

18.677 Final Project: Critical Ising and Nested CLE_3

Andrew Lin and Michael Ren

Abstract

This paper aims to discuss the arguments presented in the two papers “Conformal invariance of crossing probabilities for the Ising model with free boundary conditions” [1] and “The scaling limit of critical Ising interfaces is CLE_3 ” [2]. We explain the definitions of fundamental objects required to set up the scaling limit, and we present the lemmas and proofs that are relevant to the material we have learned in 18.677 this semester. Our goal is to connect definitions of objects like the CLE, CDE, and FAE to those that have been presented in similar settings and bring the intuition we’ve developed over the course to this particular proof.

1 Fundamental definitions: the FK-Ising model and CLE_κ

The central question we are interested in is the convergence of the spin interfaces of the **Ising model** at critical temperature to certain random loops which form a **conformal loop ensemble**. We’ll begin by defining these terms and explaining their important characteristics. Throughout, let Ω be a domain (a strict open subset of \mathbb{C}), and let Ω^ε be its ε -grid discretization.

Definition 1 (Ising model)


For a discrete grid Ω^ε , the (two-dimensional) **Ising model** is a statistical mechanics model which assigns a spin $s_x \in \{\pm 1\}$ to each vertex x of Ω^ε with probability proportional to $\exp(-\beta H)$, where $H = -\sum_{x \sim y \in \Omega^\varepsilon} s_x s_y$.

In other words, at large β (low temperature), the system tends to favor configurations with aligned spins, while for small β , the system exhibits disorder and the spins are near-independent. In all cases, the configuration of spins can be described as collections of clusters of +1s and -1s, and this motivates an alternate way to think about generating such configurations:

Definition 2 (Fortuin-Kasteleyn model)

On a graph $G = (V, E)$, in this paper taken to be the subgraph of $\varepsilon\mathbb{Z}^2$ cut out by some domain, the **Fortuin-Kasteleyn model**, also known as the FK model or random cluster model, is a probability measure on partitions of V depending on two parameters $q > 0$ and $p \in [0, 1]$. First, we perform **bond percolation** on G with probability p , in which we independently erase each edge in E with probability $1 - p$. The connected components of the resulting subgraph produce a random partition of V . We then weight each such configuration arising from percolation by q^C , where C is the number of connected components of the subgraph (that is, the number of parts in the partition of V).

Figure 1: An example configuration of the FK model. Edges that are not erased are in black, and the corresponding dual edges are in green (edges to the boundary ∂D not shown). On the right, + and - spins are indicated in red and blue, corresponding to the coupling with the Ising model.



The FK model unifies different models of statistical physics through the parameter q . When $q = 1$, we get bond percolation (all configurations equally likely). When $q > 1$ is an integer, we obtain the **Potts model** by drawing from the FK model and then randomly assigning each component a uniform random spin from $\{1, \dots, q\}$. In particular, we recover the Ising model by taking $q = 2$. This coupling of the random cluster and Ising models is known as the **FK-Ising model**, and the inverse temperature β in the Ising model corresponds to a bond percolation probability of $p = 1 - e^{-2\beta}$.

A fundamental result in statistical mechanics due to Onsager [5] is that the 2D Ising model exhibits a **phase transition** at some value β_c in which, roughly speaking, the macroscopic properties of the typical configuration changes from a regime of (exponential) correlation decay for $\beta < \beta_c$ to a regime of long-range (nonzero correlation) order for $\beta > \beta_c$. The situation at $\beta = \beta_c$ is more delicate and (as with many 2D statistical physics models) displays a conformally invariant scaling limit. This was long conjectured by physicists, and it is proven in the case of **positive boundary conditions** (where the boundary of our domain is fixed to have positive +1 spins) by the papers which we survey.

Every bond percolation configuration on a planar graph G corresponds to a dual bond percolation configuration on the dual of G by taking the dual of every erased edge. For G a grid graph, the dual of G is also a grid graph (with boundary given by points adjacent to dual vertices), so it makes sense to speak of the FK model on the dual of G . It turns out that this natural dual correspondence between bond percolation gives a coupling between the FK model on G with parameters q and p and the FK model on the dual of G with parameters q and p^* , where $\frac{p}{1-p} \cdot \frac{p^*}{1-p^*} = q$. In particular, for $p_c = \frac{\sqrt{q}}{1+\sqrt{q}}$ the model exhibits **self-duality** (since $\frac{p_c}{1-p_c} = \sqrt{q}$). As is common for statistical physics models, self-duality and criticality for the FK-Ising model are intricately related, and for $q = 2$ these values β_c and p_c do indeed satisfy $p_c = 1 - e^{-2\beta_c}$. Henceforth, we will always take the FK-Ising model to have these critical self-dual parameters and describe the scaling limit mentioned above.

The main subject of concern in this survey is that of **Ising loops** and their scaling limits. On a grid graph G in a Jordan domain D , we consider the critical Ising model with positive boundary conditions. Given a spin configuration on G , we define an Ising loop to be an oriented closed path on the dual G^* of G such that every edge in G^* appears at most once and separates a +1 spin on its left with a -1 spin on its right, as we see in Fig. 2. When we impose positive boundary conditions, the collection of all Ising


loops in a configuration is well-defined. Given a fixed Jordan domain D , we can then consider positive boundary condition Ising model on $D \cap \epsilon\mathbb{Z}^2$. As $\epsilon \rightarrow 0$, our question is **whether the collection of Ising loops have a scaling limit**; the answer is that this collection converges to a nested CLE_3 .

