0) \rightarrow \phi'_{S_\tau}(0). \tag{11}$$

For n large enough, $L'_{\tau_n} = L'_\tau$ since for any neighborhood of $\text{Im } \gamma(\tau)$ there is an n sufficiently large such that $D_{\tau_n} \setminus D_\tau$ is contained within this neighborhood. Since $L'_\tau < \text{Im } \gamma(\tau)$, by making the neighborhood sufficiently small we get that $D_{\tau_n} \setminus D$ and $D_\tau \setminus D$ must have the same infimum; that is $L'_{\tau_n} = L'_\tau$. Hence, A_{τ_n} and S_{τ_n} are only decreasing as $\gamma[0, \tau_n]$ decreases, and again since $\gamma[0, \tau_n]$ is a simple curve that shrinks to $\gamma[0, \tau]$ it follows that g_{τ_n} converges uniformly to g_τ on all subcompacts of H_τ , from which the convergences of (11) follow (by Cauchy’s derivative formula and the Schwarz reflection principle, see [8]).

The second possibility is to have $L'_\tau = \text{Im } \gamma(\tau)$. On the one hand, the conditioning on K' going below $\text{Im}(\gamma(\tau))$ is trivial so that $\mathbb{P}^\tau_\alpha = \mathbb{P}^\alpha(H_\tau, \gamma(\tau), \infty)$. On the other hand, L'_{τ_n} is greater than L'_τ so that one can strengthen the conditioning of $\mathbb{P}^{\tau_n}_\alpha$ by requiring that the future hull goes below L'_τ . Since γ is a simple curve shrinking to 0, one again has that g_{τ_n} converges uniformly to g_τ on all subcompacts of H_τ , which proves that the conditioning becomes trivial.

CASE 2: τ IS A BRIDGE TIME

In this case note that A is a hull in the domain $\mathbb{H} + i \text{Im } \gamma(\tau) = \mathbb{H} + i L_\tau$; hence it is simply a translate of a hull in \mathbb{H} . Moreover g_τ is simply the shift map $z \rightarrow z - \gamma(\tau)$, from which it follows that $A_\tau = A - \gamma(\tau)$ and $S_\tau = \mathbb{H}$. Since $\mathbb{P}^\tau_\alpha(\cdot) = \mathbb{P}_\alpha(\theta_\tau \cdot)$, proving (9) amounts to showing that

$$\mathbb{P}^{\tau_n}_\alpha(K' \cap A = \emptyset) \rightarrow \phi'_{A_\tau}(0).$$

We use (10) to rewrite the left hand side. Define $U_t = \phi_{A_t}(S_t \cap A_t^c)$ so that

$$\phi_{A_t \cup S_t} = \phi_{U_t} \circ \phi_{A_t},$$

from which it follows that

$$\phi'_{A_t \cup S_t}(0) = \phi'_{U_t}(0) \phi'_{A_t}(0).$$

Therefore

$$\mathbb{P}^{\tau_n}_\alpha(K' \cap A = \emptyset) = \phi'_{A_{\tau_n}}(0)^\alpha \frac{1 - \phi'_{U_{\tau_n}}(0)^\alpha}{1 - \phi'_{S_{\tau_n}}(0)^\alpha}.$$

The convergence of $\phi'_{A_{\tau_n}}(0)$ to $\phi'_{A_\tau}(0)$ is simple since it only involves the map g_{τ_n} . Note that $L_\tau \leq L_{\tau_n} \leq \text{Im } \gamma(\tau_n)$, so that the domains H_{τ_n} converge to H_τ , and since γ is a simple curve it once again follows that g_{τ_n} converges uniformly to g_τ on all subcompacts of A_τ . As before, this implies the convergence of $\phi'_{A_{\tau_n}}(0)$ to $\phi'_{A_\tau}(0)$.

