The limit of random graphs: a survey of Benjamini--Schramm convergence

Zhengbo Qu
2000011611@stu.pku.edu.cn

Contents

1. Introduction

One of the most powerful concepts in modern mathematics is the notion of a limit. If we want to study a large sequence of objects, we can often learn a lot by understanding what happens when we zoom out and look at the big picture. When it comes to random graphs, the situation is quite similar, i.e., we can learn a lot by understanding the limit of a sequence of random graphs. In this article, I will survey a seminal paper by Benjamini and Schramm[1], which introduced the Benjamini–Schramm convergence (or local convergence) of random graphs.

We organize the paper as follows:
In Section 2, I will introduce the notion of the Benjamini–Schramm limits of (maybe random) graphs, and state the main theorem of the paper, which says that such limit graphs are almost surely recurrent. (To be honest, I never thought that we can have such a strong result about the limit of random graphs! And this is the reason why I chose this topic). Furthermore, I will prove prove some nontrivial properties and give some examples to illustrate the concept of Benjamini–Schramm convergence since it may not be so easy to understand at first glance.
In Section 3, I will give the detailed proof of the main theorem, which is the most difficult and technical part in this paper. I divided the proof into three parts: the circle packing theory part, the case reduced to triangulation part, and the proof of an important lemma (Lemma 3.2.2). We can see from this section that how the techniques of circle packing and random walks will be used, which really impressed me when I first read the proof.
In Section 4, I will give some remarks on the proof (for example, how its relevant to the topics we have learned at courses), and try to offer some ideas and intuition behind the difficult techniques in the proof.
In Section 5, I will summarize the main results and discuss some further topics and future directions.

2. Benjamini–Schramm limits of graphs

2.1. Space of rooted graphs

However, the general idea of limits of graphs is not so easy to define since there are so much information in a graph. In order to keep track of the local structure of a graph, we need to define a notion of convergence that captures the local structure of a graph, i.e., a *root* and a *rooted graph*.

Definition 1: Let G be a graph and $o \in V(G)$.
1. A *rooted graph* is a pair (G, o), the vertex o is called the *root* of the rooted graph (G, o).
2. Two rooted graphs are *isomorphic* if there is an isomorphism between them that maps the root to the root.

In this article, we will only consider locally finite connected graphs with countable vertices. And we have the following spaces with an appropriate metric.

Definition 2:
1. Let $\mathcal{G}_\bullet$ be the space of isomorphism classes of locally finite connected rooted graphs.
2. Given a rooted graph $(G, o) \in \mathcal{G}_\bullet$, the finite graph $B_G(o, R)$ is defined as the subgraph of G induced by the set of vertices at distance at most R from o.
3. We can define a metric on $\mathcal{G}_\bullet$ to capture the local structure of a rooted graph. For any two rooted graphs $(G, o), (G', o') \in \mathcal{G}_\bullet$, we define the *local distance (or Benjamini–Schramm distance)* between them as:

$$d_{BS}((G, o), (G', o')) := 2^{-R},$$

where R is the largest integer such that $B_G(o, R)$ is isomorphic to $B_{G'}(o', R)$.

And we will see that $\mathcal{G}_\bullet$ together with the Benjamini–Schramm distance has some nice properties.

Proposition 1: d_{BS} is a metric on $\mathcal{G}_\bullet$.

The proof of this proposition needs *Kőnig's infinity lemma* and is not so trivial, but it is not much relevant to our main theorem. Due to the length limitaion, we will not give the proof here.
A first consequence of the metric properties is that $\mathcal{G}_\bullet$ is actually a Polish space.

Theorem 2.1.1: $\mathcal{G}_\bullet$ is a Polish space, i.e., it is a separable and complete metric space together with the metric d_{BS}.

Proof: The separability is easy to verify, we only prove that $\mathcal{G}_\bullet$ is complete. Let $\{(G_n, o_n)\}_{n=1}^{\infty} \subseteq \mathcal{G}_\bullet$ be a Cauchy sequence. For each $r \in \mathbb{N}_0$, there exists $n_r \in \mathbb{N}$ such that $B_{G_n}(o_n, r) \cong B_{G_{n_r}}\left(o_{n_r}, r\right), \forall n \geq n_r$. Note that the sequence n_r can be chosen to be strictly increasing with respect to r, and we can define a rooted graph (G, o) as follows: for simplicity, we denote $B_{G_{n_r}}\left(o_{n_r}, r\right)$ by (H_r, v_r) for all $r \in \mathbb{N}_0$. Then by the monotonicity of n_r, we have:

$$B_{H_s}(v_s, r) \cong B_{H_r}(v_r, r), \forall s \geq r.$$

Thus, these graphs are compatible as rooted graphs, i.e., for $s \geq r$, we can assume that they have the same set of vertices and edges when restricted to r-ball. Let $(G, o) \in \mathcal{G}_\bullet$ be the graph with $V(G) = \cup V(H_r)$ and $E(G) = \cup E(H_r)$, and o be the vertex $o_0 \in V(H_0)$. Then we have:

$$B_G(o, r) \cong B_{H_r}(v_r, r) \cong B_{G_{n_r}}\left(o_{n_r}, r\right) \cong B_{G_n}(o_n, r), \forall n \geq n_r,$$

which implies that $d_{BS}((G, o), (G_n, o_n)) \leq 2^{-r}, \forall n \geq n_r$, and this completes the proof. □

