Walking through the Brownian zoo, JFLG60
3rd-7th June 2019

The critical random transposition random walk

Christina Goldschmidt (Oxford)

joint work with Dominic Yeo (Technion until recently, now Oxford)

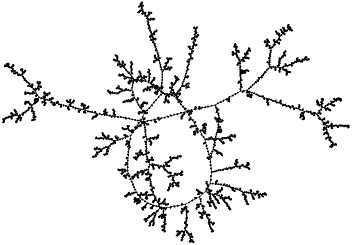

The random transposition random walk

Let $(P^n_t)_{t\geq 0}$ be a continuous-time random walk on the symmetric group $\mathfrak{S}_n$, which starts from the identity permutation and at rate $n/2$ composes the current state with an independent uniform random transposition. This is the random transposition random walk (RTRW).

(1)(2)(3)(4)(5)(6)(7)(8)

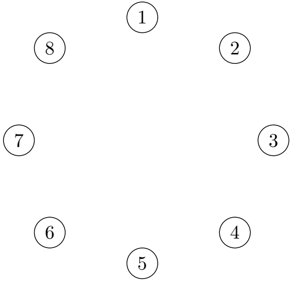

The random transposition random walk

Let $(P_t^n)_{t \geq 0}$ be a continuous-time random walk on the symmetric group $\mathfrak{S}_n$, which starts from the identity permutation and at rate $n/2$ composes the current state with an independent uniform random transposition. This is the random transposition random walk (RTRW).

$(1)(2)(3)(4)(5)(6)(7)(8) \circ (1, 2)$

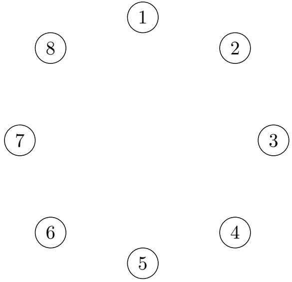

The random transposition random walk

Let $(P^n_t)_{t\geq 0}$ be a continuous-time random walk on the symmetric group $\mathfrak{S}_n$, which starts from the identity permutation and at rate $n/2$ composes the current state with an independent uniform random transposition. This is the random transposition random walk (RTRW).

$(1,2)(3)(4)(5)(6)(7)(8)$

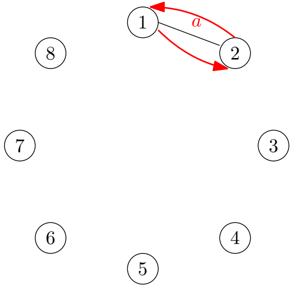

The random transposition random walk

Let $(P^n_t)_{t\geq 0}$ be a continuous-time random walk on the symmetric group $\mathfrak{S}_n$, which starts from the identity permutation and at rate $n/2$ composes the current state with an independent uniform random transposition. This is the random transposition random walk (RTRW).

$(1,2)(3)(4)(5)(6)(7)(8) \circ (3,4)$

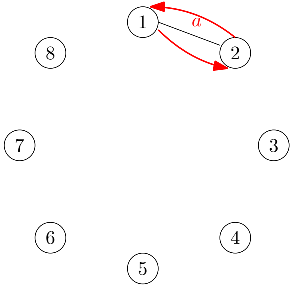

The random transposition random walk

Let $(P^n_t)_{t\geq 0}$ be a continuous-time random walk on the symmetric group $\mathfrak{S}_n$, which starts from the identity permutation and at rate $n/2$ composes the current state with an independent uniform random transposition. This is the random transposition random walk (RTRW).

$(1,2)(3,4)(5)(6)(7)(8)$

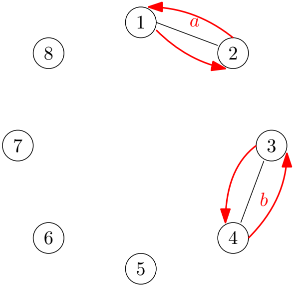

The random transposition random walk

Let $(P^n_t)_{t\geq 0}$ be a continuous-time random walk on the symmetric group $\mathfrak{S}_n$, which starts from the identity permutation and at rate $n/2$ composes the current state with an independent uniform random transposition. This is the random transposition random walk (RTRW).

