Journal of Statistical Physics, Vol. 114, Nos. 1/2, January 2004 (© 2004)

Conformal Invariance and Stochastic Loewner Evolution Predictions for the 2D Self-Avoiding Walk—Monte Carlo Tests

Tom Kennedy[1]

Received May 2, 2003; accepted June 24, 2003

Simulations of the two-dimensional self-avoiding walk (SAW) are performed in a half-plane and a cut-plane (the complex plane with the positive real axis removed) using the pivot algorithm. We test the conjecture of Lawler, Schramm, and Werner that the scaling limit of the two-dimensional SAW is given by Schramm's stochastic Loewner evolution (SLE). The agreement is found to be excellent. The simulations also test the conformal invariance of the SAW since conformal invariance implies that if we map infinite length walks in the cut-plane into the half plane using the conformal map $z \to \sqrt{z}$, then the resulting walks will have the same distribution as the SAW in the half plane. The simulations show excellent agreement between the distributions.

KEY WORDS: Self-avoiding walk; pivot algorithm; SLE.

1. INTRODUCTION

Lawler, Schramm, and Werner[7] have conjectured that the scaling limit of the two-dimensional self-avoiding walk (SAW) is given by Schramm's[11] stochastic Loewner evolution (SLE). SLE is a two dimensional conformally invariant random process which depends on a parameter κ, and so is denoted SLE_κ. Chordal SLE refers to the version of SLE in which the random curve or set connects two points on the boundary of a simply connected domain. It is usually defined first for the case where the domain is the half-plane and the two boundary points are 0 and ∞. Its definition is then extended to other simply connected domains D and boundary points

[1] Department of Mathematics, University of Arizona, Tucson, Arizona 85721; e-mail: tgk@math.arizona.edu

0022-4715/04/0100-0051/0

using a conformal map from the half-plane to D which maps the two boundary points appropriately. If $\kappa < 4$, chordal SLE gives a probability measure on simple curves, i.e., curves that do not intersect themselves.[10] The conjecture of Lawler, Schramm, and Werner is that for any simply connected domain D and points z and w on its boundary, chordal $\mathrm{SLE}_{8/3}$ is the scaling limit of SAW's that go from z to w and stay inside D.

For $\kappa = 8/3$, Lawler, Schramm, and Werner[8] have a theorem that makes it possible to explicitly compute the distributions of many random variables associated with the SLE random curve. For the scaling limit of the SAW, these random variables can be studied by simulation. Thus one can numerically test their conjecture that the scaling limit of the SAW is $\mathrm{SLE}_{8/3}$ by comparing the distributions from simulations of the SAW with the exact distributions for $\mathrm{SLE}_{8/3}$. This test was carried out for two such random variables for the SAW in the upper half-plane in ref. 2, and excellent agreement was found. In this paper we consider more random variables for which the exact distribution can be computed for $\mathrm{SLE}_{8/3}$. We compare their exact distributions with the numerical distributions of the same random variables for the SAW in the half-plane. We also simulate the SAW in the cut-plane consisting of the complex plane minus the non-negative real axis. The map $z \to \sqrt{z}$ takes the cut-plane onto the half-plane, and by composing the random variables for the half-plane with this map we obtain corresponding random variables for the cut-plane. We compare their distributions for the SAW from our simulations for the cut-plane with the exact distributions for $\mathrm{SLE}_{8/3}$. We also consider the probability that the walk passes to the right of a given point in the half-plane (or the cut-plane) and compare this probability for the SAW simulations with an exact formula of Schramm[12] for SLE. In all our simulations we use a square lattice. Other lattices, e.g., triangular or hexagonal, should have the same scaling limit, but we have not done any simulations to test this.

We consider a variety of random variables which have the advantage that their distribution for $\mathrm{SLE}_{8/3}$ may be explicitly computed. We test the conjecture by comparing these explicit distributions with the distributions obtained by simulating the self-avoiding walk. For both of the domains we consider, the terminal point of the walk is at infinity. In such cases it is expected that we can construct the scaling limit by considering all SAW walks with a fixed length N which start at the origin, taking the limit $N \to \infty$ and then taking the limit that the lattice spacing goes to zero. We simulate walks with a fixed number of steps with the pivot algorithm, a Markov chain Monte Carlo algorithm. In the appendix we show that the Markov chain in this algorithm is irreducible for the half and cut planes. It trivially satisfies detailed balance with respect to the uniform probability measure, so by a standard theorem of Markov chains the distributions for the random variables that one obtains by running the simulation will converge to their distributions under the desired uniform measure on self-avoiding walks. The simulations are of course only run for a finite amount of time. We will refer to the error resulting from only running the simulation a finite amount of time as "statistical error." These statistical errors can be estimated and are shown in our plots with error bars that give two standard deviations. Besides these statistical errors there are two other sources of error. We must take the length of the walk to infinity and then we must take the lattice spacing to zero. We will pay particular attention to these two sources of error when we discuss the results of the simulation. We will see that the differences we find between the self-avoiding walk simulations and the $SLE_{8/3}$ explicit distributions are small and consistent with these three sources of error. We conclude that our simulations support the conjecture that the scaling limit of the self-avoiding walk is $SLE_{8/3}$.

For a domain D and two finite points z and w on its boundary, the scaling limit should be constructed as follows. We introduce a lattice and consider all self-avoiding walks which start at z and end at w. (The number of steps in the walks is not fixed.) The probability of such a walk is taken to be proportional to β^{-N} where N is the number of steps in the walk, and β is the constant such that the number of SAW's in the plane starting at the origin grows with the number of steps, N, as β^N. (The value of the connectivity constant β depends on the particular lattice. Our simulations are only for cases with the terminal point at infinity, so the value of β does not enter into our simulations.) The measure is normalized so that it is a probability measure. We then take the limit of this measure as the lattice spacing goes to zero. The construction of the scaling limit in the case of SAW's with infinite terminal point is rather different from the case of a finite terminal point, so it would be interesting to test the conjecture that the scaling limit is given by $SLE_{8/3}$ in the case of a finite terminal point.

In addition to describing the scaling limit of the SAW, SLE is conjectured to describe the scaling limit of a large number of other two dimensional models. Many of these conjectures have been proved recently. Schramm showed that if the loop-erased random walk has a conformally invariant scaling limit, then that limit must be SLE_2.[11] He also conjectured that the scaling limit of percolation should be related to SLE_6, and the scaling limit of uniform spanning trees (UST) is described by SLE_2 and SLE_8. The conjectures for the loop-erased random walk and the UST have been proved by Lawler, Schramm, and Werner[6] Smirnov has proved the conformal invariance conjecture for critical percolation on the triangular lattice and that SLE_6 describes the limit.[13] Lawler, Schramm, and Werner used SLE_6 to rigorously determine the intersection exponents for Brownian motion and proved a conjecture of Mandelbrot that the outer boundary of a Brownian path has Hausdorff dimension $4/3$.[3–5] The random cluster representation of the Potts model for $0 < q < 4$ was conjectured by Rohde and Schramm to be related to the SLE process as well.[10]

In Section 2 we define the random variables that we will use to test the conjecture, and work out their distributions for $SLE_{8/3}$ using a theorem of Lawler, Schramm, and Werner. In Section 3 we describe the results of the simulations. Some details about how the simulations were done are given in Section 4. Appendix A gives a proof that the pivot algorithm is irreducible in the half and cut planes.

2. SLE PREDICTIONS

The random variables we consider are defined for curves in the upper half-plane as follows. Note that these random variables are defined both for the SAW and for SLE. We use γ to denote the random curve in both cases. Consider a horizontal line at a height of c above the horizontal axis. The curve γ will intersect it, possibly more than once, and we look for the left-most intersection. The random variable X_e is the x-coordinate of this intersection, divided by c. So

$$X_e = \frac{1}{c} \min\{x\colon x + ic \in \gamma\} \tag{1}$$

We can also consider the first intersection of the curve with the horizontal line. ("First" means the first intersection as we traverse the curve starting at the origin.) We let X_f be the x-coordinate of this intersection, divided by c. (The subscripts e and f are for "extreme" and "first," respectively.) The next random variable is defined using a vertical line at a distance c to the right of the origin. The curve will intersect it, and we look for the lowest intersection. The random variable Y_e is the y-coordinate of this intersection, divided by c. So

$$Y_e = \frac{1}{c} \min\{y\colon c + iy \in \gamma\} \tag{2}$$

The random variable Y_f is the y-coordinate of the first intersection, divided by c. For the final random variable, consider a semi-circle of radius c centered at the the point cd on the real axis where $|d| < 1$. So the origin where the random curve starts is inside the semicircle. The intersections of the random curve with the semicircle are of the form $c(d + e^{i\theta})$ and we look for the intersection with the smallest θ. The random variable Θ_e is this smallest angle, normalized so that it ranges between 0 and 1. So

$$\Theta_e = \frac{1}{\pi} \min\{\theta: c(d+e^{i\theta}) \in \gamma\} \tag{3}$$

The random variable Θ_f is the angle of the first intersection, again normalized so that it ranges from 0 to 1. If the probability measure is invariant under dilations, then the distributions of all of these random variables are independent of c. This is true for SLE and is expected to be true for the scaling limit of the SAW.