2.2. Benjamini–Schramm convergence

Since $\mathcal{G}_\bullet$ is a Polish space, with the help of the Prokhorov's theorem, we can talk about convergence in distribution of random variables $\{X_n\}_{n=1}^{\infty}$ taking values in $\mathcal{G}_\bullet$.

Definition 3: Given a sequence of random variables $\{X_n\}_{n=1}^{\infty}$ taking values in $\mathcal{G}_\bullet$, we say that $\{X_n\}$ converges in distribution to a random variable X, and denote it by $X_n \xrightarrow{d} X$, if for every bounded function $f : \mathcal{G}_\bullet \to \mathbb{R}$, we have that $\mathbb{E}[f(X_n)] \to \mathbb{E}[f(X)]$ as $n \to \infty$.

In the following, we will focus on the situation where each X_n is a *finite* random rooted random graph (G_n, o_n) such that given G_n, the root is chosen uniformly at random from $V(G_n)$. We will see this setting is quite natural and justifies the definition of the *Benjamini–Schramm convergence*.

Definition 4: Let $\{G_n\}_{n=1}^{\infty}$ be a sequence of (possibly random) finite graphs. We say that G_n *converges in the Benjamini–Schramm sense (or converges locally)* to a random rooted graph $(G, o) \in \mathcal{G}_\bullet$, and denote it by $G_n \xrightarrow{BS} (G, o)$, if for every $r \in \mathbb{N}$,

$$B_{G_n}(o_n, r) \xrightarrow{d} B_G(o, r),$$

where o_n is a uniformly random vertex in $V(G_n)$.
We say that (G, o) is the *Benjamini–Schramm limit (or local limit)* of the sequence $\{G_n\}_{n=1}^{\infty}$.

Remark:
1. Note that this is equivalent to say, for every $r \in \mathbb{N}_0$ and every finite rooted graph $(H, v) \in \mathcal{G}_\bullet$, we have:

$$\mathbb{P}\left[B_{G_n}(o_n, r) \cong B_H(v, r)\right] \to \mathbb{P}[B_G(o, r) \cong B_H(v, r)], \text{as } n \to \infty.$$

2. Roughly speaking, when taking Benjamini–Schramm convergence, we are looking at the local structure of the graph, we choose the r-balls at random for each G_n and asking what will them eventually look like.
3. We clarify that whether G_n is deterministic or random does not matter, as long as the root is chosen uniformly at random from $V(G_n)$.

2.3. Some Examples

Let's consider some simple examples to illustrate the concept of Benjamini–Schramm convergence. Since the purpose of this subsection is to offer some intuition, we will not give the detailed proof here and only provide some ideas.

Example 1: *(Convergence of cycles)*

Let $\{C_n\}_{n=3}^{\infty}$ be a sequence of cycles with n vertices. Then we have $C_n \xrightarrow{BS} (\mathbb{Z}, o)$, where $\mathbb{Z}$ is the set of integers and o is the vertex 0.

Example 2: *(Convergence of grids)*

Let $\{G_n\}_{n=1}^{\infty}$ be a sequence of $n \times n$-grids with n^2 vertices. Then we have $G_n \xrightarrow{BS} (\mathbb{Z}^2, o)$, where o is the vertex $(0, 0)$.

Note that in both examples above, the roots of the limit do not matter, since we can choose any other vertices as the roots and we get the same elemen in $\mathcal{G}_\bullet$ up to isomophism.

We now consider an example which is not so trivial.

Example 3: *(Canopy tree)*

let $\{G_n\}_{n=1}^{\infty}$ be a sequence of truncated binary trees with n levels, i.e., G_n is the unique binary tree with n levels and $(2^n - 1)$ vertices. However, the Benjamini–Schramm limit of this sequence is not the infinite binary tree, but the *canopy tree*. We will not give the precise definition of the canopy tree here, but we can offer a picture of it and some intuition. We now describe the canopy tree as the Benjamini–Schramm limit as follows:

Let T_C be the canopy tree, for each $k \in \mathbb{N}_0$, we choose a vertex x_k of "height" k in $V(T_C)$ as in the picture below (here height means the distance from the leaves). Then if we have another vertex y_k of the same height k, we can see from the picture that (T_C, x_k) and (T_C, y_k) are isomorphic as rooted graphs. Therefore, it is well defined to say $(T_C, x_k) \in \mathcal{G}_\bullet$. Let (G, o) be the Benjamini–Schramm limit of the sequence $\{G_n\}_{n=1}^{\infty}$ as above. It can be shown that:

$$\mathbb{P}[(G, o) \cong (T_C, x_k)] = \frac{1}{2^{k+1}}.$$

Furthermore, unlike the infinite binary tree, it can be shown that it is recurrent.