To formulate this result rigorously, we need the notion of a leftmost Ising loop.

Definition 3

Given an Ising loop, the set of $+1$ spins on the left of its edges form what we call a **weak path** (that is, a cycle of vertices in G such that every pair of adjacent vertices are orthogonally or diagonally adjacent in the grid). Meanwhile, a **strong path** is a cycle of vertices in G in which adjacent vertices may only be orthogonally adjacent. If the weak path corresponding to an Ising loop is contained in a strong path of $+1$ spins, then we say that the loop is **leftmost**. Analogously, we can consider the weak path of -1 spins on the right of the edges in an Ising loop, and we say that the path is **rightmost** if this weak path is contained in a strong path of -1 spins.

Figure 2: Positive boundary conditions imposed on the configuration from the above figure. An Ising loop is shown in purple which is **not** leftmost, because (for instance) at the marked purple point it could have continued downward and looped around the other blue cluster.



Remark 4. An alternative interpretation of the “leftmost” and “rightmost” loops is as follows: imagine that we are walking along the Ising loop. The loop may only fail to be leftmost or rightmost at a 90 degree turn (such as the purple point above), and only if the four (ordinary) vertices around the dual vertex alternate in color. So at every dual vertex, if we always take the leftmost edge with a $+1$ spin on the left and a -1 spin on the right, then the path is leftmost. In other words, at every dual vertex we first turn left to see if we can go in that direction, i.e. that direction has a $+1$ spin on the left and a -1 spin on the right, then we turn straight forward, then we turn right (and one of these will always be possible). For the rightmost loop, we do the same but prioritize turning right.

The reason that we have defined leftmost and rightmost loops as above is that given a dual edge dividing ± 1 spins, there could feasibly be many Ising loops using this edge. For example, if the spins were in an antiferromagnetic checkerboard pattern, then any loop that turns at every step is an Ising


loop. However, the leftmost and rightmost loops through a given dual edge are uniquely defined (by the exploration process above), and furthermore it is clear that every Ising loop through this edge will be contained between the leftmost and rightmost loops. We will only need to consider leftmost loops for the scaling limit, because it will turn out (as $\varepsilon \rightarrow 0$) that all macroscopic Ising loops (whose lengths are bounded away from 0) are, with high probability, close enough to a leftmost loop that their scaling limit will be the same as that of just the leftmost loops.

For the proof approach, we will also need the notion of an **FK loop**. In the FK-Ising model, we have a bond percolation on our grid graph G (the black edges), as well as a corresponding bond percolation on the dual graph G' . The interface between the selected (hereafter “open”; deleted edges will be called “closed”) edges in the primal and dual percolations consist of a collection of cycles that are **vertex-disjoint** (share no common vertices), which we call FK loops.

Remark 5. *Another perspective on these FK loops is that if we divide every $\varepsilon \times \varepsilon$ cell in our grid into four $\varepsilon/2 \times \varepsilon/2$ smaller cells, then the primal and dual open edges form a subgraph of this finer grid, and the FK loops are the dual graph of this subgraph. These loops will be vertex-disjoint because at any vertex of this doubly dual graph (which is contained in $\frac{\varepsilon}{2}\mathbb{Z} \times \frac{\varepsilon}{2}\mathbb{Z} + (\frac{\varepsilon}{4}, \frac{\varepsilon}{4})$), at least two of its neighbors must be blocked by a primal or dual open edge (look at the two edges of the original grid closest to the vertex).*

Because FK loops do not touch primal or dual open edges and Ising loops consist of dual open edges, and there is always an FK loop “hugging the boundary” (see Fig. 3 for a visual guide), Ising loops are always nested inside FK loops. The main result crucially relies on this interplay between Ising and FK loops.

Figure 3: One FK loop, drawn in orange, in the configuration from above.



Since the FK loops are vertex-disjoint cycles, there is a natural integer we can assign to each loop corresponding to the nesting order of the loop, which we call the **level** of the loop. Specifically, we say that an FK loop is of level k if it is contained inside exactly $k - 1$ other FK loops. Hence the level k loops cut out disjoint regions of our domain. We refer to level 1 FK loops as **outermost** loops. Note that all outermost loops touch the boundary of our domain FK loops satisfy a nice Markov-type property: in the

FK-Ising model with positive boundary conditions, conditioned on the set of outermost FK loops, the model inside each outermost loop is distributed as an FK-Ising model with **free boundary conditions** (meaning no contribution to the weighting factor due to boundary effects), since clusters inside an FK loop are disconnected from those outside the loop.

Similarly, we assign a level to each Ising loop to capture the nesting order of the loop. Here, our task is slightly trickier as the loops are not vertex-disjoint. We say that a (not necessarily leftmost) Ising loop is of level 1, or outermost, if it is not strictly contained inside any other Ising loop. For $k > 1$, an Ising loop is of level k if it is contained strictly inside an Ising loop of level $k - 1$ and there is no closed weak path of $(-1)^k$ spins between the two loops. Given a loop of level k , we can repeatedly find its “outer” loop to obtain a chain of loops of level 1 to k . The picture here is roughly that this chain of loops will alternate between going clockwise and counterclockwise, and spins are essentially all the same between two consecutive loops in the chain.

One of the ideas behind the proof is that to understand Ising loops, we only need to understand outermost Ising loops. This is because inside an outermost Ising loop (which has positive boundary conditions outside) we have an Ising model with negative boundary conditions. By symmetry of the spins, the loops inside have the same law as if we had an Ising model with positive boundary conditions inside. Thus, if we have a conformally invariant scaling limit of the outermost Ising loops, then the collection of all Ising loops will be obtained by iteratively nesting the limiting outermost Ising loop collection inside the innermost Ising loops obtained so far. Our claim is that the limiting law on the whole Ising loop collection is **nested CLE₃**, and we’ll explain that term now.