$(1,2)(3,4)(5)(6)(7)(8) \circ (1,4)$

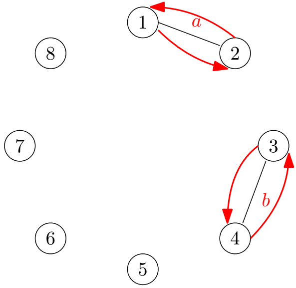

The random transposition random walk

Let $(P^n_t)_{t\geq 0}$ be a continuous-time random walk on the symmetric group $\mathfrak{S}_n$, which starts from the identity permutation and at rate $n/2$ composes the current state with an independent uniform random transposition. This is the random transposition random walk (RTRW).

$(1,2,4,3)(5)(6)(7)(8)$

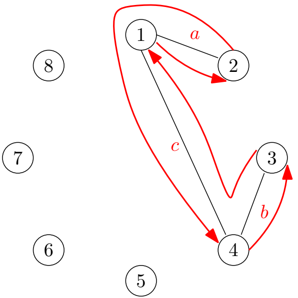

The random transposition random walk

Let $(P_t^n)_{t \geq 0}$ be a continuous-time random walk on the symmetric group $\mathfrak{S}_n$, which starts from the identity permutation and at rate $n/2$ composes the current state with an independent uniform random transposition. This is the random transposition random walk (RTRW).

$$(1, 2, 4, 3)(5)(6)(7)(8) \circ (6, 7)$$

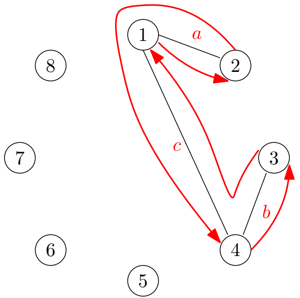

The random transposition random walk

Let $(P^n_t)_{t\geq 0}$ be a continuous-time random walk on the symmetric group $\mathfrak{S}_n$, which starts from the identity permutation and at rate $n/2$ composes the current state with an independent uniform random transposition. This is the random transposition random walk (RTRW).

$(1,2,4,3)(5)(6,7)(8)$

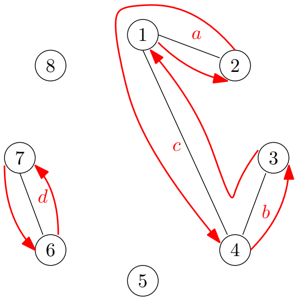

The random transposition random walk

Let $(P^n_t)_{t\geq 0}$ be a continuous-time random walk on the symmetric group $\mathfrak{S}_n$, which starts from the identity permutation and at rate $n/2$ composes the current state with an independent uniform random transposition. This is the random transposition random walk (RTRW).

$(1, 2, 4, 3)(5)(6, 7)(8) \circ (1, 3)$

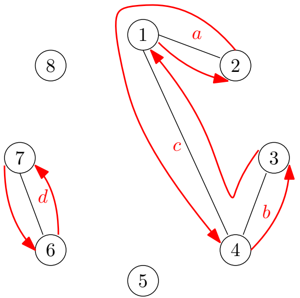

The random transposition random walk

Let $(P^n_t)_{t\geq 0}$ be a continuous-time random walk on the symmetric group $\mathfrak{S}_n$, which starts from the identity permutation and at rate $n/2$ composes the current state with an independent uniform random transposition. This is the random transposition random walk (RTRW).

$(1,2,4)(3)(5)(6,7)(8)$

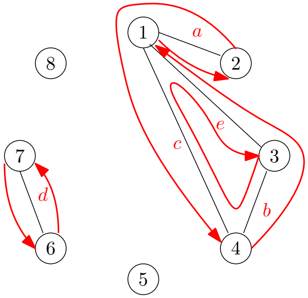

Cycle lengths

Our primary interest is going to be in the lengths of the cycles. For an arbitrary permutation π, write $C_1(\pi), C_2(\pi), \ldots$ for the cycle-lengths, written in decreasing order.