The distributions of X_e, Y_e, Θ_e are all easily computed using the following theorem of Lawler, Schramm, and Werner. Let $\mathbb{H}$ be the upper half-plane. Let A be a compact subset of the closure of $\mathbb{H}$ such that $\mathbb{H} \setminus A$ is simply connected and 0 is not in A. Let Φ_A be the conformal map from $\mathbb{H} \setminus A$ onto $\mathbb{H}$ which fixes 0 and ∞ and has $\Phi'_A(\infty) = 1$. We continue to denote the random curve generated by SLE, the SLE "trace," by γ.

Theorem 1 (Lawler, Schramm, and Werner[8]). For $\kappa = 8/3$, chordal SLE in the half plane has

$$P(\gamma \cap A = \varnothing) = \Phi'_A(0)^{5/8} \tag{4}$$

Our next step is to use this theorem to compute the distributions of X_e, Y_e, and Θ_e.

2.1. Hitting the Horizontal Line

It is convenient to take $c = \pi$ to compute the distribution of X_e. Let L_t be the horizontal ray which starts at $t + \pi i$ and goes to the left. Let $\Phi_{L_t}(z)$ be the conformal map which maps $\mathbb{H} \setminus L_t$ onto $\mathbb{H}$ and satisfies the conditions in the theorem. Note that $X_e \leqslant t/\pi$ if and only if γ hits L_t. So by the theorem

$$P(X_e \leqslant t/\pi) = 1 - \Phi'_{L_t}(0) \tag{5}$$

The map $w(z) = z + \ln(z) + 1 + t$ maps $\mathbb{H}$ onto $\mathbb{H} \setminus L_t$. We need the inverse of this map but it cannot be explicitly found. The inverse should be normalized so that it fixes 0 and ∞ and has derivative 1 at ∞. The above map does not fix 0, but meets the other two conditions. Fixing 0 is not necessary since we can achieve this condition by just adding a constant to the inverse map, and this which will not change its derivative. So we have

$$\Phi'_{L_t}(0) = \frac{dz}{dw}(0) = \left(\frac{dw}{dz}(z_0)\right)^{-1} \tag{6}$$

where z_0 is the image of 0 under the inverse map, i.e., $0 = z_0 + \ln(z_0) + 1 + t$.

Define

$$g(x) = x + \ln(x) \tag{7}$$

This is an increasing function which maps $(0, \infty)$ onto the real line, so it has an inverse that maps the real line to $(0, \infty)$. Note that $z_0 = g^{-1}(-t-1)$. We have

$$\frac{dw}{dz}(z_0) = 1 + \frac{1}{g^{-1}(-t-1)} \tag{8}$$

So (5) and a trivial change of variables gives

$$P(X_e \leqslant t) = 1 - \left(\frac{g^{-1}(-\pi t - 1)}{g^{-1}(-\pi t - 1) + 1}\right)^{5/8} \tag{9}$$

Although g^{-1} cannot be explicitly computed, it can be trivially computed numerically. The graph of the above distribution is the solid line in Fig. 1. The open circles in the figure are the results of the simulation for the SAW.

We can find the asymptotic behavior of the distribution in (9) as t goes to $\pm\infty$. For large positive t, $g(t) = t + \ln(t) \approx t$. So as $t \to -\infty$, $g^{-1}(-\pi t - 1) \approx -\pi t$, and so

$$P(X_e \leqslant t) \approx 1 - \left(\frac{-\pi t}{-\pi t + 1}\right)^{5/8} \approx -\frac{5}{8\pi t}, \qquad \text{as} \quad t \to -\infty \tag{10}$$

As $t \to 0$, $g(t) \approx \ln(t)$. So as $t \to \infty$, $g^{-1}(-\pi t - 1) \approx e^{-\pi t - 1}$. So

$$P(X_e \leqslant t) \approx 1 - \left(\frac{e^{-\pi t - 1}}{e^{-\pi t - 1} + 1}\right)^{5/8} \approx 1 - e^{-5(\pi t + 1)/8}, \qquad \text{as} \quad t \to \infty \tag{11}$$

As $t \to -\infty$, the probability goes to zero slowly, but as $t \to \infty$, the probability goes to one exponentially fast. This is reasonable since when X_e is

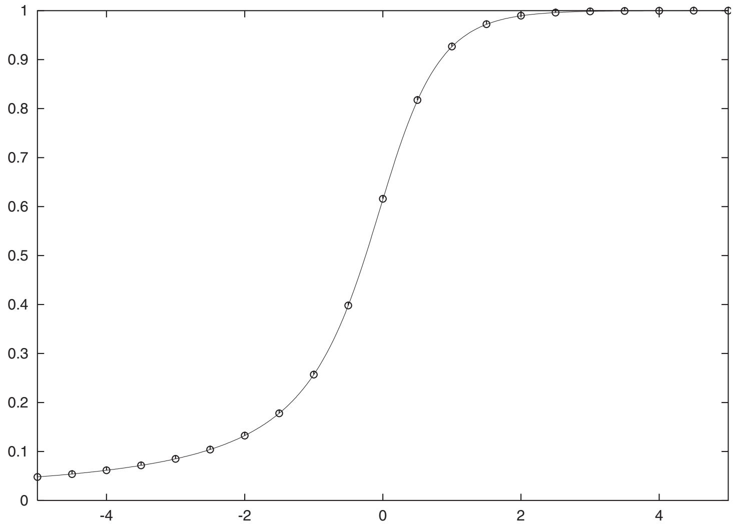

Fig. 1. The distribution, $P(X_e \leqslant t)$, of X_e for the half-plane. The solid line is the distribution for $\mathrm{SLE}_{8/3}$, and the open circles are the results of the simulation of the SAW.

very negative it only means there is at least one intersection with the horizontal line far to the left of the origin, but when X_e is very positive it means that all intersections with the horizontal line are far to the right of the origin.

2.2. Hitting the Vertical Line

The distribution of Y_e was studied in ref. 2. We take A_t to be the line segment from 1 to $1+it$. The conformal map that maps $\mathbb{H} \setminus A_t$ onto $\mathbb{H}$ with the required normalizations is

$$\Phi_{A_t}(z) = i\sqrt{-(z-1)^2 - t^2} \tag{12}$$

where the square root has its branch cut along the negative real axis. Thus the distribution of Y_e is

$$P(Y_e \leqslant t) = P(\gamma[0,\infty) \cap A_t \neq \varnothing) = 1 - \Phi'_{A_t}(0)^{5/8} = 1 - (1+t^2)^{-5/16} \tag{13}$$

Figure 2 shows this distribution and the results of the simulation for the SAW.

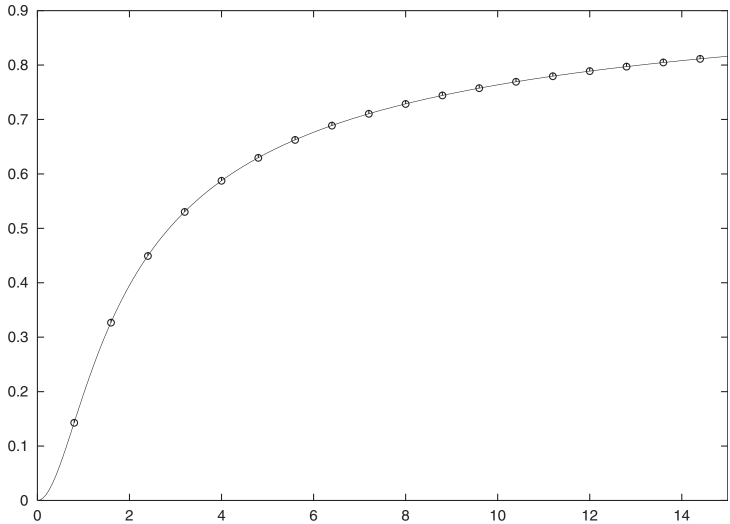

Fig. 2. The distribution of Y_e for the half-plane. The solid line is $\mathrm{SLE}_{8/3}$, and the open circles are the SAW.