Definition 6

The conformal loop ensemble CLE_κ for $\kappa \in (\frac{8}{3}, 4]$ is a random collection of disjoint simple loops in a simply connected domain D . One way it can be constructed is by taking a Brownian loop soup of intensity $c = \frac{(3\kappa-8)(6-\kappa)}{2\kappa}$ and taking the outer boundaries of the clusters formed [6].

Like SLE, the dimension of CLE_κ is $1 + \frac{\kappa}{8}$ (because the loops are “SLE-type”). In fact, the conformal loop ensembles are the unique family of random loop collections satisfying certain natural conformal invariance and restriction axioms:

Theorem 7 (Sheffield-Werner [6])

Suppose that a random loop ensemble L in a domain D satisfies

- conformal invariance: for any other domain D' , the ensemble in D' is the pushforward measure of L by any conformal map from D to D' ,
- local finiteness: if D is the unit disk then for any $\epsilon > 0$ there are finitely many loops in the ensemble of diameter greater than ϵ ,
- restriction: given a subdomain D' of D , conditioned on the loops in the ensemble that do not lie completely in D' , the loops that do lie completely in D' are distributed as the ensemble in D' , which is D' and the interiors of loops that do not lie completely in D' removed from D .

Then it is CLE_κ for some $\kappa \in (\frac{8}{3}, 4]$.

(It is derived in [6] that as long as the intensity of the loop-soup is small enough, so that there are multiple clusters with high probability, conformal invariance through a Markov property, plus a calculation with a particular SLE martingale, gives c as a factor which shows up in a Radon-Nikodym derivative.) Nested CLE_κ can be obtained by repeatedly sampling CLE_κ inside the innermost loops already obtained, iterating until infinity. Using this axiomatic characterization, we can identify the scaling limit of Ising interfaces.

Theorem 8 (Benoist-Hongler)

Given a Jordan domain D , the collection of leftmost Ising loops in the Ising model on the discretization D^ε converges in law to CLE_3 as $\varepsilon \rightarrow 0$.

Remark 9. We briefly describe the topology on loop collections in which this convergence in law occurs. We define a metric on loop collections in which the distance between two loops is given by an infimum over supremum distance over reparametrizations (in other words, find a way to travel along both loops simultaneously so that the maximum distance over the travel time is minimized). Then, the distance between two loop collections I and J is given by the following idea: pair up a subset of the loops in I and a subset of the loops in J , and take the maximum over the diameters of the unmatched loops in either I or J , along with the maximum distance between the paired loops. Essentially, if this collection distance goes to 0, then every loop in our collection is either negligibly small or close up to reparameterization to our target collection.

The main idea behind the proof of this result is to consider the outermost Ising loops and show that they converge to CLE_3 (and then iterate). To do so, we follow an elaborate exploration scheme in which we **first discover outermost FK loops from the boundary** and then discover Ising loops conditioned on these outermost FK loops. To do so, we will need to introduce the **continuous dipolar explorer** (CDE) and **free arc ensemble** (FAE) in a domain and show how to glue together Ising loops from these objects. In Section 2, we define a variant of SLE needed to define the CDE and FAE, which we do in Section 3. In Section 4, we demonstrate a recursive discovery of outermost Ising loops through FK loops, which we use to prove the desired convergence in Section 5.

2 Driven Loewner chains and convergence of explorers

We first define $\text{SLE}(\kappa, \rho_1, \rho_2)$, a variant of SLE_κ in which we have “force points”. Heuristically, force points are points on the boundary of the domain that will encourage the SLE-like curve to turn left or right more. We will define this curve in the domain \mathbb{H} , which will give by conformal invariance the definition in any domain with two specified boundary points.

Definition 10

Recall that SLE_κ is defined by a Loewner chain with driving function $W_t = \sqrt{\kappa}B_t$ for a Brownian motion B_t . We define an **SLE(κ, ρ)** in the same way, except we use for our driving function W_t the solution to the following system of SDEs:

$$\begin{aligned} dW_t &= \sqrt{\kappa}dB_t + \frac{\rho}{W_t - O_t}dt, \\ dO_t &= \frac{2}{O_t - W_t}dt, \end{aligned}$$

where we impose some initial condition $O_t = x$ for $x \geq 0$ called the **force point**.

Thus, the driving function (which we may think of as a particle moving along the real axis) now evolves similarly to a Brownian motion, except there is another “force particle” to its right that repels or attracts it and is repelled by the driving function particle. This makes the resulting Loewner chain tend to turn left or right, depending on the value of ρ . Thus, we can think of κ as controlling the speed of the Brownian motion and thus roughness of the resulting path, while ρ controls the strength and direction of the effect from the other forcing particle. For $\rho > -2$, it turns out that there is a solution to this system of SDEs for all $t \geq 0$ such that $W_t \leq O_t$ for all t and W_t is instantaneously reflected off of O_t whenever they collide, though the stochastic calculus details are beyond the scope of this survey.