It remains to be shown that, as $n \rightarrow \infty$,

$$\frac{1 - \phi'_{U_{\tau_n}}(0)^\alpha}{1 - \phi'_{S_{\tau_n}}(0)^\alpha} = \frac{\mathbb{P}_\alpha(K'' \cap U_{\tau_n} \neq \emptyset)}{\mathbb{P}_\alpha(K'' \cap S_{\tau_n} \neq \emptyset)} \rightarrow 1.$$

Observe that

$$\begin{aligned} \mathbb{P}_\alpha(K \cap U_{\tau_n} \neq \emptyset) &= \mathbb{P}_\alpha(K \cap \phi_{A_{\tau_n}}(S_{\tau_n} \cap A_{\tau_n}^c) \neq \emptyset) \\ &= \mathbb{P}_\alpha^{(\mathbb{H} \setminus A_{\tau_n}, 0, \infty)}(K \cap S_{\tau_n} \neq \emptyset) \\ &\sim \mathbb{P}_\alpha^{(\mathbb{H} \setminus A_\tau, 0, \infty)}(K \cap S_{\tau_n} \neq \emptyset). \end{aligned}$$

The last relation follows since g_{τ_n} converges uniformly to g_τ on all subcompacts of H_{τ_n} , to which A eventually belongs, so that A_{τ_n} converges to A_τ . Next recall that $S_{\tau_n} = g_{\tau_n}(R_{\tau_n})$, and

$$0 < \sup \text{Im } R_{\tau_n} \leq L'_{\tau_n} - \text{Im } \gamma(\tau),$$

with the right hand side going to zero as $n \rightarrow \infty$. Since the distance of A_τ from zero is positive, for n sufficiently large the probability that a restriction hull intersects S_{τ_n} is of the

order of $\sup \text{Im } R_{\tau_n}$ and dominated by hulls that intersect S_{τ_n} near zero. Since the set S_{τ_n} is the same near zero in both \mathbb{H} and $\mathbb{H} \setminus A_\tau$, the ratio

$$\frac{\mathbb{P}_\alpha^{(\mathbb{H} \setminus A_\tau, 0, \infty)}(K \cap S_{\tau_n} \neq \emptyset)}{\mathbb{P}_\alpha(K \cap S_{\tau_n} \neq \emptyset)}$$

tends to 1. □

Remark 3 Theorem 3 is most useful when τ is a bridge time, meaning it almost surely takes values in G . In that case $\gamma(\tau)$ is a bridge point for K , and the corresponding bridge line separates the future hull from the past. Shifting the future hull back to the origin by subtracting off $\gamma(\tau)$, we have the following:

Corollary 1 *At \mathcal{G}_t -stopping times τ that almost surely take values in G , the shifted future hull $\theta_\tau K$ obeys the law \mathbb{P}_α .*

Corollary 1 will be the key element in proving that the restriction hulls can be decomposed into a Poisson Point Process, which is the subject of the next section. Before doing that, we immediately apply the corollary to Theorem 1, part (2) by showing that C and D almost surely have no isolated points.

Proof (Proof of Theorem 1, part (2)) We have already shown that C and D are closed, we prove that C has no isolated points. Almost surely, zero is not isolated in C because of the scale invariance and the fact that bridge points exist. For a rational number r , let τ_r be the first bridge time after time r . Then by the previous corollary, we deduce that the law of $\theta_{\tau_r} K$ obeys the law \mathbb{P}_α . Since $\gamma(\tau_r)$ shifts to zero under θ_{τ_r} , the previous remark shows that $\gamma(\tau_r)$ is almost surely not isolated. From these facts we deduce that the event $\{\gamma(\tau_r) \text{ is not isolated in } C \text{ for all rational } r\}$ has probability one. If a point $\gamma(t) \in C$ were isolated then there would have to be an interval of time around t which contains no other bridge times, but since this interval contains a rational time we arrive at a contradiction. □

5 Local Time of the Decomposition

In this section we will show that there exists a natural local time on the bridge heights that we use to decompose the restriction hulls into a Poisson Point Process of irreducible bridges. All the results of this section derive from the theory of subordinators and regenerative sets, which is well described in [3]. We briefly recall the definition of regenerative sets, which is taken from [3, Chapter 2].