2.3. Hitting the Circle

It is convenient to translate so the semicircle is centered at the origin. Setting $c=1$, this means the random curves start at $-d$. So $P(\Theta_e \leqslant t) = 1-\Phi'(-d)^{5/8}$, where Φ is a conformal map which takes the half-plane minus the arc $A_\phi = \{e^{i\theta}\colon 0 \leqslant \theta \leqslant \pi t\}$ onto the half-plane with the normalizations that the map fixes ∞ and has derivative 1 at ∞. (As in the previous case, we ignore the condition that the map fixes the origin since it does not affect the derivative.)

The conformal map

$$z \to \frac{z-1}{z+1} \tag{14}$$

sends the upper half-plane (including ∞) onto itself, and it sends the upper half of the unit circle to the upper half of the imaginary axis. Let

$$a = \frac{\sin(\pi t)}{1+\cos(\pi t)} \tag{15}$$

Then the arc A_t is mapped onto the line segment from 0 to ia. We can then map $\mathbb{H}$ with this line segment removed onto $\mathbb{H}$ as we did in the previous section. Composing these two maps and multiplying by a factor of $\sqrt{1+a^2}$ for later convenience, we define

$$\psi(z) = i\sqrt{1+a^2}\left[-\frac{(z-1)^2}{(z+1)^2} - a^2\right]^{1/2} \tag{16}$$

with the branch cut for the square root being the negative real axis. The map ψ sends $\mathbb{H} \setminus A_t$ onto $\mathbb{H}$.

The map ψ does not send ∞ to itself. For z near ∞,

$$\psi(z) = (1+a^2)\left(1 - \frac{2}{(1+a^2)\,z} + \cdots\right) \tag{17}$$

In particular, $\psi(\infty) = (1+a^2)$. Now let

$$\Phi_{A_t}(z) = \frac{2}{1+a^2-\psi(z)} \tag{18}$$

For large z, $\Phi_{A_t}(z) \approx z$, so the derivative at ∞ is 1 as required.

For real x with $-1 < x < 1$, the choice of branch cut leads to

$$\psi(x) = -\sqrt{1+a^2}\left[\frac{(x-1)^2}{(x+1)^2} + a^2\right] \tag{19}$$

Define

$$s = (1+a^2)^{-1} = \frac{1+\cos(\pi t)}{2} \tag{20}$$

This will prove to be a natural variable to use. We have

$$\psi(x) = -\frac{1}{\sqrt{s}}\left[\frac{(x-1)^2}{(x+1)^2} + \frac{1}{s} - 1\right]^{1/2} = -\frac{1}{s}\left[1 - \frac{4xs}{(x+1)^2}\right]^{1/2} \tag{21}$$

so

$$\Phi(x) = \frac{2}{\frac{1}{s} - \psi(x)} = \frac{2s}{1 + [1 - \frac{4xs}{(x+1)^2}]^{1/2}} \tag{22}$$

Computing the derivative $\Phi'(-d)$ then yields

$$P(\Theta_e \leqslant t) = 1 - \left(\frac{4s^2(1+d)}{(1-d+[(1-d)^2+4ds]^{1/2})^2\,((1-d)^2+4ds)^{1/2}}\right)^a \tag{23}$$

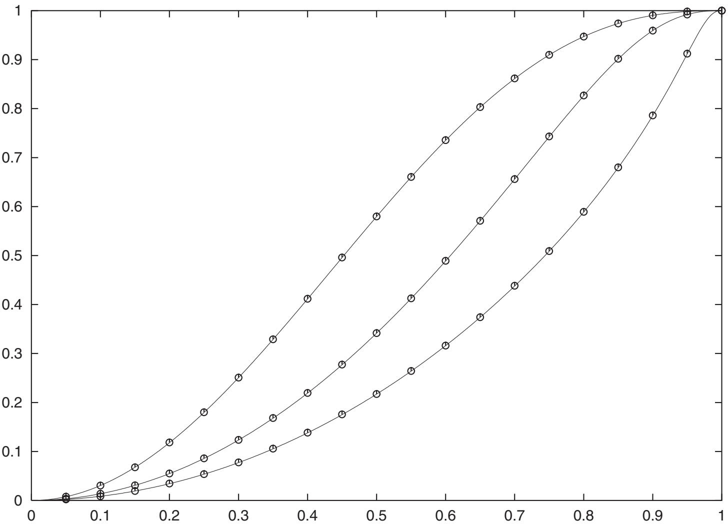

Fig. 3. The distribution of Θ_e for the half-plane for $d=0, 0.5, 0.9$. (d increases from left to right.) The solid lines are $\mathrm{SLE}_{8/3}$, and the open circles are the SAW.

For $d=0, 0.5$, and 0.9, this distribution and the results of the simulation for the SAW are shown in Fig. 3.

2.4. Passing Right

In addition to the distributions of the random variables X_e, Y_e, and Θ_e, we also consider the following probability. Fix a point in the upper half-plane. One can then ask if the random curve passes to the right or left of this point. For SLE this probability only depends on the polar angle of the point since SLE is invariant under dilations. This should also be true for the scaling limit of the SAW, since it is expected to be invariant under dilations. Schramm[12] rigorously derived an explicit formula for this probability for $\kappa<8$. For general κ it is given by a hypergeometric function, but for $\kappa=8/3$, his formula is quite simple. Denoting the probability that the curve passes to the right of a point with polar angle θ by $p(\theta)$, he showed that for $\kappa=8/3$

$$p(\theta)=\tfrac{1}{2}\,(1-\cos(\theta)) \tag{24}$$

In our simulations we study this probability by fixing a radius c and computing the probability the path passes to the right of $ce^{i\theta}$ for a large set of values of θ. The above function and the results of the SAW simulation are shown in Fig. 4. (Note that the horizontal axis in the figure is θ/π.)

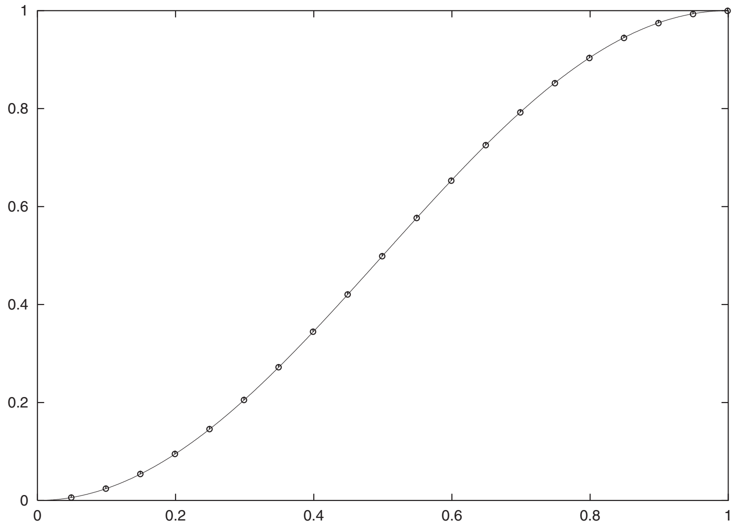

Fig. 4. The probability that the walk passes to the right of a point as function of its polar angle for walks in the half-plane. The horizontal axis is the angle divided by π, so that it ranges from 0 to 1. The solid line is the exact result for $\mathrm{SLE}_{8/3}$, and the open circles are the results of the simulation of the SAW.

2.5. The Cut-Plane

The cut-plane we consider is the plane with the non-negative real axis removed. Let $f(z) = \sqrt{z}$, with the branch cut along the positive real axis. Then f maps the cut-plane onto the upper half-plane. We will continue to denote curves in the upper half-plane by γ, and use $\hat{\gamma}$ to denote curves in the cut-plane. Given a curve $\hat{\gamma}$ in the cut-plane, $\gamma = f \circ \hat{\gamma}$ is a curve in the upper half-plane. So we can define the various random variables for the cut-plane by applying their definitions in the half-plane to $f \circ \hat{\gamma}$. We will put a $\hat{}$ on top of random variables defined on curves in the cut-plane. For the simulations it is useful to work out these definitions explicitly in terms of the curve $\hat{\gamma}$ in the cut-plane, rather than map each SAW in the cut-plane to the half-plane.

First consider $\hat{\Theta}_e$ and $\hat{\Theta}_f$ for $d = 0$. The map f simply divides the polar angle by 2, so for curves $\hat{\gamma}$ in the cut-plane,

$$\hat{\Theta}_e = \frac{1}{2\pi} \min\{\theta \colon ce^{i\theta} \in \hat{\gamma}\} \tag{25}$$

The random variable $\hat{\Theta}_f$ is the polar angle of the first intersection of $\hat{\gamma}$ with the circle, divided by 2π. If $d \neq 0$, the image of the semicircle under $z \to z^2$ is not a circle. We have not simulated $\hat{\Theta}_e$ or $\hat{\Theta}_f$ in this case.