We can introduce a second force point to the left of the driving function as well. This gives the process $\text{SLE}(\kappa, \rho_1, \rho_2)$, whose driving function W_t is given by the solution to the following system of SDEs:

$$\begin{aligned} dW_t &= \sqrt{\kappa}dB_t + \left(\frac{\rho_1}{W_t - O_t^L} + \frac{\rho_2}{W_t - O_t^R} \right) dt, \\ dO_t^L &= \frac{2}{O_t^L - W_t}dt, \\ dO_t^R &= \frac{2}{O_t^R - W_t}dt, \end{aligned}$$

where we impose initial conditions $O_t^L = x^L \leq 0$ and $O_t^R = x^R \geq 0$ for the force points O^L and O^R . Similar properties and heuristic interpretations from our discussion of $\text{SLE}(\kappa, \rho)$ hold for $\text{SLE}(\kappa, \rho_1, \rho_2)$, and in fact we can relate the two kinds of processes when $\rho_1 + \rho_2 + 6 = \kappa$ (they agree in law, if we choose our force points appropriately in each case, until hitting the boundary). This gives some intuition for the parameters we pick next:


Definition 11

The **continuous dipolar explorer** (CDE) in \mathbb{H} with boundary points is defined as $\text{SLE}(3, -1.5, -1.5)$ with initial conditions 0^\pm for the force points. For a simply connected domain D , the CDE with boundary points (a, b) is a continuous simple path from a to b given by the image of the CDE in \mathbb{H} under a conformal map sending $(0, \infty) \mapsto (a, b)$.

One property of the CDE is that it can hit the boundary of the domain, but it will never hit itself. Hitting the boundary of the domain can be interpreted as the driving function “catching up” to one of

the force points that it repels. However, the specific parameters of the process ensure that the driving function is instantaneously reflected off of the force point (essentially, the hull is instantly filled in much like in SLE). Some calculations involving Bessel processes give this property as well as the fact that the resulting Loewner chain will give a simple path, but we omit the details. The CDE also satisfies a Markov property like SLE: conditioned on the path up to a certain time, the rest of the path is distributed as a CDE in the slit domain, with force points given by the two images of the starting point under the conformal map, as in Fig. 4.

Figure 4: A sketch of the Markov property for the continuous dipolar explorer.



Recall that for the critical Ising model with **positive** boundary conditions, we defined Ising loops as closed interfaces between positive and negative spins. We now consider an analogue for the critical Ising model with **free** boundary conditions in which we draw a path, called an **exploration**, between two marked boundary points. This is a simple path on the dual graph between two given dual boundary points with the property that the left of every edge is either $+1$ or outside of the domain and the right of every edge is either -1 or outside of the domain. (When we hit the boundary, we continue along it until we can re-enter the domain.) We can analogously define the leftmost and rightmost explorations and note that every exploration lies between them. (With free boundary conditions, specifying the start and end points allows us to unambiguously construct the explorations by always following a process like Remark 4, except the first and last step which are forced to go from ∂D to D .)

Since the FK-Ising Markov property involves FK clusters with free boundary conditions, we need to be able to work with Ising explorations for free boundaries, as well as Ising loops for positive boundaries.

Theorem 12 (Benoist, Duminil-Copin, Hongler)

In a simply connected domain D with marked boundary points u and v , any Ising exploration (for free boundary conditions) from u to v in a discretization of D converges in law to the CDE in D from u to v .

Note that the CDE bouncing off the boundary corresponds to the exploration needing to take an edge along the boundary due to there being, say a bunch of -1 spins on the left boundary. The scaling limit, like the Ising exploration, is still a simple path. We again do not delve into the details of the topology of convergence, though this is the same topology that we have discussed in lecture for, say, the convergence of the harmonic explorer to SLE_4 .

We briefly sketch the ideas behind this proof:

Proof sketch. First, a previous result of Hongler and Kytölä [4] shows that the scaling limit of the leftmost explorer agrees in law with CDE up to the first hitting time of the boundary. The argument there makes use of complex analysis and stochastic calculus, looking at Loewner chains on a strip of the complex plane and then computing quantities related to a local martingale – specifically, a certain term of the form $d\langle X, X \rangle - 3dt$ must vanish, which is where the $\kappa = 3$ characterization comes from.

Some technical bounding shows that the discrete explorations are tight, thus have a subsequential limit in law. Specifically, **crossing probabilities** are the probabilities that if we look at two disjoint arcs on the boundary of our Ising domain, there is a path of $+$ s that connect the two arcs. These probabilities converge in the scaling limit (and in fact it can be shown using this argument that using strong or weak neighbors does not affect the crossing probability in the limit), and verifying tightness requires checking a “compactness” argument that for any avoidable region (“topological rectangle”) that we haven’t already hit, the probability that (for any mesh size) we do in fact avoid that region is bounded from below by a positive constant. A similar check of “avoiding weird things in the limit” is done to ensure that if we hit the boundary in the limit, there are actually corresponding hitting times on the ϵ meshes, so that we can be consistent with “closing off the hull” like we did in ordinary SLE.

Now we can proceed to the proof. In a discrete exploration between u and v , at any time before we hit v , we can consider the slit domain obtained by cutting the edges along the dual path we have taken so far and taking the connected component of it containing v . (This is like Fig. 4, but imagining that we do so on a mesh.) Denote the leftmost and rightmost boundary points hit by the path in the slit domain by L and R , and conformally map the slit domain into \mathbb{H} with u, v, L, R mapping to $0, \infty, O^L, O^R$. (Remember that 0 and ∞ are our starting and ending points, and L and R start off very close to u (that is, 0 .) Then if W is the driving process for the corresponding Loewner chain, then we only need to check that W, O^L, O^R satisfy the system of SDEs along with $O^L \leq W \leq O^R$, Hölder continuity, and instantaneous reflection – one can show that these properties **uniquely characterize CDE**.