Definition 6 A random subset S of $[0, \infty)$ is a *regenerative set* with respect to a filtration \mathcal{F}_t if for every $s \geq 0$, conditionally on $M_s = \inf\{t > s : t \in S\} < \infty$, the shifted set $(S - M_s) \cap [0, \infty)$ has the same law as S and is independent of \mathcal{F}_{M_s} .

Using the results of Sects. 3 and 4, we can immediately prove:

Proposition 6 *The set D of bridge heights is regenerative with respect to $\mathcal{D}_L := \sigma(D \cap [0, L])$.*

Proof Consider $L \geq 0$. Since D is closed, $M_L \in D$ almost surely. Then M_L is a bridge height, and the time τ_L at which the curve reaches this bridge height is a \mathcal{G}_t -stopping time taking values in G . By Corollary 1, the \mathcal{G}_{τ_L} -law of $\theta_{\tau_L} K$ is the same as the original law of K . Consequently, the \mathcal{G}_{τ_L} -law of $D(\theta_{\tau_L} K) = D - M_L$ is the same as the law of D . Since $\mathcal{D}_L \subset \mathcal{G}_{\tau_L}$ this completes the proof. \square

Proposition 6 proved that the set D is regenerative, and consequently by [3, Theorem 2.1] it is the closure of the image of some subordinator (and the subordinator is unique up to a linear change of its time scale). On the other hand, Theorem 1 showed that D is scale invariant, and it is an easy step to deduce from this that the subordinator must be stable. Recall that there is a one-parameter family of stable subordinators, indexed by the real numbers between 0 and 1, and, as shown in [3, Chapter 5], the index of a stable subordinator is the same as the Hausdorff dimension of its image. Hence we have the following:

Corollary 2 *Under the law \mathbb{P}_α , the set D is the closure of the image of a stable subordinator $(\sigma_\lambda, \lambda \geq 0)$ of index $2 - 2\alpha$.*

The parameter λ can be thought of as the local time corresponding to the subordinator. Recall that the local time for σ is the function $\lambda : [0, \infty) \rightarrow [0, \infty)$ defined by $\lambda(s) := \inf\{t \geq 0 : \sigma_t > s\}$, and it is well known in the subordinator literature that λ is an increasing, continuous function which increases only on D . This means that if we run the restriction hulls on the λ time scale, then the hull grows only when it is crossing bridge lines. For $\lambda \geq 0$ we define

$$\tau_h := \inf\{t \geq 0 : \sup \text{Im}(K_t) = h\},$$

and

$$t(\lambda) := \tau_{\sigma_\lambda}.$$

Note that σ_λ is the bridge height at which λ units of local time are first accumulated, and then $t(\lambda)$ is the time, in the original parameterization of the restriction hull, at which the local time first reaches λ . It follows that $t(\lambda)$ is an increasing, right-continuous process for which the closure of its image is precisely the set of bridge times G . Intervals of λ on which the process $t(\lambda)$ is flat correspond to times at which the restriction hull is between bridge heights. Using the $t(\lambda)$ time-scale, we are able to define a Poisson Point Process taking values in the space of irreducible bridges rooted at the origin. Let δ be the curve which starts and ends at zero in zero time (i.e. $\delta : \{0\} \rightarrow \{0\}$). For $\lambda \geq 0$, define e_λ by

$$e_\lambda = \begin{cases} \theta_{t(\lambda-), t(\lambda)} K, & t(\lambda) > t(\lambda-) \\ \delta, & t(\lambda) = t(\lambda-). \end{cases} \tag{12}$$