Cycle lengths

Our primary interest is going to be in the lengths of the cycles. For an arbitrary permutation π, write $C_1(\pi), C_2(\pi), \ldots$ for the cycle-lengths, written in decreasing order.

It's not hard to see that the stationary distribution is uniform on $\mathfrak{S}_n$. For a uniform permutation Π^n, we have

$$\frac{1}{n}(C_1(\Pi^n), C_2(\Pi^n), \ldots) \xrightarrow{d} \mathrm{PD}(0,1).$$

A celebrated result of Schramm says that $\mathrm{PD}(0,1)$ turns up long before the chain is mixed.

The phase transition

Theorem.

- If $t < 1$, $C_1(P_t^n) = O_{\mathbb{P}}(\log n)$.
- If $t = 1$, $C_1(P_t^n) = \Theta_{\mathbb{P}}(n^{2/3})$.
- (Schramm, 2005) If $t > 1$, there exists a random set $A_t^n \subseteq [n]$ such that $|A_t^n| \sim \theta(t)n$ and

$$\frac{1}{|A_t^n|}(C_1(P_t^n), C_2(P_t^n), \ldots) \xrightarrow{d} \mathrm{PD}(0,1),$$

  as $n \to \infty$, where $\theta(t)$ is the survival probability of a Poisson(t) Galton–Watson process i.e. the smallest non-negative solution to the equation $1 - \theta = e^{-t\theta}$.

The phase transition

Theorem.

- If $t < 1$, $C_1(P_t^n) = O_{\mathbb{P}}(\log n)$.
- If $t = 1$, $C_1(P_t^n) = \Theta_{\mathbb{P}}(n^{2/3})$.
- (Schramm, 2005) If $t > 1$, there exists a random set $A_t^n \subseteq [n]$ such that $|A_t^n| \sim \theta(t)n$ and

$$\frac{1}{|A_t^n|}(C_1(P_t^n), C_2(P_t^n), \ldots) \xrightarrow{d} \mathrm{PD}(0,1),$$

as $n \to \infty$, where $\theta(t)$ is the survival probability of a Poisson(t) Galton–Watson process i.e. the smallest non-negative solution to the equation $1 - \theta = e^{-t\theta}$.

(See also Berestycki (2011) for a simpler proof that there are giant cycles for $t > 1$.)

A coupling with the Erdős–Rényi random graph

A key point in the proof of this theorem is a coupling with the Erdős–Rényi random graph process.

Let G_t^n be the graph with vertex-set $[n]$ and with edges given by the set of pairs $\{i, j\}$ such that the transposition (i, j) has been applied in the construction of P_t^n. Then $G_t^n \sim \mathbb{G}(n, 1 - e^{-t/n})$ with $1 - e^{-t/n} \approx t/n$.

A coupling with the Erdős–Rényi random graph

Key fact: any cycle of the permutation is wholly contained within a component of the graph.

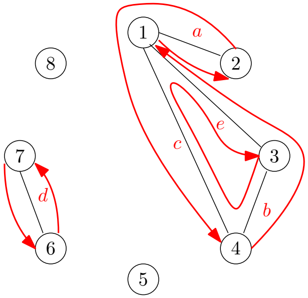

So the size of the largest component of G_t^n is an upper bound on $C_1(P_t^n)$. This easily gives the claimed results for $t < 1$ and $t = 1$.

The giant component

For $t > 1$, there is a giant component, with vertex-set A_t^n, of size $\sim \theta(t)n$. Necessarily, any giant cycles must be contained within it, because all other components have size $O_{\mathbb{P}}(\log n)$.

The giant component

For $t > 1$, there is a giant component, with vertex-set A_t^n, of size $\sim \theta(t)n$. Necessarily, any giant cycles must be contained within it, because all other components have size $O_{\mathbb{P}}(\log n)$.