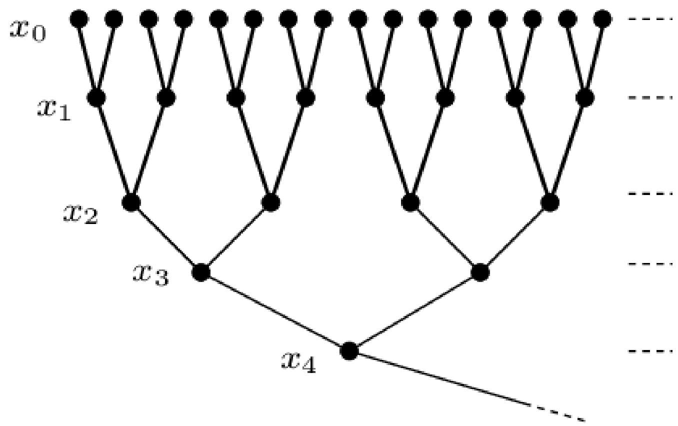

Figure 1: A picture of part of the canopy tree, which is the Benjamini–Schramm limit of the sequence of truncated binary trees.

2.4. Statement of the main theorem

However, the Benjamini–Schramm limit does not always exist. For example, we can consider the sequence $\{G_n\}$ of star-shaped graphs with n leaves emananating from the center. Thus, to get a better result, we need to impose some restrictions.

Definition 5: Let $M \in \mathbb{N}$, define $\mathcal{G}_M$ to be the set of locally finite, connected rooted graphs (up to isomorphism) with degrees bounded by M, clearly we have $\mathcal{G}_M \subseteq \mathcal{G}_\bullet$.

Clearly, $\mathcal{G}_M$ is a closed subset of $\mathcal{G}_\bullet$, and together with the metric d_{BS}, it is not hard to verify that $\mathcal{G}_M$ is a complete and totally bounded space and we omit the detailed proof here. From the knowledge of analysis, we conclude that:

Proposition 2: The metric space $\mathcal{G}_M$ is compact.

It is not had to see, by the compactness of $\mathcal{G}_M$, that if $\{G_n\}_{n=1}^{\infty}$ is a sequence of random finite graphs with $(G_n, o_n) \in \mathcal{G}_M$ almost surely, $M < \infty$, then there is always a subsequence that converges in the Benjamini–Schramm sense to a random rooted graph $(G, o) \in \mathcal{G}_M$.

Now, we can state the main result of Benjamini and Schramm in their seminal paper[1], which is also the main focus of this article.

Theorem 2.4.1: Let $M < \infty$, and let $(G, o) \in \mathcal{G}_\bullet$ be a Benjamini–Schramm limit of a sequence of rooted random finite planar graphs $\{G_n\}_{n=1}^{\infty}$ with degrees bounded by M (almost surely). Then (G, o) is almost surely recurrent.

3. The proof of the main theorem

The proof we will present here basically follows the original proof by Benjamini and Schramm[1], and also the lecture notes by Asaf Nachmias[2].

3.1. Circle packing

We alwasy want to reduce the problem to the case where it is easier to handle. Luckily, we have the *circle packing theorem*, which gives us a way to approximate a finite planar graph by a circle packing. We first give a statement of the circle packing theorem (*without proof*).

Theorem 3.1.1: (*Circle packing theorem*) Given any connected simple planar graph G there is a circle packing P in the plane whose tangency graph is (isomorphic to) G.

Moreover, we need the following ring lemma.

Lemma 3.1.2: (*Ring lemma*) For every integer $N \geq 3$, there exists $B > 0$, such that if C_0 is a circle or radius r_0 and is completely surrounded by N circles of radius $r_i, i = 1, ..., N$, then $r_0 / r_i \leq B$, for every $i = 1, ..., N$.

The reason that things become easier in the situation of circle packing is that we actually have a recurrence-transience dichotomy result on the circle packing. We first define the accumulation point in a circle packing, which is quite similar to the definition in the metric space case.

Definition 6: We say that a point is an *accumulation point* of a circle packing if in every neighborhood of the point there are infinitely many circles of the packing.

And we have the following theorem.

Theorem 3.1.3: If a circle packing has no accumulation point, then its tangency graph is recurrent.

We can assume without loss of generality that the circle packing has infinitely many circles, by the uniformization theorem of (discrete) circle packing, we can conclude that the circle packing is supported either on the plane or the Poincaré disk, while the circle packing has accumulation points in the latter case. So we only need the following theorem to complete our proof.

Theorem 3.1.4: Let R be either the plane or the disk and G the tangency graph of a circle packing on R. Then R is plane iff G is recurrent.