To find the definition of $\hat{X}_e$, we first take $c = 1$. The image of the horizontal line $\{i + t : -\infty < t < \infty\}$ under $z \to z^2$ is a parabola whose axis is the horizontal axis and which opens to the right,

$$x = t^2 - 1, \qquad y = 2t \tag{26}$$

In the half-plane, X_e is the smallest t such that $i + t$ is on the curve. In the cut-plane, we consider all intersections of $\hat{\gamma}$ with the parabola and find the intersection with the smallest y-coordinate. Since $t = y/2$, $\hat{X}_e$ is one half of the y-coordinate of this "lowest" intersection. Equivalently,

$$\hat{X}_e = \min\{t\colon (t^2 - 1, 2t) \in \hat{\gamma}\} \tag{27}$$

SLE is invariant under dilations of the cut-plane, and the scaling limit of the SAW in the cut-plane is expected to have this invariance as well. So for $c \neq 1$ we can take the parabola to be

$$x = c(t^2 - 1), \qquad y = 2ct \tag{28}$$

and let

$$\hat{X}_e = \min\{t\colon (c(t^2 - 1), 2ct) \in \hat{\gamma}\} \tag{29}$$

$\hat{X}_f$ is the y-coordinate of the first intersection of $\hat{\gamma}$ with the parabola divided by $2c$.

To find the definition of $\hat{Y}_e$, we consider the image of $\{1 + it : 0 < t < \infty\}$ under $z \to z^2$. It is the upper half of a parabola whose axis is the horizontal axis and which opens to the left:

$$x = 1 - t^2, \qquad y = 2t, \qquad t > 0 \tag{30}$$

In the half-plane, Y_e is the smallest t such that $1 + it \in \gamma$, so in the cut-plane $\hat{Y}_e$ is one half of the y-coordinate of the lowest intersection of $\hat{\gamma}$ and the half parabola.

$$\hat{Y}_e = \min\{t\colon (1 - t^2, 2t) \in \hat{\gamma}, t > 0\} \tag{31}$$

More generally, we can let

$$\hat{Y}_e = \min\{t\colon (c(1 - t^2), 2ct) \in \hat{\gamma}, t > 0\} \tag{32}$$

$\hat{Y}_f$ is the y-coordinate of the first intersection with the parabola divided by $2c$.

We have defined the random variables in the cut-plane so that if the probability measure is conformally invariant, then they will have the same distribution as their counterparts in the half-plane. Rather than compare the distributions of the random variables X_e, Y_e, and Θ_e with those of $\hat{X}_e$, $\hat{Y}_e$, and $\hat{\Theta}_e$, we will compare all these distributions with the $\mathrm{SLE}_{8/3}$ predictions, Eqs. (9), (13), and (23). This tests both the conjecture that the scaling limit of the SAW is $\mathrm{SLE}_{8/3}$ and the conformal invariance of the SAW. For the random variables X_f, Y_f, and Θ_f, we do not know their distributions for $\mathrm{SLE}_{8/3}$. So we will directly compare the distributions of X_f, Y_f, and Θ_f with those of $\hat{X}_f$, $\hat{Y}_f$, and $\hat{\Theta}_f$. This tests the conformal invariance of the SAW.

3. THE SIMULATIONS

In all of our simulations the walks had one million steps. For the half-plane we ran the pivot algorithm for 10 billion iterations of the Markov chain. For the cut-plane we ran for 11.4 billion iterations. The simulation of Θ_e for $d \neq 0$ in the half-plane was done separately and consisted of $6.8B$ iterations. For walks with a million steps only about 5% of the proposed pivots are accepted. Of course, accepted pivots do not produce independent walks and for the random variables considered here most accepted pivots do not even change the value of the random variables. So the number of effectively independent samples is considerably less than the number of accepted pivots. Each of the simulations requires about a month on a 1.5 GHz PC. The exact speed of the simulation depends on the choice of the half-plane vs. cut-plane and how many random variables are simulated.

A walk with N steps is typically of size $N^{3/4}$, so to study the various random variables we take $c = lN^{3/4}$, where l is fairly small. Note that if we rescaled to make c equal to 1, the lattice spacing would be $(lN^{3/4})^{-1}$. We will refer to this quantity as the "effective lattice spacing." Note that l is the ratio of the scale used to define the random variable to the scale of the walk. So we must take l small to make the effect of the finite length of our walks negligible. But as l gets smaller, the effective lattice spacing gets larger. There is a second effect as l gets smaller. For smaller l, the fraction of the pivots that change the values of the random variables is smaller. So the statistical errors get larger as l gets smaller. We do not know *a priori* what value of l will be optimal, so we compute the distributions of each random variable for four different values of l in our simulations. The particular values of l that we use are determined by some experimentation with much shorter simulation runs. We do not use the same four values of l for the different random variables.

In Figs. 1 to 3 we show the distributions of X_e, Y_e, and Θ_e. (Throughout this paper we work with the cumulative distributions of our random variables rather than their densities since any simulation computes cumulative distributions. Computing densities requires taking numerical derivatives of the cumulative distributions, and so the densities would have larger statistical errors.) The solid curves are the exact distributions for $SLE_{8/3}$. The circles are the results of the simulation of the SAW. Figure 4 studies the probability that the walk passes to the right of a point in the upper half-plane as a function of the polar angle of the point. The solid curve is Schramm's exact result for $SLE_{8/3}$, and the circles are the results of the SAW simulation. In all of Figs. 1 to 4, one cannot see any difference between the SAW simulations and the exact curves for $SLE_{8/3}$. In Figs. 5 to 9 we plot the same four quantities, except that now we plot the result of the SAW simulation minus the $SLE_{8/3}$ functions. The first thing that should be observed in these figures is the scale of the vertical axis. It is quite small. In all but one of these figures the total vertical range shown is 0.007 or 0.7%. In Fig. 8 it is 0.008.

In Figs. 5 to 9 several values of l are shown. The nonzero effective lattice spacing means that we are simulating discrete random variables. So their distributions will be discontinuous. After subtracting off the continuous SLE distribution, the jumps will appear in the difference as rapid

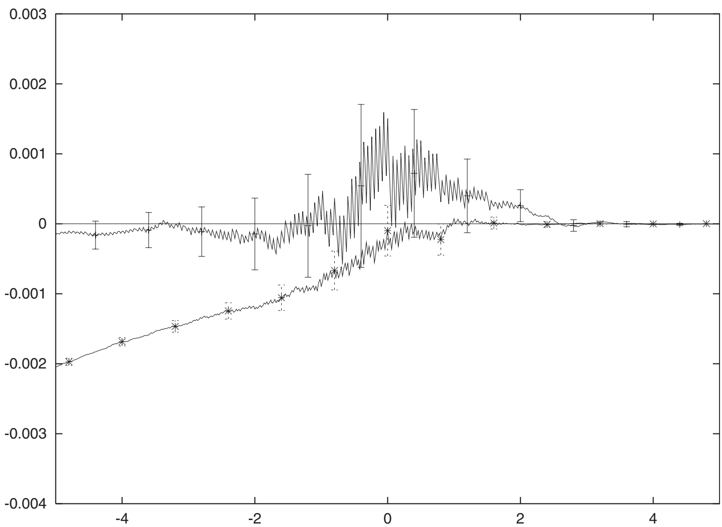

Fig. 5. Half-plane: The distribution of X_e for the SAW minus the distribution of X_e for $SLE_{8/3}$. The top curve, with the larger error bars drawn with solid lines, has $l = 0.01$, and the bottom curve has $l = 0.05$.

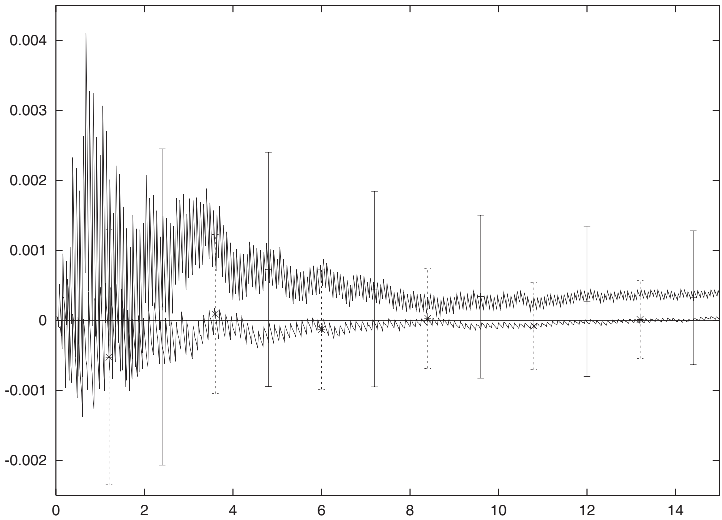

Fig. 6. Half-plane: The distribution of Y_e for the SAW minus the distribution of Y_e for $SLE_{8/3}$. The top curve, with the larger error bars drawn with solid lines, has $l = 0.002$, and the bottom curve has $l = 0.005$.