Again by Hongler and Kytölä’s result, in the subsequential limit, the law of the leftmost explorer comes from the $SLE(3, -1.5, -1.5)$ process until the first hitting time. Thus, the SDE holds so long as $O^L < W < O^R$, since this (along with the domain Markov property) corresponds to when the path in the hull has not hit the boundary yet. That W is always between O^L and O^R is clear, because the current point on the path is always between the leftmost and rightmost points hit on the boundary. Hölder continuity follows from the same result of Hongler and Kytölä (because of tightness). Finally, instantaneous reflection is a property of Loewner chains with continuous driving functions. This shows the result for the leftmost explorer, and we can get this for all explorers by showing that the leftmost and rightmost explorers agree on a dense subset of edges because they both must converge to the same limiting object, and we can **check crossing probabilities** along various points on the paths that they travel to see that the two explorers cannot diverge with nonzero probability. \square

3 The CDE and FAE

Now that we have identified the scaling limit of Ising explorations between two fixed boundary points in a domain (the CDE), we can identify the joint scaling limit of Ising explorations between all pairs of boundary points. This is an object called the **free arc ensemble** (FAE), consisting of a union of non-intersecting chordal paths between boundary points. This section will be dedicated to describing the construction of the FAE. We will give heuristics for why the scaling limit of all Ising explorations is the FAE as we will have constructed, though we will skip the formal proof as it consists of technical details that do not add much heuristically.

The first observation is that given points u, v, w on the boundary of the domain, **the CDE between u and v and the CDE between u and w should agree** until the first time that it hits the boundary point and separates v and w , i.e. the first time it hits the boundary arc between v and w . This may seem surprising (since the two explorers are aiming for different boundary points), but here we can use the scaling limit of explorers to CDE. Indeed, on a discretized version of the domain, the leftmost Ising explorer between u and v and the leftmost Ising explorer between u and w will agree until they hit the boundary arc between v and w , after which they will separate.

Figure 5: Starting the construction of the FAE from the first few CDEs.



Thus, conditioned on the CDE between u and v , we can obtain the CDE between u and w by coupling their driving functions until the CDE hits the boundary arc between v and w . At this point, depending on which point v or w gets mapped to ∞ in \mathbb{H} , the Loewner chain hull will swallow up different regions of the domain. (For example, in Fig. 5 above, if the first hitting point were left of v instead of between v and w , then both CDEs must continue in the same direction because the swallowed hull would be identical.) This accounts for the divergence of the two CDE paths at this point, and the rest of the CDE path from u to w is determined by the SDE solution in the partially-swallowed domain.

With this property, we will start to build up our set of arcs. First, we can simultaneously construct all CDEs from a fixed boundary vertex u to ∂D . Take a countable dense subset $\{v_1, v_2, \dots\}$ of ∂D . Then by the previous paragraph, we can inductively construct the CDEs from u to v_k for all k : in the union of the CDEs from u to v_i for $i < k$ there is a unique arc separating v_k from the rest of u, v_1, \dots, v_{k-1} . We follow the CDE until that arc has been traversed, and then use the construction from the previous paragraph to finish the exploration to v_k . (Alternatively, our set of arcs only branches when we hit a boundary point, at which point we turn towards our new point, repeatedly until there are no arcs left to follow.) Finally, by density we can reach any point on ∂D by taking approximation by our dense set v_k and looking at

the limiting path.

Note that every arc comes with an orientation, as the Ising explorer has $+1$ on the left and -1 on the right, so a given arc must always be traversed in the same direction. (In particular, the path from a to b is not the same as the path from b to a .) So suppose now that we want to construct a path between two points a and b not in our countably dense set – see Fig. 6 for reference. It is easy to see that the arcs between boundary points separating a and b , highlighted in blue, come with a linear ordering (based on how close they are to a), and there exists a order-preserving bijection between these arcs and \mathbb{Z} . Thus, since we know the direction of traversal of these separating arcs, all we need is to figure out how to traverse the part between consecutive separating arcs. Notice that (for example) in our figure, the middle arc is traversed upward, while the other two are traversed downward (these must alternate; otherwise we would have crossed between the separating arcs). Thus, to get from a to b , we will trace out all three of those blue arcs, and we will also follow the largest arc possible in each “vertical slice” that disconnects the two components of the boundaries – in other words, get as far away from the side of the slice we’re on while still being contained within the blue arcs. (This is marked in light green.)

Figure 6: Describing the separating arcs between two points in the next step of the FAE.