Remark: This theorem is quite natural from the perspective of rough isometry, but I'd like to give a direct proof of one direction of this theorem: if G is transient, then R is the disk, which is the case we need.

Proof: We first have a criterion for the transience of G: G is transient iff there exists a vertex x_0 and a constant A, such that for all $\varphi \in C(R)$, we have $\varphi(x_0)^2 \leq A \sum_{(x,y)\in E(G)} (\varphi(x) - \varphi(y))^2$. And we have a criterion for the transience of a Riemannian manifold R: R is transient iff there exists $U \subseteq R, K > 0$,such that for all $\varphi \in C^2_{c(R)}$, we have $\left(\int_U \varphi\right)^2 \leq K \int_R \|\nabla\varphi\|^2$.
Now, fix the x_0 in the first criterion and let $\varphi \in C^2_c(R)$. For every $x \in V(G)$, let $U(x)$ be the union of all faces of G that meet x. We then take some f such that $f(x) = \frac{1}{|U(x)|} \int_{U(x)} \varphi(x) \,\mathrm{d}x, \forall x \in V(G)$. In particular, we can assume that $|U(x_0)| = 1$, and we have:

$$\int_{U(x_0)} \varphi(x) \,\mathrm{d}x = f(x_0) \leq \left(A \sum_{(x,y)\in E(G)} (\varphi(x) - \varphi(y))^2\right)^{\frac{1}{2}}.$$

Moreover, let $x, y \in V(G)$ and $D_{x,y}$ be the smallest disk containing $U(x) \cup U(y)$. We have:

$$\begin{aligned}
(|U(x)| \times |U(y)|)(f(y) - f(x)) &= \int_{U(x)} \int_{U(y)} \varphi(a) - \varphi(b) \,\mathrm{d}a \,\mathrm{d}b \\
&\leq \int\int_{D_{x,y}\times D_{x,y}} |\varphi(a) - \varphi(b)| \,\mathrm{d}a \,\mathrm{d}b \\
&\leq \int\int_{D_{x,y}\times D_{x,y}} \int_0^1 \|\nabla\varphi(a + t(b-a))\| \,\mathrm{d}t \,\mathrm{d}a \,\mathrm{d}b \\
&\leq \operatorname{diam}(D_{x,y})|D_{x,y}| \int_{D_{x,y}} \|\nabla\varphi(v)\| \,\mathrm{d}v.
\end{aligned}$$

By the ring lemma, $|U(x)|, |U(y)|, |D_{x,y}|$ are all within constant factor, so we have the following inequality:

$$(f(x) - f(y))^2 \leq A' \frac{1}{|D_{x,y}|} \left(\int_{D_{x,y}} \|\nabla\varphi(v)\| \,\mathrm{d}v\right)^2 \leq A' \int_{D_{x,y}} \|\nabla\varphi(v)\| \,\mathrm{d}v$$

for some contant A'. Combining theses inequalities, we have:

$$\left(\int_{U(x_0)} \varphi(x)\,\mathrm{d}x\right)^2 \leq A \sum_{(x,y)\in E(G)} (\varphi(x)-\varphi(y))^2 \leq K \int_R \|\nabla\varphi(v)\|\,\mathrm{d}v.$$

While the last inequality also comes from the ring lemma. By the second criterion, R is transient. Since the disk is transient and the plane is recurrent, we conclude that the disk is the only possibility. □

3.2. The triangulation case and the main theorem

And we have the theorem holds in the triangulation case.

Proposition 3: (*The triangulation case*) Let $M < \infty$, and let (T, o) be a Benjamini-Schramm limit of a sequence of rooted random finite triangulations of the sphere $\{T_j\}_{j=1}^{\infty}$ with degrees bounded by M (almost surely). Then (T, o) is almost surely recurrent.

Proof: Assume that T is a.s an infinite graph. Let P^j be the circle packing of T_j, and o_j be the random root of T_j. Without loss of generality, we can assume that $P^j_{o_j}$, the disk corresponding to the root o_j, is the unit disk $B(0, 1)$. We want to find an appropriate limit of $\{P^j\}$, a circle packing P with tangency graph T.

There is a unique triangle t_j in T_j whose vertices correspond to the three disks of P^j which intersect the unbounded part in $\mathbb{R}^2 \setminus P^j$. And for vertex $v \in V(T_j) \setminus t_j$, the corresponding circle is completely surrounded by other circles, Since the vertex degree is bounded by M, we can apply Lemma 3.1.2 to show that: for any circle with (combinatorial) distance at most d from o_j, provided that the distance between o_j and t_j is larger than d, there exists a constant $c = c(d, M)$, such that its radius is at most c. Since $|V_j| \to \infty$ as $j \to \infty$ and the vertex degree is bounded, when d is fixed, we can always suppose that the o_j and its surrounding circles with distance at most d are far away from t_j. And we conclude that: there exists a constant $c = c(d)$, such that for all circles with distance at most d from o_j, their radii are in $[1/c, c]$.