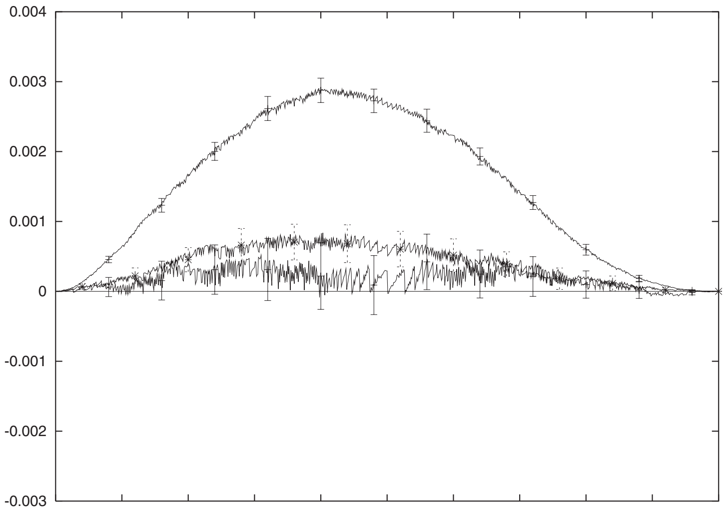

Fig. 7. Half-plane: For $d = 0$, the distribution of Θ_e for the SAW minus the distribution of Θ_e for $SLE_{8/3}$. The three curves shown are for $l = 0.2, 0.1, 0.05$, in order from top to bottom. As l decreases the finite length effects decrease, but the error bars and lattice effects grow larger.

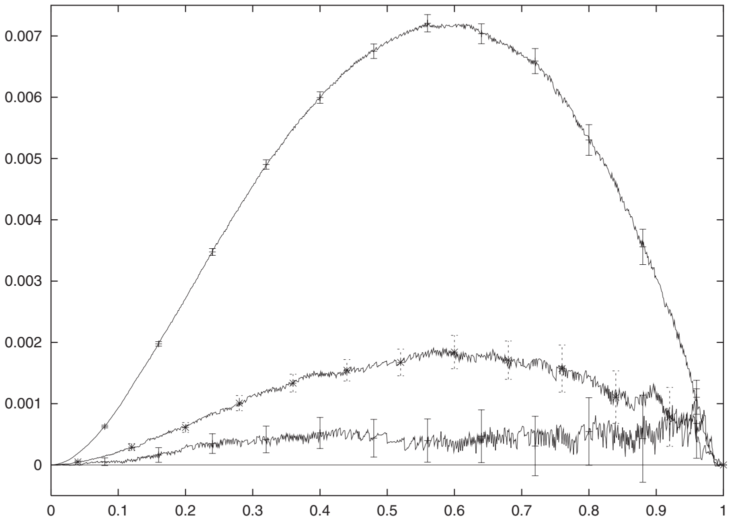

Fig. 8. Half-plane: For $d = 0.9$, the distribution of Θ_e for the SAW minus the distribution of Θ_e for $SLE_{8/3}$. The three curves shown are for $l = 0.2, 0.1, 0.05$, in order from top to bottom.

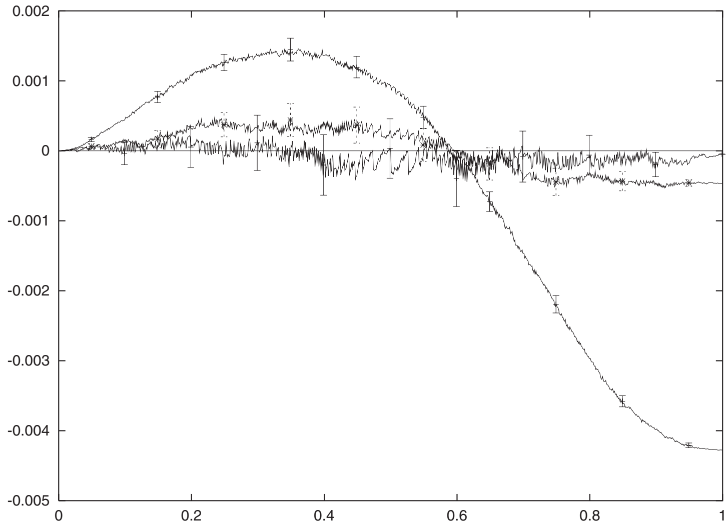

Fig. 9. Half-plane: The probability that the SAW passes to the right of a point as function of the polar angle of the point. The corresponding function for $SLE_{8/3}$ has been subtracted off. Going from top to bottom on the left half of the figure, the curves are $l = 0.2, 0.1, 0.05$.

oscillations. As l increases, the effective lattice spacing decreases, and so the oscillations are usually "faster" but smaller in amplitude. Also, for larger l a larger fraction of the pivots change the values of the random variables, and so a larger l typically produces smaller statistical errors. Both of these effects can be seen in all four of the plots.

However, as l becomes larger, the effect of the finite length of the walk will begin to be seen. The effect of the finite length is well illustrated by Fig. 7, which shows the distribution of Θ_e for $d = 0$. For the largest value of l shown, $l = 0.2$, the effect of the finite length of the walk is clearly seen—the curve differs from zero by many times the size of the statistical errors. This curve is the smoothest of the three curves and has the smallest statistical error bars. For $l = 0.1$ the finite length effect is greatly reduced, but is still statistically significant. The $l = 0.05$ curve seems to be the best of the values of l that were simulated. The maximum difference of the SAW and $\mathrm{SLE}_{8/3}$ distributions is only about 0.05%. Our simulations included a fourth value of l which is not shown, $l = 0.02$. For this value the larger effective lattice spacing and larger statistical errors produce a difference curve that is rougher and larger than the $l = 0.05$ curve. The behavior in Fig. 8 for the distribution of Θ_e for $d = 0.9$ is quite similar to Fig. 7, except that the nonzero l effects appear to be larger. Note that the vertical scales in the two figures are not the same.

In Fig. 5 the finite length effect is clearly seen in the $l = 0.05$ curve; for large negative values of t the deviation of this curve from zero is caused by the walk being too short. In Fig. 6 there are no obvious finite length effects; the deviation of the curve from zero appears to be primarily caused by the nonzero effective lattice spacing. The deviation is of the same order as the error bars and the oscillations. In Fig. 9 the finite length effects and nonzero effective lattice spacing effects are similar to those seen in Fig. 7. Note that the $l = 0.2$ and $l = 0.1$ curves are significantly different from zero at the right, corresponding to a polar angle of π. This effect is a result of the nonzero probability that the walk does not reach the semi-circle or that it crosses it, but ends inside the semi-circle. In both of these cases it is unclear whether the walk will pass to the right or left of the points on the semicircle. The algorithm must make some arbitrary choices in these cases.

Figures 10 through 13 show the same quantities as Figs. 5 to 7 and 9, but for the cut-plane. For the random variable $\hat{\Theta}_e$ (Fig. 12) and the probability of passing right of a point (Fig. 13), the agreement is again excellent. In both of these figures the vertical scale is 0.007, the same as in the corresponding figures for the half-plane. For the random variables $\hat{X}_e$ and $\hat{Y}_e$, Figs. 10 and 11, the agreement is not quite as good, but the deviations from the SLE results are still small. (In these two figures the vertical scale is two to three times larger than in the other figures.) For these two random variables it is harder to do accurate simulations for the following reason. In the cut-plane, $\hat{X}_e$ and $\hat{Y}_e$ depend on the intersections of the random curve with parabolas. It typically takes a longer length of curve to attain these intersections than for the lines involved in the definition of X_e and Y_e in the half-plane. So in the cut-plane we must use smaller values of l. For $\hat{X}_e$ in the cut-plane, the curves shown use $l = 0.002$ and $l = 0.005$ as compared to $l = 0.01$ and $l = 0.05$ for X_e in the half-plane. Even with these small values of l, the finite length effects are still quite visible in Fig. 10. The deviation of the curves from 0 for the most negative values of t is pronounced. This is the part of the distribution that is particularly sensitive to the need for very long walks to hit the parabola. Of course, small values of l mean a large effective lattice spacing and large statistical errors. For $\hat{Y}_e$ in the cut-plane, the values of l shown are 0.0005 and 0.001, as compared to 0.002 and 0.005 for Y_e in the half-plane. The finite length effects in Fig. 11 can be seen in the substantial deviation of the curves from 0 for large t, again a reflection of the need for long walks to reach the parabola.