Roughly speaking, the it is the dynamics on the giant component which produce the $\mathrm{PD}(0,1)$ relative cycle lengths.

The giant component

For $t > 1$, there is a giant component, with vertex-set A_t^n, of size $\sim \theta(t)n$. Necessarily, any giant cycles must be contained within it, because all other components have size $O_{\mathbb{P}}(\log n)$.

Roughly speaking, the it is the dynamics on the giant component which produce the $\mathrm{PD}(0, 1)$ relative cycle lengths.

Why? Whenever we add an edge inside the giant component, it selects two size-biased permutation cycles. If these cycles are distinct, adding the edge merges them. If they were the same to start with, the selected permutation cycle splits at a uniform point into two pieces.

The split-merge process

Let

$$\Delta = \left\{ \mathbf{x} = (x_1, x_2, \dots) : x_1 \geq x_2 \geq \dots \geq 0, \sum_{i \geq 1} x_i = 1 \right\}.$$

Consider the following (continuous) version of the split-merge dynamics on Δ. At rate $1/2$, take two independent size-biased picks from among the blocks:

- if the blocks are distinct, merge them;
- if they are the same, split the block into two blocks with relative sizes U and $1 - U$, where $U \sim \mathrm{U}[0, 1]$.

The split-merge process

Let

$$\Delta = \left\{ \mathbf{x} = (x_1, x_2, \ldots) : x_1 \geq x_2 \geq \ldots \geq 0, \sum_{i \geq 1} x_i = 1 \right\}.$$

Consider the following (continuous) version of the split-merge dynamics on Δ. At rate $1/2$, take two independent size-biased picks from among the blocks:

- if the blocks are distinct, merge them;
- if they are the same, split the block into two blocks with relative sizes U and $1 - U$, where $U \sim \mathrm{U}[0, 1]$.

Theorem. (Diaconis, Meyer-Wolf, Zeitouni & Zerner, 2004) PD$(0, 1)$ is the unique invariant distribution for the split-merge dynamics on Δ.

The critical window

Question (Schramm): what precisely happens in the critical window, where $t = 1 + \lambda n^{-1/3}$, $\lambda \in \mathbb{R}$?

The critical window

Question (Schramm): what precisely happens in the critical window, where $t = 1 + \lambda n^{-1/3}$, $\lambda \in \mathbb{R}$?

Aim for this talk: to give an answer, and to convince you that the solution lives (at least partly) in the Brownian zoo!

The critical window

One can also make sense of a split-merge process $(Y(t), t \geq 0)$ taking values in

$$\ell_2^{\downarrow} = \left\{ \mathbf{x} = (x_1, x_2, \ldots) : x_1 \geq x_2 \geq \ldots \geq 0, \sum_{i \geq 1} x_i^2 < \infty \right\}.$$

Informally,

- a pair of blocks of sizes x and y merge at rate xy
- a block of size x splits at rate $x^2/2$ into blocks of sizes xU and $x(1-U)$, where $U \sim \mathrm{U}[0,1]$.

After a jump, we re-sort the blocks into decreasing order of size.

The critical window

One can also make sense of a split-merge process $(Y(t), t \geq 0)$ taking values in

$$\ell_2^{\downarrow} = \left\{ \mathbf{x} = (x_1, x_2, \ldots) : x_1 \geq x_2 \geq \ldots \geq 0, \sum_{i \geq 1} x_i^2 < \infty \right\}.$$

Informally,

- a pair of blocks of sizes x and y merge at rate xy
- a block of size x splits at rate $x^2/2$ into blocks of sizes xU and $x(1-U)$, where $U \sim \mathrm{U}[0,1]$.

After a jump, we re-sort the blocks into decreasing order of size.

Without the splitting, this is the multiplicative coalescent. We extend a result of Aldous (1997) to show that the split-merge process possesses the Feller property.