By compactness and passing to a subsequence, we can assume that there is a limit random circle packing P with tangency graph T. Now, we know from Theorem 3.1.3 that if P has no accumulation point, then T is recurrent. Actually, it will also be fine if P has only one accumulation point. Let p be this accumulation point, let G_1 be the subgraph of T spanned by vertices in $B(p, 1)$, G_2 be the subgraph of T spanned by vertices in the complement of $B(p, 1)$, since G_2 is recurrent and the bouondary separating G_1 and G_2 in T is finite, we only need to show that G_1 is recurrent. In this situation, the accumulation point p is actually not so important, since we can alwasy invert in the circle to make p go to infinity, then it is not an accumulation point anymore, so G_1 is also recurrent. Thus, we only need to show the following fact to complete our proof.

Proposition 4: With probability 1, there is at most one accumulation point in $\mathbb{R}^2$ of the circle packing P.

However, this proposition is not so easy to prove, we will need some technical lemmas to prove it, the next section will be devoted to the proof of this proposition. Now, suppose that this proposition is true, then we can conclude that T is recurrent, which completes the proof. □

We are now ready to prove the main theorem.

Proof: (of Theorem 2.4.1)

With the help of Proposition 3, we only need to show that all case can be reduced to the triangulation case, which can be done by the following lemma.

Lemma 3.2.1: Let $\{G_n\}$ as in the statement of Theorem 2.4.1. Then there is a constant c (which only depends on M), such that for all $j = 1, 2, \ldots$, there exists a triangulation T_j of the sphere with degrees bounded by cM, which contains a subgraph isomorphic to G_j, and such that $|V(T_j)| \leq c|V(G_j)|$.

Proof: Let f be a face of G, first suppose that there is no edges between non-consecutive vertices of f, then we can use the *zig-zag* construction to get a triangulation (see the left of Figure 2). If there are edges between non-consecutive vertices of f, then we can add a cycle with the same number of vertices in the interior of f and connect it in a similar zig-zag fashion to get a triangulation (see the right of Figure 2).

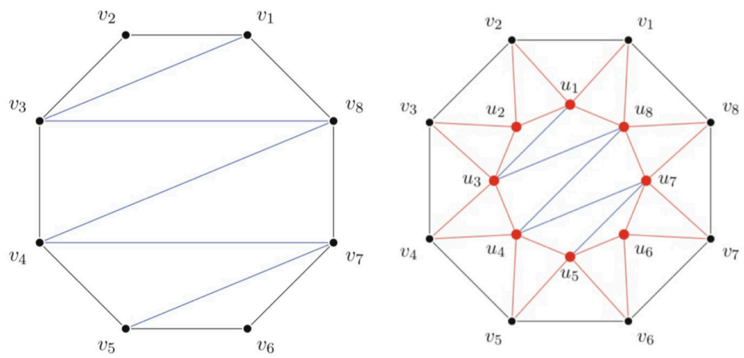

Figure 2: The zig-zag construction picture from [2]

In both cases, for each vertex in each face, there are at most two edges added, so the maximal degree of resulting graph is at most $3M$, and the number of vertices in this new graph is at most M times the number of vertices in the original graph, which completes the proof.
□

By Proposition 3, a subsequential limit of $\{(T_j)\}$ is almost surely recurrent. Let o_j be a vertex chosen uniformly form $V(T_j)$, we have $\mathbb{P}(o_j \in V(G_j)) \geq \frac{1}{c}$. By Rayleigh's monotonicity principle, we conclude that G is almost surely recurrent.

□

3.3. The magic lemma and the proof of Proposition 4

To prove the main theorem, we need to do some preparation.

Definition 7:
1. Suppose that $C \subseteq \mathbb{R}^2$ is finite. For each $w \in C$, define the *isolation radius* of w in C as:
$$\rho_w = \min\{|v - w| : v \in C \setminus \{w\}\}.$$
2. Given $\delta \in (0, 1)$, $s \geq 2$ and $w \in C$, we say that w is (δ, s)*-supported* in C if:
$$\inf_{p \in \mathbb{R}^2} \left|C \cap B(w, \delta^{-1}\rho_w) \setminus B(p, \delta\rho_w)\right| \geq s.$$

In other word, w is (δ, s)*-supported* if in the disk of radius $\delta^{-1}\rho_w$ around w there are at least s points of C outside any given disk of radius $\delta\rho_w$.

The proof of Theorem 2.4.1 is based on the following lemma, which is called the *magic lemma.*

Lemma 3.2.2: There exists $A > 0$ such that for every $\delta \in \left(0, \frac{1}{2}\right)$, every finite $C \subseteq \mathbb{R}^2$ and every $s \geq 2$, the number of (δ, s)-supported points in C is at most
$$\frac{A|C|\delta^{-2}\ln(\delta^{-1})}{s}.$$

Proof: The proof of the magic lemma is a little bit technical, and we will do this step by step.