The scaling limits for the SAW in the half and cut-planes are conjectured to be related by the conformal transformation, but there is no

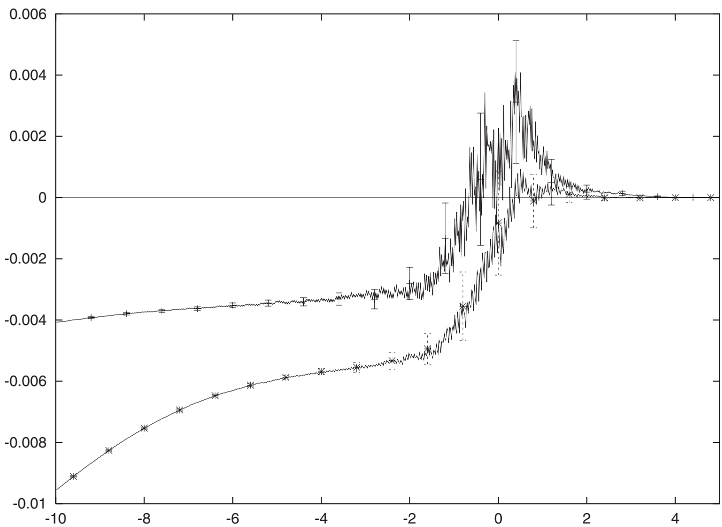

Fig. 10. Cut-plane: The distribution of $\hat{X}_e$ for the SAW minus the distribution of $\hat{X}_e$ for $SLE_{8/3}$. The top curve, with the larger error bars drawn with solid lines, has $l = 0.002$, and the bottom curve has $l = 0.005$.

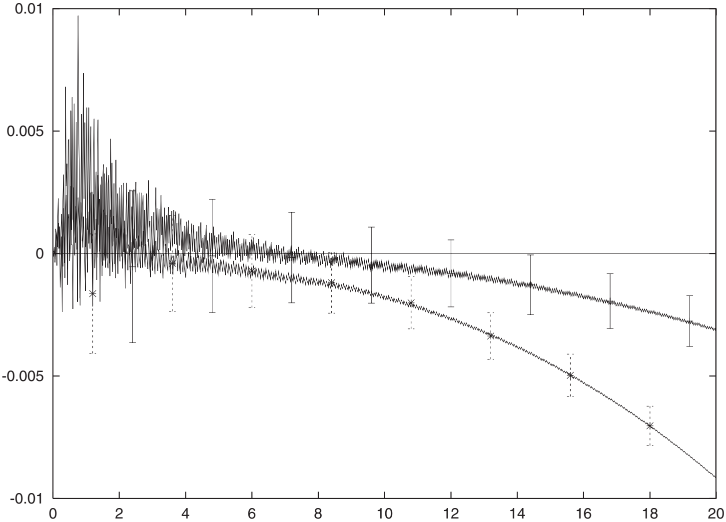

Fig. 11. Cut-plane: The distribution of $\hat{Y}_e$ for the SAW minus the distribution of $\hat{Y}_e$ for $SLE_{8/3}$. The top curve, with the larger error bars, has $l = 0.0005$, and the bottom curve has $l = 0.001$.

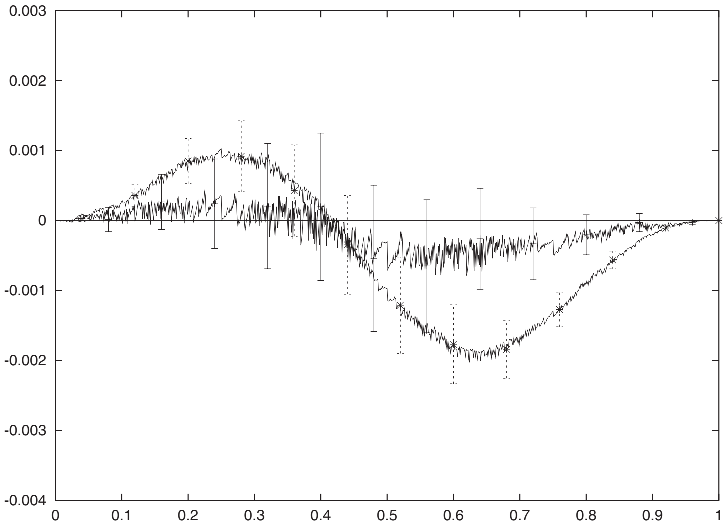

Fig. 12. Cut-plane: The distribution of $\hat{\Theta}_e$ for the SAW minus the distribution of $\hat{\Theta}_e$ for $SLE_{8/3}$. The curve with the greater deviation from the horizontal axis and the error bars drawn with dashed lines has $l = 0.05$. The other curve has $l = 0.02$.

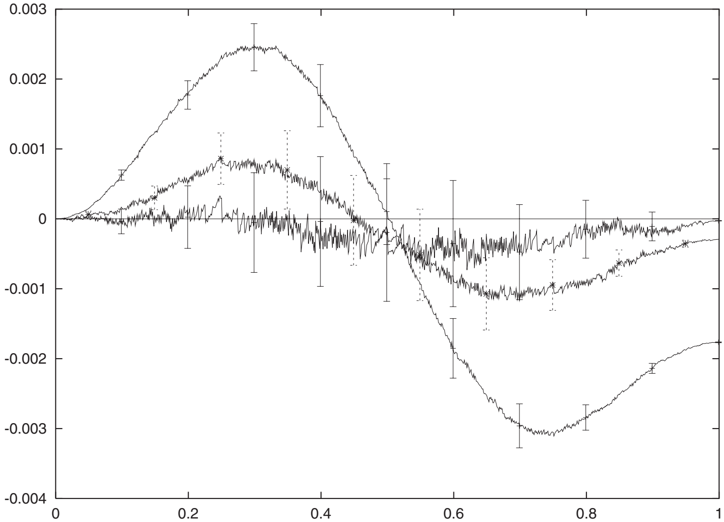

Fig. 13. Cut-plane: The probability that the walk passes to the right of a point as function of the polar angle of the point. Going from top to bottom on the left half of the figure, the curves have $l = 0.1, 0.05, 0.02$.

reason that the finite length effects in the two cases should be related. Indeed, the simulations show they are quite different. For example, compare the curves for the largest values of l in Figs. 7 and 12. The curve in Fig. 7 is always positive, looking roughly like the first half of a sine wave, while the curve in Fig. 12 is both positive and negative.

Finally, we consider the random variables X_f, Y_f, and Θ_f in the half and cut-planes. We don't know the exact distributions of these random variable for $SLE_{8/3}$, but we can still compare the distributions we get from the simulations of the SAW in the half-plane with the simulations for the cut-plane. Recall that $\hat{X}_f$, $\hat{Y}_f$, and $\hat{\Theta}_f$ (the random variables in the cut-plane) were defined so that they will have the same distribution as their counterparts in the half-plane if the SAW is conformally invariant. If we simply plot the distributions themselves, they agree so well that the difference cannot be seen in the plots. So instead of plotting the distributions, we plot the distributions minus various reference functions. These reference functions are quite *ad hoc*. They are chosen to be simple functions that are relatively good approximations to the distributions. They are defined as follows. For X_f and $\hat{X}_f$ we use the function

$$F(t) = \tfrac{1}{2}\left(\tanh(1.16t) + 1\right) \tag{33}$$

For Y_f and $\hat{Y}_f$ we use the distribution of Y_e for $\mathrm{SLE}_{8/3}$, i.e.,

$$F(t) = 1 - (1 + t^2)^{-5/16} \tag{34}$$

For Θ_f and $\hat{\Theta}_f$ we use

$$F(t) = t - 0.12 \sin(2\pi t) - 0.009 \sin(4\pi t) \tag{35}$$

We emphasize that these are not meant to be highly accurate approximations of the distributions of X_f, Y_f, and Θ_f. One could find better approximations with more complicated functions. The only purpose of these functions is to provide a convenient reference with respect to which we can plot the distributions for the half and cut-planes and compare them.

Figures 14 to 16 compare the distributions of X_f, Y_f, and Θ_f in the half-plane with their analogs for the cut-plane. Again, the most important features of these graphs is the small scale of the vertical axis. For X_f and Θ_f the difference between the distributions in the half and cut-planes is very small. For Y_f the difference is somewhat larger for large values of t, but still small. We attribute this greater difference to the larger finite length effects in the cut-plane. It can take a walk in the cut-plane a long time to reach the parabola involved in the definition of $\hat{Y}_f$.