The critical window

Theorem. (G. & Yeo, 2019+) There exists an eternal version $(Y^*(\lambda), \lambda \in \mathbb{R})$ of the split-merge process on $\ell_2^{\downarrow}$ such that, in the Skorokhod sense,

$$\left(n^{-2/3}\left(C_1(P^n_{1+\lambda n^{-1/3}}), C_2(P^n_{1+\lambda n^{-1/3}}), \ldots\right), \lambda \in \mathbb{R}\right) \xrightarrow{d} (Y^*(\lambda), \lambda \in \mathbb{R}).$$

Moreover, as $\lambda \to \infty$,

$$\frac{1}{2\lambda} Y^*(\lambda) \xrightarrow{d} \mathrm{PD}(0, 1)$$

for the product topology.

The critical window

Theorem. (G. & Yeo, 2019+) There exists an eternal version $(Y^*(\lambda), \lambda \in \mathbb{R})$ of the split-merge process on $\ell_2^{\downarrow}$ such that, in the Skorokhod sense,

$$\left(n^{-2/3}\left(C_1(P^n_{1+\lambda n^{-1/3}}), C_2(P^n_{1+\lambda n^{-1/3}}), \ldots\right), \lambda \in \mathbb{R}\right) \\ \xrightarrow{d} (Y^*(\lambda), \lambda \in \mathbb{R}).$$

Moreover, as $\lambda \to \infty$,

$$\frac{1}{2\lambda} Y^*(\lambda) \xrightarrow{d} \mathrm{PD}(0,1)$$

for the product topology.

We describe the entrance law $(\nu_\lambda)_{\lambda \in \mathbb{R}}$ such that $Y^*(\lambda) \sim \nu_\lambda$ below.

The critical random graph

The critical random graph is well-understood, so the main point in the proof of this theorem is to understand the relationship between the structure of the graph and that of the permutation.

The critical random graph

The critical random graph is well-understood, so the main point in the proof of this theorem is to understand the relationship between the structure of the graph and that of the permutation.

It is clearly sufficient to consider the different components separately.

The critical random graph

The critical random graph is well-understood, so the main point in the proof of this theorem is to understand the relationship between the structure of the graph and that of the permutation.

It is clearly sufficient to consider the different components separately.

Key facts:

- Conditionally on having given vertex-set and e edges, a component of the Erdős–Rényi random graph is uniform on set of connected graphs with those properties.
- Conditionally on the component, its edges arrived in a uniform random order.

Connected graphs and permutations

For an arbitrary finite connected graph G, let $\Pi(G)$ be the permutation obtained by applying the transpositions represented by the edges in a uniformly random order.

Connected graphs and permutations

For an arbitrary finite connected graph G, let $\Pi(G)$ be the permutation obtained by applying the transpositions represented by the edges in a uniformly random order.

The surplus of a connected graph G is defined to be $\#\text{edges} - \#\text{vertices} + 1$. The surplus is always non-negative, and is 0 for a tree. If G has surplus 2 or more, we call it complex.

Connected graphs and permutations

Tree

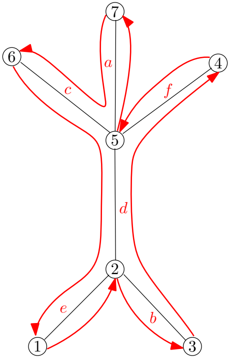

Lemma. (Dénes, 1959) A tree always gives rise to a single permutation cycle.

Connected graphs and permutations

Tree

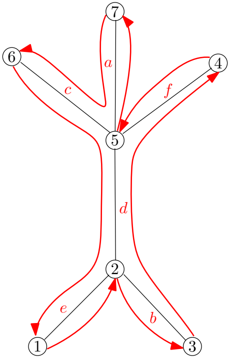

Lemma. (Dénes, 1959) A tree always gives rise to a single permutation cycle.

(Indeed, there is a bijection between minimal factorisations of a permutation cycle and labelled trees.)