From this we have the following:

Proposition 7 *e_λ is an $(\mathcal{F}_{t(\lambda)})_{\lambda \geq 0}$ Poisson Point Process on the space of irreducible bridges.*

Proof Take a subset U of the set of irreducible bridges that doesn't contain δ , and an interval $I := [\lambda_1, \lambda_2]$. As in [15, Chapter XII], one needs to show that the number of times that e_λ belongs to U for $\lambda \in I$ is independent of $\mathcal{F}_{t(\lambda_1)}$ and has the same law as the number of times that e_λ belongs to U for $\lambda \in [0, \lambda_2 - \lambda_1]$. But this is essentially a property of Corollary 1. \square

We denote by ν_α the intensity measure of the Poisson Point Process e_λ , and we call it the **continuum irreducible bridge measure**. It conveniently encodes all the behavior of continuum irreducible bridges. For a set of irreducible bridges E , $\nu_\alpha(E)$ is simply the expected number of elements of E that occur in $e[0, 1]$, which may or may not be finite. For instance, if E_L is the set of irreducible bridges with height greater than L , then a simple consequence of Corollary 2 is that $\nu_\alpha(E_L) = c_\alpha L^{2\alpha-2}$ for some fixed constant c_α , and furthermore,

$$\mathbf{P}_\alpha^L(\cdot) := \frac{\nu_\alpha(\cdot \cap E_L)}{\nu_\alpha(E_L)} \tag{13}$$

is exactly the law of the first irreducible bridge with height greater than L . To make the analogy with other well-known decompositions of stochastic processes, ν_α is the equivalent of Itô’s measure on 1-dimensional Brownian excursions, or Bálint Virág’s measure on 2-dimensional Brownian Beads. Compared to half-plane SAWs, ν_α is the analogue of the measure $\mathbf{P}(\omega) = \beta^{-|\omega|}$ on SAW irreducible bridges, although we point out that \mathbf{P} is a probability measure (by Kesten’s relation), whereas ν_α is infinite but σ -finite.

In the case of half-plane SAWs, the measure on paths is realized by concatenating together an i.i.d. sequence of irreducible bridges, each distributed according to \mathbf{P} , and in the continuum a similar statement holds. If $(e_\lambda)_{\lambda \geq 0}$ is a Poisson Point Process of irreducible bridges with intensity measure ν_α , then the concatenation

$$K = \bigoplus_{\lambda \geq 0} e_\lambda$$

has the law of an index α restriction hull. Note, however, that we are not attempting to show that the irreducible bridges can be concatenated together in such a way as to reconstruct the sequence of growing hulls $(K_t)_{t \geq 0}$, even though this should be possible with enough care. Recall though that the time parameterization we are using for the restriction hulls is completely artificial to begin with, and therefore attempting to reconstruct it would mostly be an uninteresting and unuseful exercise.

6 Open Questions

In this final section we present some open questions that were raised by our work.

Question 1 What other properties of the irreducible bridge measure ν_α can be derived?

Our work has essentially determined only one main property of bridges: that the distribution of their vertical height is the same as the jump distribution for a stable subordinator of index $2 - 2\alpha$ (up to a multiplicative constant). Ultimately we hope that much more can be said about irreducible bridges than this. It may be naturally difficult to say anything more, since even in the case of half-plane SAWs there is not much known about irreducible bridges (although in the “off-critical” case there are some results, see [14, Chapter 4]). For other two-dimensional decompositions, notably Virág’s Brownian Beads, it appears similarly difficult to say anything about the bead measure.

Question 2 Is there a constructive way of building irreducible bridges?

In the case of SLE(8/3), for example, is there a driving term for the Loewner equation that outputs irreducible bridges (perhaps with at least some specified vertical height)? And for general restriction measures with $\alpha < 1$, can some driving term for the Loewner equation be combined with the Brownian loop soup to produce irreducible bridges for restriction hulls?