First, let $k \geq 3$ be an integer. Let $\mathfrak{G}_0$ be a tiling of $\mathbb{R}^2$ by 1×1 squares, rooted at some point $p \in \mathbb{R}^2$, and for every $n \in \mathbb{Z}$, let $\mathfrak{G}_n$ be the tiling of $\mathbb{R}^2$ by $k^n \times k^n$ squares, such that each square in $\mathfrak{G}_n$ is divided into k^2 squares in $\mathfrak{G}_{n-1}$. Without loss of generality, we may choose p appropriately so that none of the points of C lie on the edges of the squares in $\mathfrak{G}_0$.

We say that a square $S \in \mathfrak{G}_n$ is s*-supported* if for every smaller $S' \in \mathfrak{G}_{n-1}$, we have that
$$|C \cap (S \setminus S')| \geq s.$$

Actually, we will see that the number of such squares is bounded.

Proposition 5: For any $s \geq 2$, the total number of s-supported squares in $\mathfrak{G} = \cup_{n \in \mathbb{Z}} \mathfrak{G}_n$, is at most $2|C|/s$.

Proof: We define a *"flow"* on $\mathfrak{G}$, $f : \mathfrak{G} \times \mathfrak{G} \to \mathbb{R}$, as follows:
$$f(S', S) = \begin{cases} \min\left(\frac{s}{2}, |C \cap S'|\right), \text{ if } S' \subseteq S, S' \in \mathfrak{G}_n, S \in \mathfrak{G}_{n+1}; \\ -f(S, S'), \text{ if } S \subseteq S', S \in \mathfrak{G}_n, S' \in \mathfrak{G}_{n+1}; \\ 0, \text{ otherwise.} \end{cases}$$

We can have three observations:

1. $$\sum_{S' \in \mathfrak{G}_b} \sum_{S \in \mathfrak{G}_{b+1}} f(S', S) \geq 0.$$
2. $$\sum_{S' \in \mathfrak{G}} f(S', S) \geq 0, \forall S \in \mathfrak{G}.$$

This is because if there exists some S' such that $f(S',S)=s/2$, then $\sum_{S'\in\mathfrak{G},f(S',S)>0} f(S',S)\geq s/2=\sum_{S'\in\mathfrak{G},f(S',S)<0} f(S',S)$. If none of the S' satisfies this, then $\sum_{S'\in\mathfrak{G}} f(S',S)\geq 0$ automatically.

3. If S is a s-supported square, then

$$\sum_{S'\in\mathfrak{G}} f(S',S)\geq \frac{s}{2}.$$

   We only need to show that $\sum_{S'\subsetneq S} f(S',S)\geq s$. This is because if there are at least 2 squares $S'\subsetneq S$ such that $f(S',S)=s/2$, then the above inequality automatically holds true. If the number is at most 1, then we can directly apply the definition of s-supported to get this lower bound.

Next, we choose $a\in\mathbb{Z}$ to be small enough, such that no squares in $\mathfrak{G}_a$ contains more than 2 points in C. Since $s\geq 2$, there are no s-supported squares in $\bigcup_{n\leq a}\mathfrak{G}_n$, and we have:

$$\sum_{S'\in\mathfrak{G}_a}\sum_{S\in\mathfrak{G}_{a+1}} f(S',S)=|C|.$$

Together with the observations, we have:

$$\begin{aligned}\sum_{n=a+1}^{b}\sum_{S\in\mathfrak{G}_n}\sum_{S'\in\mathfrak{G}} f(S',S) &= \sum_{n=a+1}^{b}\sum_{S\in\mathfrak{G}_n}\left(\sum_{S'\in\mathfrak{G}_{n-1}} f(S',S)+\sum_{S'\in\mathfrak{G}_{n+1}} f(S',S)\right)\\ &= \sum_{S'\in\mathfrak{G}_a}\sum_{S\in\mathfrak{G}_{a+1}} f(S',S)+\sum_{S'\in\mathfrak{G}_{b+1}}\sum_{S\in\mathfrak{G}_b} f(S',S)\\ &\leq |C|\end{aligned}$$

Now, let $b\to+\infty$, using the observations above, we know that the number of s-supported squares is at most $2|C|/s$. □

However, it still needs some work to see that the number of (δ,s)-supported points is also bounded.

Let $k:=\lceil 20\delta^{-2}\rceil$, $\beta\sim\mathrm{Unif}([0,\ln(k)])$. Let $\mathfrak{G}_0$ be a tiling with side lenght e^{β} that bases at the origin, i.e.,

$$\mathfrak{G}_0=\left\{e^{\beta}([x,x+1]\times[y,y+1]):x,y\in\mathbb{Z}\right\}.$$

Then the squares in $\mathfrak{G}_n$ have side length $k^{n+1}e^{\beta}$. Suppose we have the tiling for $\mathfrak{G}_n$, we can define the tiling for $\mathfrak{G}_{n+1}$ by choosing base uniformly at one of the k^2 squares in $\mathfrak{G}_n$. Since dilating C will not affect our expected conclusion, we may assume that $\rho_w\geq k$ for every $w\in C$.