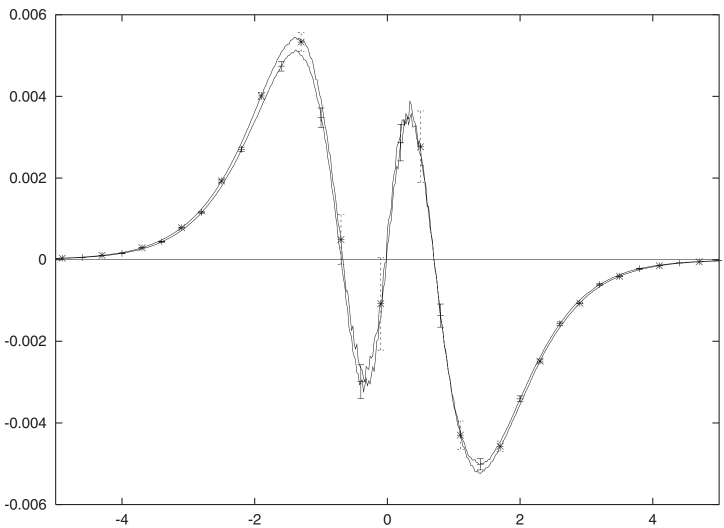

Fig. 14. For the half and cut planes the distribution of X_f for the SAW simulation minus the reference function (33) is shown. The half-plane simulation used $l = 0.05$, and the cut-plane used $l = 0.02$. The finite length effects are small for X_f, resulting in excellent agreement between the two curves.

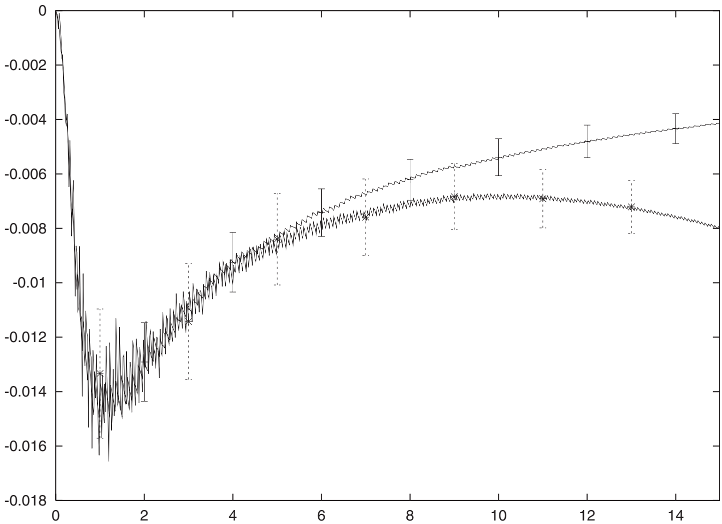

Fig. 15. For the half and cut planes the distribution of Y_f for the SAW simulation minus the reference function (34) is shown. The half-plane simulation used $l = 0.005$ and the cut-plane simulation used $l = 0.001$. Even with these small values of l, the finite length effects produce a noticeable difference between the curves for large t.

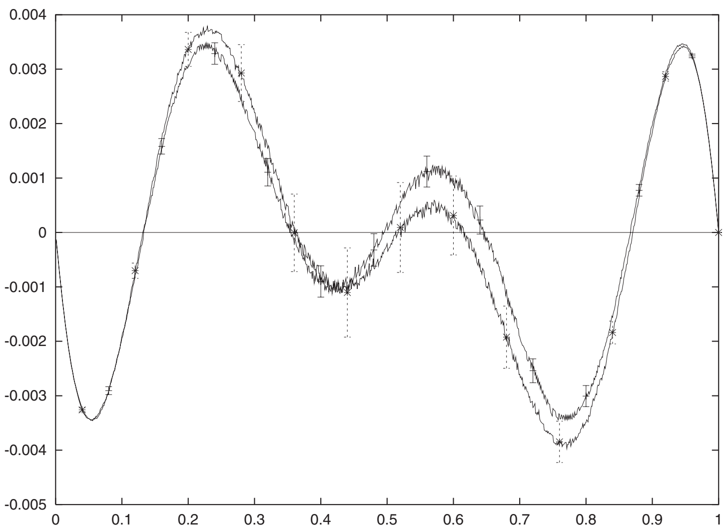

Fig. 16. For the half and cut planes the distribution of Θ_f for the SAW simulation minus the reference function (35) is shown. The half-plane simulation used $l = 0.1$ and the cut-plane simulation used $l = 0.05$. The half-plane curve has error bars drawn with solid lines, while the cut-plane uses dashed error bars.

4. ALGORITHMIC CONSIDERATIONS

The pivot algorithm is used for our simulations. (This algorithm is discussed in ref. 9.) The algorithm picks a site at random along the walk, called the pivot point, and picks a random element of the group of symmetries of the lattice about the point. This group element is applied to the part of the walk after the pivot point. The result is a new nearest neighbor walk, but it need not be self-avoiding or lie in our domain (the upper half-plane or the cut-plane). The walk is accepted only if both of these conditions are meet. Otherwise the proposed walk is rejected and the current walk is counted as another state in the Markov chain. The Markov chain trivially satisfies detailed balance with respect to the uniform measure on all self-avoiding walks with N steps that start at the origin and stay inside our domain. In the appendix we show that it is irreducible.

The speed of the pivot algorithm is typically measured by considering the average time needed to produce an accepted pivot. The algorithm may be implemented[(1)] so that this time grows with the number of steps, N, as $O(N^q)$ with $q < 1$. The exact value of q is not known and probably depends on details of the implementation, but simulations indicate the implementation in ref. 1 has $q < 0.57$ in two dimensions. (This estimate is based on simulations of the walk in the full plane, not the half or cut planes.)

There are two main steps in the pivot algorithm, and both would seem to require a time $O(N)$ per accepted pivot. The first is the test for self intersections to see if the new walk should be accepted. The second is actually carrying out the pivot. To test for self-intersections quickly, we take advantage of the fact that the walk ω only takes nearest neighbor steps. Rather than simply checking if $\omega(i) = \omega(j)$, we compute the distance $d = \|\omega(i) - \omega(j)\|_1$. If d is nonzero then we can conclude not just that $\omega(i) \neq \omega(j)$, but also that

$$\omega(i') \neq \omega(j'), \qquad \text{if} \quad |i - i'| + |j - j'| < d \tag{36}$$

Thus we can rule out a large number of potential self intersections if d is large. Since it takes a time $O(N)$ to simply write down a walk with N steps, the second step of carrying out the pivot would seem to require a time that is $O(N)$ per accepted pivot. To do better, the key idea is to not carry out the pivot each time a pivot is accepted. Instead we keep track of which pivots have been accepted and only carry them out after a certain number have been accepted. Details of this implementation of the pivot algorithm may be found in ref. 1.

In the usual implementation of the pivot algorithm one chooses the pivot point by giving equal probability to all the points on the walk. One can, however, take the probability of picking the ith site along the walk to be $p(i)$, where $p(i)$ is a function whose sum is 1. The only constraint is that the $p(i)$ must be positive. If one is interested in the distribution of the end-point of the walk, then every accepted pivot changes this random variable. For this random variable it does not appear that anything could be gained by making $p(i)$ non-uniform. However, there is a substantial benefit to using a non-uniform $p(i)$ for the random variables in this paper. All of our random variables typically depend only on a short segment of the walk near the origin. (The smaller l is, the shorter the segment.) So most accepted pivots do not produce any change in the random variable. This suggests that it might be worthwhile to choose pivot locations near the start of the walk more often than pivot locations far from the start. For the simulations in this paper we define $p(i)$ as follows

$$p(i) = \begin{cases} 8c, & \text{if } 0 \leqslant i < \frac{1}{5} N \\ 4c, & \text{if } \frac{1}{5} N \leqslant i < \frac{2}{5} N \\ 2c, & \text{if } \frac{2}{5} N \leqslant i < \frac{3}{5} N \\ c, & \text{if } \frac{3}{5} N \leqslant i < N \end{cases} \tag{37}$$

where $c = \frac{5}{16} N^{-1}$ so that the sum of the $p(i)$ is 1. This is a rather *ad hoc* choice, but a crude test indicates that for a given number of iterations of the algorithm, it typically reduces the standard deviation of the random variable by a factor of two. A systematic study of the effect of $p(i)$ would be useful.

For each of the six random variables we consider four different values of l. We also consider four values of l for the probability of passing to the right of a given point. Thus there are 28 different observables to be computed, and some care is necessary to be sure that the time required for this part of the simulation does not dominate the simulation. All of these observables require finding intersections of the walk with a given curve (a line, parabola, or circle). Searching through the walk one step at a time for these intersections would be disastrous, since it would require a time $O(N)$. Such a search is easily avoided. At a given site in the walk we do not simply check if the next step intersects the curve. Instead we compute the distance from the site to the curve. The walk must take at least this many steps before it can intersect the curve, so we can jump ahead this many steps in the walk before we check again for an intersection.