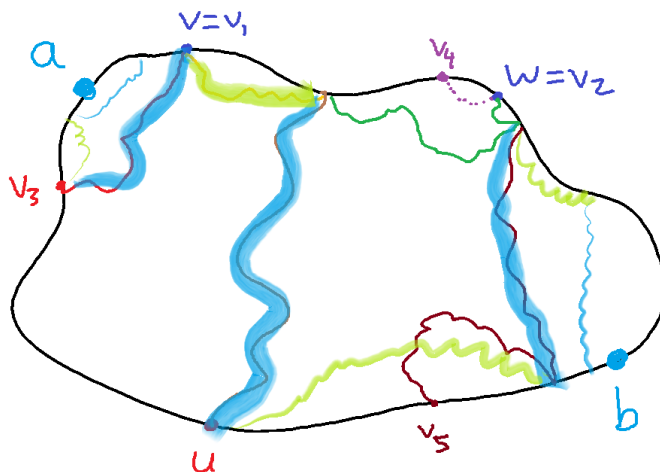
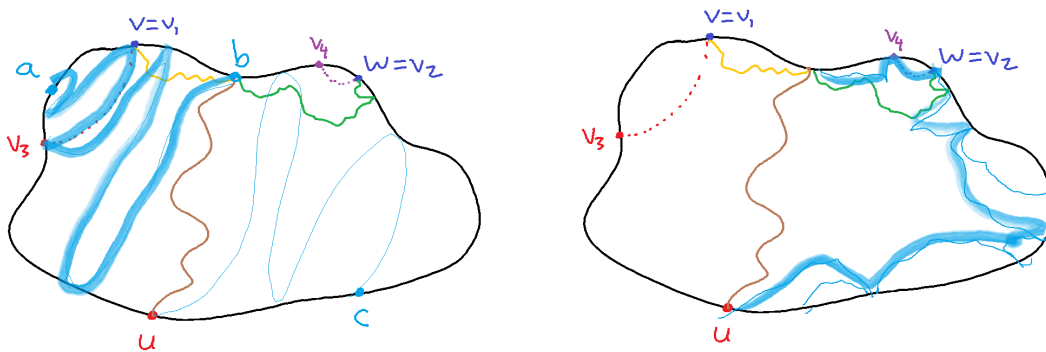


Figure 7: Special cases when points are on a separating arc – the paths γ are shown in highlighted blue.



There are now just a few more specific cases to check. The above argument works as long as we’re not on the boundaries of the separating arcs; if exactly one of our points is on an arc but the other is not, as in the left side of Fig. 7, we can pick some arbitrary non-arc-endpoint on the other end (point c

here), and follow the subset of the arc from c to a that starts at b . Also, if both of our points are on the same arc but the arc leads in the right way, we just follow that arc. However, since each arc comes with an orientation, we cannot backtrack arcs, so we must describe how to traverse arcs to “undo” an arc traversal, i.e. go from v to u if $u \rightarrow v$ is an arc. We do this as shown in the right side of Fig. 7: there will be both an arc on the left of the arc and also on the right side of the arc, and in both sides we basically take the largest arcs that separate the two sides of the boundaries while staying on one side of our arc.

This concludes the description of the free arc ensemble, which is the union of all arcs taken above. Note here that our construction above was guided by the properties of the explorations on the discretized domain, and all the things said about paths between boundary vertices hold when the CDE is replaced with leftmost Ising explorations. It should not be surprising, then, that the set of all Ising explorations converges to the FAE as we have constructed.

4 Recursive discovery of FK and Ising loops

Now, we are ready to describe our exploration process of FK and Ising loops. Recall that our goal here is to characterize the scaling limit of the outermost leftmost Ising loops as CLE_3 , using the axiomatic characterization of CLE with conformal invariance, Markov property, and local finiteness. We will use the important fact that each **Ising loop touches an FK loop** (or more specifically, is within $\frac{\varepsilon}{4}$ of one), because FK loops are the interface between the primal and dual open edges in the FK-Ising bond percolation. (Specifically, we already know that Ising loops are nested within an FK loop, and FK loops are always nested inside a cycle of vertices of the same cluster in the original model, or inside a loop of edges in the dual model. Therefore, either at least one of the vertices outside the Ising loop are connected to the outer FK loop cluster, meaning the FK loop does touch the Ising loop, or all vertices are disconnected, meaning that there is a further green/dual loop inside our FK loop and thus another FK loop right inside, so we aren’t dealing with outermost FK loops.) The key idea is to do this in reverse – given an FK loop, we will try to find all leftmost Ising loops that touch it.


In the following discussion, all Ising loops are leftmost. Suppose that we could find all Ising loops within an FK loop that touch its boundary. Note then that unless some Ising loop surrounds the FK loop, these Ising loops must all be outermost (crossings between loops aren’t allowed, and the Ising loop is minimal distance from its touching FK loop at at least one point). Once we have found these loops, it remains to look within FK loops that are outside of the outermost Ising loops we have found already for more outermost Ising loops.

This suggests the following algorithm, illustrated in Fig. 8:

1. Initialize L to be the set of all outermost FK loops.
2. Start a new set of FK loops L' . For each loop ℓ in L , find all Ising loops that touch the boundary of ℓ . Add to L' the outermost (with respect to the interior of ℓ rather than the whole domain) FK loops that are outside of the Ising loops discovered in ℓ .
3. Set L to be L' and repeat step 2 until no further iteration is possible.

To see why this algorithm will find all outermost Ising loops, note that all outermost Ising loops touch some FK loop. Thus, we should iterate through all FK loops and find the Ising loops inside that

Figure 8: Exploring the set of outermost Ising loops. On the left, we consider just one of the outermost FK loops (but in exploration we consider all of them; they're disjoint because of vertex-disjointness of the lattice FK loops). Within the yellow shaded region, we find the Ising loops that touch it, and we no longer need to check the blue regions. However, in the (yellow \setminus blue) regions, there will be some level 2 FK loops – any Ising loops contained in those will be nested in multiple FK loops but still be outermost Ising loops, so we must do this exact exploration process again with the level 2 loops (in red). All blue loops are outermost, but the numbers indicate the level of their touching FK loop.



touch its boundary. Once we have all outermost FK loops with respect to some region, we only have to look inside those FK loops for more FK loops; by definition all other FK loops in the region are inside the outermost FK loops we have found. Within an FK loop, once we have found the Ising loops touching the boundary, we don't need to search within those loops for more FK loops, as any Ising loops we find inside will not be outermost. Putting all of this together implies the correctness of the algorithm, since eventually we will run out of FK loops to search through and we only skip searching through FK loops contained inside Ising loops we found.