Next, we want to transform the statement in Proposition 5, which is about squares, to the statement in Lemma 3.2.2, which is about circle. To do this, we first introduce the following definition.

Definition 8: A point $w\in C$ is called a *city* in a square $S\in\mathfrak{G}$, if:
1. the side length of S is in the range $[4\delta^{-1}\rho_w,5\delta^{-1}\rho_w]$;
2. the distance from w to the center of S is at most $\delta^{-1}\rho_w$.

We can see from the following proposition that, in some sense, a (δ, s)-supported city $w \in C$ represents a s-supported square $S \in \mathfrak{G}$.

Proposition 6: If $w \in C$ is a (δ, s)-supported city in $S \in \mathfrak{G}_n$, then S is s-supported.

Proof: First we have $B(w, \delta^{-1}\rho_2) \subseteq S$. Moreover, any little square $S' \in \mathfrak{G}_{n-1}$ with $S' \subseteq S$ has side length at most $\frac{\delta^2}{20} \times \frac{5\rho_w}{\delta} = \frac{\delta\rho_w}{4}$, and is contained in a disk of radius $\delta\rho_w$. Thus, for every such S', there exists a point p such that $S' \subseteq B(p, \delta\rho_w)$. So we have:

$$|C \cap (S \setminus S')| \geq \left|C \cap \left(B(w, \delta^{-1}\rho_w) \setminus B(p, \delta\rho_w)\right)\right| \geq s,$$

which holds for every $S' \in \mathfrak{G}_{n-1}$ with $S' \subseteq S$. □

We only need to estimate the number of cities now.

Proposition 7: The probability that any given $w \in C$ is a city is $O(\ln^{-1}(\delta^{-1}))$.

Proof: We can see from the definition of a city that we have the following conditions:

1. There exists $n \in \mathbb{Z}$ such that $k^n e^\beta \in [4\delta^{-1}\rho_w, 5\delta^{-1}\rho_w]$, i.e., $\beta + n \ln k \in [\ln(\delta^{-1}\rho_w) + \ln 4, \ln(\delta^{-1}\rho_w + \ln 5)]$. Since $\beta \sim \mathrm{Unif}([0, \ln k])$, easy computation shows that the probability is about $\left(\frac{\ln(\frac{5}{4})}{\ln(k)}\right) = O(\ln^{-1}(\delta^{-1}))$.
2. The second condition will not affect our desired conclusion, since we have a positive probability to make w close to the center of S by choosing the base of $\mathfrak{G}_n$ over $\mathfrak{G}_{n-1}$, which is independent of δ and β.

□

Now let N be the number of (δ, s)-supported points in C, then the expected number of (δ, s)-supported cities is $cN \ln^{-1}(\delta^{-1})$, where c is a constant independent of N and δ. Also note that given a square S, the number of cities in S is at most $c'\delta^{-2}$. By Proposition 6, we have that the expected number of s-supported squares is at least $\frac{c}{c'}N\delta^2 \ln^{-1}(\delta^{-1})$, then by the estimation in Proposition 5, we have that:

$$N \leq \frac{A|C|\delta^{-2}\ln(\delta^{-1})}{s},$$

completing the proof. □

Now we turn to prove Proposition 4, which is the last step to complete the proof of the main theorem.

Proof: (of Proposition 4) Without loss of generality, we may assume that for every $j \in \mathbb{N}$, the center of the disk corresponding to the root o_j is the origin 0 and has $\rho_{o_j} = 1$ in the set C^j consisting of the centers of the disks of the circle packing P^j.

Now suppose that with positive probability, there are at least two accumulation points in the circle packing P. Then there exists $\delta \in (0, 1)$ small enough and $\varepsilon > 0$, such that with probability at least ε, there are two accumulation points $p_1, p_2 \in \mathbb{R}^2$ which is contained in $B(0, 1/\delta)$ and $|p_1 - p_2| \geq 3\delta$.

When such two points exist, let s be large enough, there are infinitely many j such that o_j is (δ, s)-supported in C^j, since the little circle $B(p, \delta)$ can cover only one of the accumulation points.
Since o_j is chosen uniformly, it means that a proportion of C^j is always (δ, s)-supported for s large enough. However, by the magic lemma, the number of (δ, s)-supported points is at most $K|C^j|\delta^{-2}\frac{\ln(\delta^{-1})}{s}$, which is a contradiction. □

4. Some remarks and explanations about the proof

As we have seen, the proof of the main theorem is quite technical and really takes some dirty work, so I'd like to give some remarks and explanations about the ideas behind the proof. I think the most important parts in the proof are the idea of using circle packing, and the magic lemma, so I will focus on these two parts.

4.1. Why we need circle packing

I think that the meaning of using circle packing is to find a canonical way to present a graph, which is given by Theorem 3.1.1. In this way, we can take advantage of the properties of the circle packing and better understand the structure of the graph. For example, we can easily see the quasi–isometry structure behind the graph, thus we can reduce the problem to the recurrence–transience dichotomy of the Riemannn surfaces, which is a well–known result.