APPENDIX A. PROOF OF IRREDUCIBILITY

In this appendix we prove that the pivot algorithm is irreducible in the half-plane and cut-plane that we have been considering. The proof is very similar to the proof for the full plane.[9] We show that for any self-avoiding walk in the half-plane (cut-plane, respectively), there is a sequence of pivot operations which "unfold" the walk into a straight line and such that each walk produced in this unfolding process is self-avoiding and remains in the half-plane (cut-plane, respectively).

We first consider the half-plane. The restriction is that except for the starting point of the walk at the origin, the walk must remain strictly above the horizontal axis. We will show that the number of turns in the walk can be decreased by one. We denote the sites in the walk by $\omega(i)$ where $i = 0, 1,\ldots, N$. We will say there is a turn at $\omega(i)$ if $\omega(i-1), \omega(i)$ and $\omega(i+1)$ are not co-linear.

We will consider cases based on the direction of the last step of the walk. If it is to the right, i.e., $\omega(N) = \omega(N-1) + (1, 0)$ we proceeds as follows. Let l be the largest integer such that the line $y - x = l$ contains a site on the walk. So the walk is entirely below or on this line. Let i be the largest integer such that $\omega(i)$ is on this line. Since the last step of the walk is to the right, $\omega(N)$ is not on this line. So $i < N$. Since the first step of the walk in the half-plane must be up, i cannot be 0. We take $\omega(i)$ as the pivot point and reflect the portion of the walk from $\omega(i)$ to $\omega(N)$ in the line $y - x = l$. The reflected portion of the walk lies entirely above the line, so the reflection does not produce self-intersections. Furthermore, since the walk was on or below the line, the reflection can only increase the y coordinates of points on the walk. So the new walk is still in the upper half plane. The walk before this reflection has a turn at $\omega(i)$ and the reflected walk does not. The reflection does not add any turns to the walk, so the total number of turns decreases by one. If the final step of the walk is to the left, i.e., $\omega(N) = \omega(N-1) - (1, 0)$, we use an analogous procedure with lines $y + x = l$ to reduce the number of turns in the walk.

Now suppose that the final step of the walk is either up or down, i.e., $\omega(N) = \omega(N-1) \pm (0, 1)$. Consider the vertical line which contains this last step. First suppose that the walk lies entirely to the right of or on this vertical line. Let $i < N$ be the largest integer such that there is a turn at $\omega(i)$. (Of course, if there are no turns the walk is a straight line and we are done.) The walk is a straight segment from $\omega(i)$ to $\omega(N)$ which lies on the vertical line. We take $\omega(i)$ as the pivot point and perform a rotation of 90 degrees (-90, respectively) if the last step of the walk is up (down, respectively). This rotates the segment from $\omega(i)$ to $\omega(N)$ to the left of the vertical line and removes the turn at $\omega(i)$. No new turns are added to the walk, so the total number of turns decreases by one. If the walk likes entirely to the right or on the vertical line containing the last step, an analogous argument shows the number of turns can be reduced by one.

Now suppose that the walk contains sites on both sides of the vertical line which contains the last step of the walk. Let d be the horizontal width of the walk:

$$d = \max\{x\colon (x, y) = \omega(i), \text{ for some } i, y\} \\ - \min\{x\colon (x, y) = \omega(i), \text{ for some } i, y\} \tag{38}$$

We will show that d can be increased. Let l be the smallest integer such that the vertical line $x = l$ contains sites in the walk. So the walk lies on or to the right of this line. Note that $\omega(N)$ is not on this line. Let $i < N$ be the largest integer such that $\omega(i)$ is on this line. We take $\omega(i)$ as the pivot point and reflect the walk from $\omega(i)$ to $\omega(N)$ in the line $x = l$. This increases the width of the walk. (The argument is the same as that given in ref. 9.) The reflection does not change the y-coordinate of points on the walk, so the new walk is still in the upper half-plane. Note that in the new walk the last step is in the same direction as before, i.e., either up or down. So we can repeat this procedure to increase d until we obtain a walk which lies entirely on or to one side of the vertical line containing the final step. When we reach such a walk we apply the procedure of the proceeding paragraph to reduce the number of turns by one. This completes the proof for the case of the half-plane.

Now consider the cut-plane. The restriction now is that the walk cannot contain sites of the form $(x, 0)$ with $x \geqslant 0$, except for the starting point at the origin. We again consider cases based on the direction of the last step of the walk. If it is to the right, we proceed as in the half-plane algorithm. Note that the line involved, $y - x = l$, must have $l \geqslant 0$ since the walk starts at the origin. (In fact l must be at least 1, but we do not need this.) The reflected portion of the walk will lie above this line while the cut, the non-negative real axis, lies below it. So the reflection produces a walk that lies in the cut-plane.

If the last step of the walk is to the left, a different procedure is needed to avoid producing a walk that intersects the cut. Consider the lines $x - y = l$ and $x + y = l$. They intersect at $(l, 0)$ and divide the plane into four quadrants which we will describe as being left, right, above and below the point $(l, 0)$. We take l to be the smallest integer such that the sites on the walk lie in the quadrant to the left of $(l, 0)$ or on the lines. (l is necessarily positive.) We then let i be the largest integer such that $\omega(i)$ is on one of the two lines. (It is not N since the last step of the walk is to the left.) Note that $\omega(i)$ cannot be $(l, 0)$. We take $\omega(i)$ as the pivot point and reflect the walk from $\omega(i)$ to $\omega(N)$ in the line containing $\omega(i)$. The reflected portion of the walk will lie either in the quadrant above or below $(l, 0)$ and so cannot intersect the cut.

If the last step of the walk goes up or down we use the algorithm for the half-plane. There is a subtle point here. Recall than when the walk has points on both sides of the vertical line containing the last step, we chose l so that the walk is to the right of or on the vertical line $x = l$. For the half plane we could have chosen it so that "right" is replace by "left." For the cut-plane this choice could result in a reflected walk that intersects the cut. To see that our choice does not produce a walk that intersects the cut, we observe that since the walk starts at the origin, l must be negative. The reflected portion of the walk will lie on or to the left of the line $x = l$, and so will not intersect the cut. This completes the proof for the cut-plane.

ACKNOWLEDGMENTS

Oded Schramm pointed out to me that one can find the conformal map from the half-plane with the circular arc removed onto the half-plane and so compute the distribution of Θ_e. Some of this work was performed at the Mathematics Department of the Université Paris-Sud and the Institute des Hautes Études Scientifiques, and the author thanks them for their hospitality. This work was supported by the National Science Foundation (DMS-9970608 and DMS-0201566).

REFERENCES

1. T. Kennedy, A faster implementation of the pivot algorithm for self-avoiding walks, *J. Stat. Phys.* **106**:407–429 (2002), archived as cond-mat/0109308 in arXiv.org.
2. T. Kennedy, Monte Carlo tests of stochastic Loewner evolution predictions for the 2D self-avoiding walk, *Phys. Rev. Lett.* **88**:130601 (2002), archived as math.PR/0112246 in arXiv.org.
3. G. Lawler, O. Schramm, and W. Werner, Values of Brownian intersection exponents I: Half-plane exponents, *Acta. Math.* **187**:237–273 (2001), archived as math.PR/9911084 in arXiv.org.
4. G. Lawler, O. Schramm, and W. Werner, Values of Brownian intersection exponents II: Plane exponents, *Acta. Math.* **187**:275–308 (2001), archived as math.PR/0003156 in arXiv.org.
5. G. Lawler, O. Schramm, and W. Werner, Values of Brownian intersection exponents III: Two-sided exponents, *Ann. Inst. H. Poincaré PR* **38**:109–123 (2002), archived as math.PR/0005294 in arXiv.org.
6. G. Lawler, O. Schramm, and W. Werner, Conformal invariance of planar loop-erased random walks and uniform spanning trees, archived as math.PR/0112234 in arXiv.org.
7. G. Lawler, O. Schramm, and W. Werner, On the scaling limit of planar self-avoiding walk, archived as math.PR/0204277 in arXiv.org.
8. G. Lawler, O. Schramm, and W. Werner, Conformal restriction: The chordal case, archived as math.PR/0209343 in arXiv.org.
9. N. Madras and G. Slade, *The Self-Avoiding Walk* (Birkhäuser, Boston/Basel/Berlin, 1993).

10. S. Rohde and O. Schramm, Basic properties of SLE, archived as math.PR/0106036 in arXiv.org.
11. O. Schramm, Scaling limits of loop-erased random walks and uniform spanning trees, *Israel J. Math.* **118**:221–288 (2000).
12. O. Schramm, A percolation formula, *Electron. Comm. Probab.* **6**:115–120 (2001), archived as math.PR/0107096 in arXiv.org.
13. S. Smirnov, Critical percolation in the plane: conformal invariance, Cardy's formula, scaling limits, *C. R. Acad. Sci. Paris Sér. I Math.* **333**:239–244 (2001).