Question 3 Is there a natural “length” that can be put on irreducible bridges?

For half-plane SAWs the length of the walk is simply the number of steps in it, and many results on SAWs are expressed in terms of this length. We expect that there is some way of defining a similar natural length on irreducible bridges, and that this length is somehow the scaling limit of the length for SAWs. However, because the irreducible bridges are fractal objects it is not an easy matter to define a non-trivial length on them. In the case of SLE(8/3) specifically, this question is closely related to the problem of the “natural time parameterization” for SLE, which has recently been considered by Lawler and Sheffield [12]. The key idea of their time parameterization is to build a length measure on the curve (that also has some other desirable properties), and then reparameterize in such a way that the length of the curve at time t is t , as with the SAWs. Their length measure should also be a natural length measure for irreducible bridges.

Question 4 Is there some sort of continuous analogue of Kesten’s relation?

This is closely related to the problem of the natural length on irreducible bridges described above. Supposing that $L(K)$ is the “natural length” of an irreducible bridge, and making an analogy with (2), we might expect that

$$\int_0^\infty \beta^{-l} \nu_\alpha(L(K) \in dl)$$

is finite for $\beta < \mu$ but infinite for $\beta > \mu$, for some universal μ , and then one can ask for the behavior at this critical μ .

Question 5 Can the restriction hulls be time parameterized in such a way that the time parameterization also refreshes itself at bridge points?

Presently we are only showing that the hulls refresh themselves as sets and *not* as time parameterized objects. But it is entirely plausible that there is some time parameterization which refreshes itself at bridge points along with the geometrical objects, especially considering that the counting parameterization for half-plane SAWs has this property (at each bridge point, one simply starts counting off the number of steps anew). It is possible that the natural time parameterization of Lawler and Sheffield will have this property for SLE(8/3) but it is not immediately clear that this will be the case, since their time parameterization has no way of seeing that it is currently at a bridge point and therefore is unlikely to refresh at such bridge times.

Question 6 Can some element of the bridge decomposition be used to prove the existence of, or at least heuristically deduce, critical exponents for half-plane SAWs or SAW bridges?

For example, it is conjectured that the number of N -step SAW bridges grows asymptotically like $N^{-\beta} u^N$ as $N \rightarrow \infty$, for the same μ as in (1) and some unknown constant

β . Recently, Neal Madras has privately communicated to us his conjecture that $\beta = 7/16$, although this quantity was likely known beforehand in the physics literature. He uses two different methods to derive this value, the first being based purely on some heuristics for half-plane SAWs, and the other making use of the relation (13) and the conjecture that the scaling limit of half-plane SAWs is SLE(8/3). Being able to answer further questions of this type would be extremely helpful for studying half-plane SAWs.

Question 7 Do bridge heights and lines exist for SLE(κ) for values of κ different from 8/3. If so, what is the Hausdorff dimension of C and D and how does it depend on κ ?

Currently we only know that at $\kappa = 0$ and $\kappa = 8/3$, the Hausdorff dimensions of C and D are 1 and 3/4, respectively (the $\kappa = 0$ result is clear from the fact that the corresponding SLE curve is a vertical line). We conjecture that the Hausdorff dimensions of C and D are always the same, and they are a strictly decreasing, continuous function of κ . When $\kappa = 4$ the Hausdorff dimension must certainly be zero since the SLE(4) curve comes arbitrarily close to the real line, but we do not know if this is the smallest κ for which the dimension is zero. We have no conjecture as to what that κ might be, other than it is somewhere between 8/3 and 4.