Now, we discuss the "implementation" of this algorithm – that is, how we actually "find all the loops." First, note that in step 2, the FK loops we add to L' (in red) are the outermost (boundary touching) FK loops in a region that consists of a union of simply connected domains (in yellow), because the Ising loops whose interiors we have removed from consideration all touch the boundary of ℓ . Furthermore, this region always has a strong path of $+1$ boundary conditions, because both the (leftmost) outermost Ising loops and the outside of the FK loop do, so this invariant property holds. Thus, if we show that FK loops have a conformally invariant scaling limit then the part where we must find all outermost (boundary touching) FK loops can be understood for any region by a conformal map if it is understood for a single region.

Indeed, FK loops do have a conformally invariant scaling limit; in fact, the scaling limit of the FK loops are non-crossing, and the interiors of the discrete outermost FK loops converge in law to the interiors of the limiting outermost FK loops. All of these are essentially just technical considerations to ensure that our algorithm can in a sense also be passed to the scaling limit. We will skip the details of this technical lemma (which is also relegated to the Appendix in the original paper), but we will mention a few of the key points:

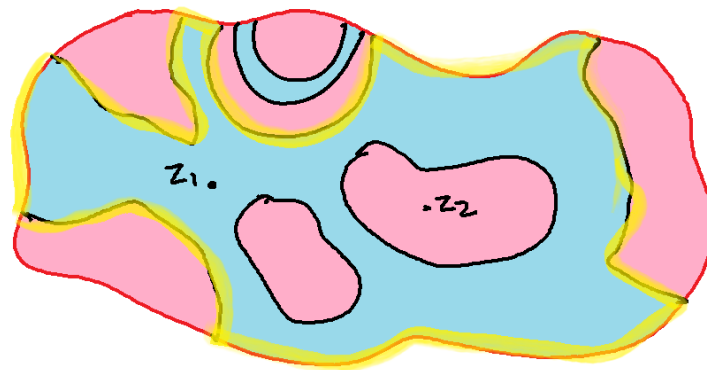
- Construct an **FK exploration** between boundary points which follows primal and dual clusters on the left and right (similarly to how an Ising exploration follows ± 1 spins on its two sides).
- Discrete FK explorations actually converge to $\text{SLE}(\frac{16}{3}, -\frac{2}{3})$, using similar arguments to [1].

- Because of the topology of convergence, we can ignore all loops of diameter less than δ (this leaves only finitely many).
- An exploration path that connects two boundary paths (more specifically, two sequences of lattice paths whose endpoints converge to the appropriate boundary points) splits up the domain into regions, and we can bound (uniformly in ϵ) from above the number of regions with diameter at least δ .
- We can look at whether the regions have wired or mixed boundary conditions; FK loops can be found within each mixed boundary region by finding the interfaces along which we change from primal to dual. Furthermore, all interfaces will continue to break up our regions into smaller subregions with either wired or free boundary conditions.
- There are only finitely many larger-than- δ wired boundary condition subregions within each region. This process decreases the maximum diameter of the unexplored region, and then we can continue FK exploring in the remaining regions until we've eliminated all possible larger-than- δ regions with high probability.

The other part of the implementation is to find the Ising loops touching the boundary of an FK loop. Recall that the FK loops we are considering all have strong $+1$ boundary conditions, so by the FK-Ising domain Markov property, the interior of the FK loops are distributed as a free boundary condition Ising model. We resample this Ising model after discovering the outermost FK loops. Consider the set of all leftmost Ising **explorations** between boundary points of this FK loop enclosed region. The key observation here is that there is a **deterministic way to piece together these exploration arcs** into Ising loops that touch the boundary of the FK loop.

Indeed, fix an interior dual vertex z . For each edge b on the boundary of this region, consider the Ising exploration arc $A(z, b)$ that separates z from b and is closest to z (these arcs are ordered because of the deterministic exploration method), where we define $A(z, b) = b$ if there are no such arcs. Then the union of $A(z, b)$ over all b gives us an innermost region $R(z)$ enclosing z – see Fig. 9.

Figure 9: In black, interfaces between $+$ and $-$ spins. In red and blue, regions of $+$ and $-$ spins, respectively. Notice that black lines cannot touch each other. The yellow highlighted boundaries enclose $R(z)$.



Now consider the sign of the spin on the “inside” edge of each arc, i.e. the spins on each arc closer to z . It is clear that all of these spins are the same, and that if the spin is -1 then $R(z)$ is an outermost

leftmost Ising loop that touches the boundary of the FK loop, like in the figure. (The alternative would only occur if z were instead in one of the regions outside of this $R(z)$ – then instead flipping the boundary conditions gives us an Ising loop.) Furthermore, given an outermost leftmost Ising loop that touches the boundary, it is easy to see that $R(z)$ for any z inside the loop will be this Ising loop. Hence, we have exhibited a deterministic construction of all Ising loops touching the boundary of the FK loop, finishing the implementation of this part of our algorithm.

What remains is to deduce from the algorithm/exploration process the scaling limit of the outermost leftmost Ising loops. Before doing so, let us summarize the results we will need from this section. To sample a set of outermost leftmost Ising loops, we can do the following:

1. Find the outermost FK loops.
2. Resample the spins inside those FK loops by the Markov property as a free boundary condition Ising model.
3. Use all leftmost Ising explorations between points on the boundary of these FK loop surrounded regions to deterministically obtain all outermost Ising loops touching the boundary of the regions.
4. Resample the FK-Ising model inside the FK loop region but outside of the Ising loops we found, noting that we can do so because we have preserved a strong path of $+1$ for the boundary of this region, which consists of simply connected components.
5. Repeat the previous steps with each resampled FK-Ising model.