(P.S. And this might be the part in this article that is most relevant to our courses, which is really about the random walks and recurrence.)

4.2. Where does the magic lemma come from

The magic lemma, which deserves its name, is the most technical part in the proof. To understand where it comes from, we first need to examine the main theorem Theorem 2.4.1 again.
The main theorem actually says that when you have finite graphs (with bounded degree) and take limits with respect to the uniformly chosen roots. Since taking uniform roots is like a kind of "averaging" process, we expect the resulting limit to have some nice and uniform properties, then we have seen that this limit is actually almost surely recurrent.
So the problem in proving the theorem is to show how nice the limit is, or in other words, how not bad the limit is. The most bad points in the limit are the *accumulation points*, which is not so easy to analyze. So we turn our attention to points that are bad but not so bad, which are the (δ, s)*-supported points*. So here comes the magic lemma, which says that the number of (δ, s)-supported points is not so much, and this is the key to show that the limit is nice enough.

The proof of the magic lemma also needs some explanation. The definition of (δ, s)-supported points is a little difficult to work directly with, so we need to find a way to transform it to a more manageable form, which is the reason why we introduce the concept of s-supported squares and cities, the squares are much easier to handle than the disks. Moreover, we have defined *flows* to help us to estimate the number of s-supported squares, and we can see that the flow is a very powerful tool in this proof. Recall what s-supported squares mean, it means that it actually contains a lot of points of C, and the way of containing these points is not so centered at a certain small square, and we cab use the flow to describe this property.
More specifically, let me explain why we need the term $\min\left(\frac{s}{2}, |C \cap S'|\right)$ in the definition of flow f. The idea is that we want to control the effect brought by the squares which contain a lot of points of

C, and the term $s/2$ cannot be chosen too big or too small (otherwise we lose the information of the term min and will not get the desired inequality), so $s/2$ is a moderate choice (but not the only choice, for example, take $s/3$ will also work).

5. Conclusion and further reading

In this article, we have surveyed the concept of Benjamini–Schramm convergence and the main theorem introduced in [1] , which says that the Benjamini–Schramm limit of a sequence of finite graphs with bounded degrees is almost surely recurrent. We have also given a detailed proof of the main theorem, which is based on the [1] and [2, Chapter 5].

Furthermore, the bounded degree assumption can be relaxed to the assumption that the degree of the root has an exponential tail [3], which says that the unbounded part can be controlled by an exponential function, i.e., $\mathbb{P}(\deg(\rho) \geq k) \leq Ce^{-\beta k}$ for some positive constants C and β. As a corollary, this implies that the uniform infinite planar triangulation and quadrangulation (UIPT and UIPQ) are almost surely recurrent, which is obtained by taking the limit of a uniform random triangulation (or quadrangulation) on n vertices.

Moreover, let G be a finite graph and consider the simple random walk $(X_t)_{t\geq 0}$ on it, where X_0 is a uniform random vertex of G. Let $\phi(T, G)$ be the probability that $X_T \neq X_0$ for all $t = 1, 2, ..., T$, i.e, the probability that the simple random walk avoids the starting point for the first T steps. We can then define for $D \geq 1$ that:

$$\phi_D(T) = \sup\{\phi(T, G) : G \text{ is planar with degrees bounded by D}\}.$$

The main theorem (Theorem 2.4.1) is equivalent to the statement that $\phi_D(T) \to 0$ as $T \to 0$ for any fixed D. A quite natural question is that what is the rate of decay of this funcion [1, Problem 1.3]. This question was also answered in [3], which says that for any fixed D and $T \geq 2$, $\phi_D(T) \leq C/\log(T)$ for a constant C.

However, these results involves much more techniques and detailed analysis, so I choose not to include it in this article. The interested readers can refer to [2, Chapter 6] and the original paper [3].

There are also some connections between this topic and other theories, such as the theory of SLE, Liouville quantum gravity and boundary theory. We again omit these topics in this article, but the interested readers can refer to [2, Chapter 8].

Bibliography

[1] I. Benjamini and O. Schramm, “Recurrence of Distributional Limits of Finite Planar Graphs,” *Electronic Journal of Probability*, vol. 6, no. none, pp. 1–13, Jan. 2001, doi: 10.1214/EJP.v6-96.

[2] A. Nachmias, *Planar Maps, Random Walks and Circle Packing: École d'Été de Probabilités de Saint-Flour XLVIII - 2018*, vol. 2243. in Lecture Notes in Mathematics, vol. 2243. Cham: Springer International Publishing, 2020. doi: 10.1007/978-3-030-27968-4.

[3] O. Gurel-Gurevich and A. Nachmias, “Recurrence of planar graph limits,” *Annals of Mathematics*, vol. 177, no. 2, pp. 761–781, Mar. 2013, doi: 10.4007/annals.2013.177.2.10.