We should briefly mention that, as a corollary of Theorem 1, we do have lower bounds on the Hausdorff dimension of C and D for $2 \leq \kappa \leq 8/3$. Since attaching loops to an SLE curve can only reduce the number of bridge points that the SLE curve has, we know

Proposition 8 *Let C and D be the set of bridge points and heights for an SLE(κ) curve, with $2 \leq \kappa \leq 8/3$. Then the Hausdorff dimensions of C and D are both almost surely constant, with $\dim_{\text{H}} C \geq 3 - \frac{6}{\kappa}$.*

This lower bound is probably far from sharp, since it is increasing with κ rather than decreasing. To prove that the Hausdorff dimensions of C and D are almost surely constant, Theorem 1 part (3) can be used without modification.

Acknowledgements We are grateful to Wendelin Werner for initially suggesting this problem to us, for many helpful and encouraging discussions along the way, and for hosting the first author at the École Normale Supérieure where most of this work was completed. We also thank Vladas Sidoravicius for hosting us at IMPA, where this work was begun, and Bálint Virág for some enlightening conversations. Finally, we thank an anonymous referee for some very helpful suggestions which greatly improved the presentation of this work.

References

1. Alberts, T., Sheffield, S.: Hausdorff dimension of the sle curve intersected with the real line. *Electron. J. Probab.* **13**, 1166–1188 (2008). (electronic) URL <http://www.math.washington.edu/~ejpecp/>
2. Beffara, V.: The dimension of the SLE curves. *Ann. Probab.* **36**(4), 1421–1452 (2008)
3. Bertoin, J.: Subordinators: examples and applications. In: *Lectures on probability theory and statistics* (Saint-Flour, 1997). *Lecture Notes in Math.*, vol. 1717, pp. 1–91. Springer, Berlin (1999)
4. Dubédat, J.: Excursion decompositions for SLE and Watts' crossing formula. *Probab. Theory Related Fields* **134**(3), 453–488 (2006)
5. Kennedy, T.: A fast algorithm for simulating the chordal Schramm-Loewner evolution. *J. Stat. Phys.* **128**(5), 1125–1137 (2007)
6. Kesten, H.: On the number of self-avoiding walks. *J. Math. Phys.* **4**, 960–969 (1963)
7. Kesten, H.: On the number of self-avoiding walks. II. *J. Math. Phys.* **5**, 1128–1137 (1964)
8. Lawler, G., Schramm, O., Werner, W.: Conformal restriction: the chordal case. *J. Am. Math. Soc.* **16**(4), 917–955 (2003) (electronic)

9. Lawler, G.F.: Hausdorff dimension of cut points for Brownian motion. *Electron. J. Probab.* **1**(2), (1996) approx. 20 pp. (electronic)
10. Lawler, G.F.: Conformally invariant processes in the plane. *Mathematical Surveys and Monographs*, vol. 114. American Mathematical Society, Providence (2005)
11. Lawler, G.F., Schramm, O., Werner, W.: On the scaling limit of planar self-avoiding walk. In: *Fractal Geometry and Applications: a Jubilee of Benoît Mandelbrot, Part 2. Proc. Sympos. Pure Math.*, vol. 72, pp. 339–364. Amer. Math. Soc., Providence (2004)
12. Lawler, G.F., Sheffield, S.: Construction of the natural parameterization for SLE curves. (2009). [arXiv:0906.3804v1](https://arxiv.org/abs/0906.3804v1) [math.PR]
13. Lawler, G.F., Werner, W.: The Brownian loop soup. *Probab. Theory Related Fields* **128**(4), 565–588 (2004)
14. Madras, N., Slade, G.: *The Self-Avoiding Walk*. Probability and Its Applications. Birkhäuser, Boston (1993)
15. Revuz, D., Yor, M.: *Continuous Martingales and Brownian Motion*. *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*, vol. 293, 3rd edn. Springer, Berlin (1999)
16. Schramm, O., Zhou, W.: Boundary proximity of SLE (2007). [arXiv:0711.3350v2](https://arxiv.org/abs/0711.3350v2) [math.PR]
17. Virág, B.: Brownian beads. *Probab. Theory Related Fields* **127**(3), 367–387 (2003)