5 Proof of the desired convergence

We are finally ready to conclude that the scaling limit of Ising interfaces is nested CLE_3 . First we will show that the scaling limit of **outermost** Ising interfaces is CLE_3 , using the axiomatic characterization of CLE.

The first step is to show that the outermost Ising loops have a conformally invariant scaling limit. This follows from the conformal invariance of the scaling limit of the outermost FK loops and our exploration process. Indeed, consider the Ising loops we discover on the first pass through the steps at the end of the previous section. These Ising loops in the scaling limit are a deterministic function of the FAE in each FK loop. Since we know that the FAE is conformally invariant and the outermost FK loops are conformally invariant, it follows that the outermost leftmost Ising loops arising from the first pass of our algorithm above are conformally invariant.

But this means that on the second pass through the steps, we apply the steps to a random set of regions whose law is conformally invariant, and within each region we produce a set of Ising loop scaling limits whose law in that region is conformally invariant. This implies that the second pass through the steps yields a set of loops that have a conformally invariant law. The same argument applies by induction to the entire process, which by a technical lemma covers all outermost Ising loops in the scaling limit.

Remark 13. *For some additional details, the full set of outermost Ising loops is the countable union of the sets we obtain through inductively performing this procedure. We can show that with high probability, there are only a*

finite number of Ising loops larger than some δ diameter because of the **FK arm exponents** [3], and then by using a sign-flip symmetry argument as described in the discussion of Fig. 9, the unexplored regions have supremum distance bounded by the Ising loop distances, so they must go to 0 and thus convergence in our topology does exist.)

Furthermore, the scaling limit of outermost Ising loops satisfy local finiteness, and they are almost surely simple, disjoint, and do not touch the boundary. (We skip the proofs of these technical lemmas.) A simplified view of the reason that the loops are simple and disjoint is because they are formed by gluing together FAE arcs in disjoint regions, and within each region FAE arcs do not intersect. Additionally, the fact that Ising interfaces converge to SLE_3 tells us that we should never hit the boundary (because $\kappa < 4$). Finally, the domain Markov property in the continuum follows from the analogue of the domain Markov property in the discretization. Indeed, it suffices to check that given a domain D and a fixed subdomain D' , conditioned on the leftmost outermost Ising loops that intersect D' , the remaining leftmost outermost Ising loops are distributed as the leftmost outermost Ising loops in the region remaining when we remove D' and the interiors of the conditioned Ising loops from D , with **positive** boundary conditions. Indeed, the sampling process for the Ising model (and the local interactions) tell us that we do have an Ising model in the remaining domain. To check the boundary conditions, any vertex that is a vertex on some conditioned Ising loop (touching D') will have $+1$ spin since outermost Ising loops have $+1$ spin on their outside. Also, any vertex on the boundary between D' and $D \setminus D'$ not part of one of those loops must also have $+1$ spin, as if they had -1 spin they would be contained inside some Ising loop and would have been removed during the conditioning. Thus from the point of view of $D \setminus D' \setminus \{\text{additional loops}\}$, the entire boundary is $+1$, as desired.

We have thus shown the outermost Ising loops have a scaling limit that satisfies the CLE axiomatic characterization, so must be CLE_κ for some κ . But note that since we have glued together Ising loops from arcs in the FAE in subdomains of D , the dimension of the Ising loops is equal to the dimension of the arcs in the FAE. These arcs in turn are made up of CDEs, which are variants of SLE_3 whose Loewner driving functions have a Brownian motion running at speed 3. Thus, Ising loops have dimension $1 + \frac{3}{8}$, which implies that $\kappa = 3$ and the outermost Ising loops have scaling limit CLE_3 .

To finish, it remains to show that all of the Ising loops give nested CLE_3 in the scaling limit. But note that inside each Ising loop in the critical Ising model, we have by the Markov property of the Ising model and the definition of an Ising loop that the spins inside are distributed as an Ising model on the region with negative boundary conditions. Since the spins are symmetric, the outermost Ising loops within each of these regions is distributed as a CLE_3 . Iterating this process (and flipping the sign at each nesting), we find that the scaling limit of all Ising loops is given exactly by nested CLE_3 , as desired.

References

- [1] Stéphane Benoist, Hugo Duminil-Copin, and Clément Hongler. Conformal invariance of crossing probabilities for the ising model with free boundary conditions. *Annales de l'Institut Henri Poincaré, Probabilités et Statistiques*, 52, 10 2014.

- [2] Stéphane Benoist and Clément Hongler. The scaling limit of critical ising interfaces is cle(3). *The Annals of Probability*, 47, 04 2016.
- [3] Dmitry Chelkak, Hugo Duminil-Copin, and Clément Hongler. Crossing probabilities in topological rectangles for the critical planar fk -ising model. 2013.
- [4] Clément Hongler and Kalle Kytölä. Ising interfaces and free boundary conditions. 2011.
- [5] Lars Onsager. Crystal statistics. i. a two-dimensional model with an order-disorder transition. *Phys. Rev.*, 65:117–149, Feb 1944.
- [6] Scott Sheffield and Wendelin Werner. Conformal loop ensembles: The markovian characterization and the loop-soup construction, 2011.