Connected graphs and permutations

Graph-cycle

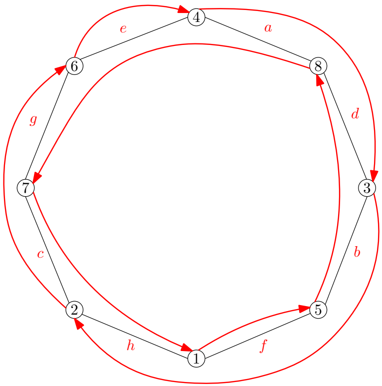

Lemma. If G_m is a graph-cycle of length m, then $\Pi(G_m)$ has precisely two cycles.

Connected graphs and permutations

Graph-cycle

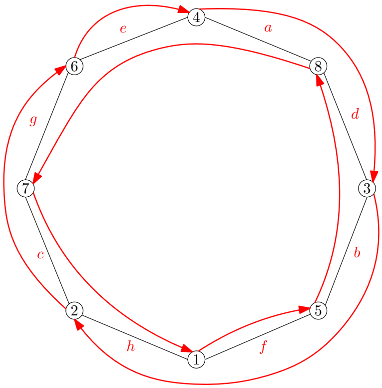

Lemma. If G_m is a graph-cycle of length m, then $\Pi(G_m)$ has precisely two cycles. Moreover,

$$\frac{1}{m}(C_1(\Pi(G_m)), C_2(\Pi(G_m))) \xrightarrow{p} \left(\tfrac{1}{2}, \tfrac{1}{2}\right), \quad \text{as } m \to \infty.$$

Connected graphs and permutations

Unicyclic component (surplus 1)

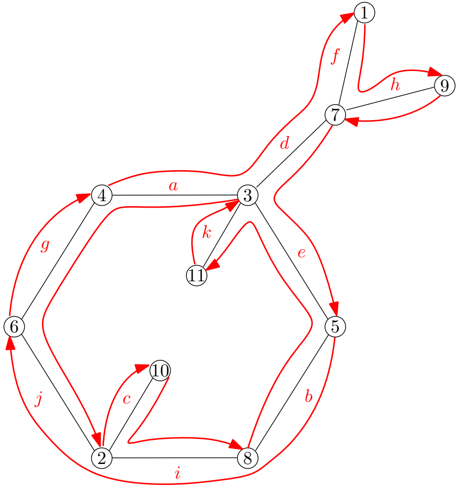

Connected graphs and permutations

Unicyclic component

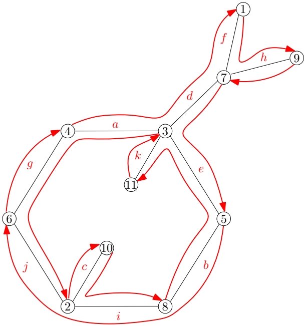

There are always two cycles. Each pendant subtree belongs entirely to one or other cycle.

Connected graphs and permutations

Unicyclic component

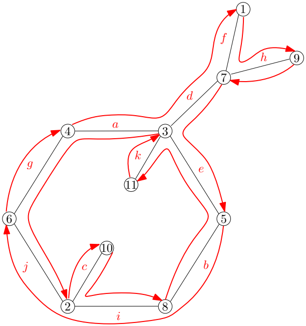

There are always two cycles. Each pendant subtree belongs entirely to one or other cycle.

Proposition. For G_m a uniform unicyclic component on m vertices, we have

$$\frac{1}{m}\left(C_1(\Pi(G_m)), C_2(\Pi(G_m))\right) \xrightarrow{d} (B, 1-B)^{\downarrow},$$

as $m \to \infty$, where $B \sim \mathrm{Beta}(1/2, 1/2)$.

Connected graphs and permutations

Unicyclic component

Sketch proof. G_m may be thought of as follows. Start from a uniform random labelled tree T_m on $[m]$ and pick two independent uniform points, u and v. Let L_m be the length of the path between the two points. Now reweight the distribution of T_m by $[2(L_m + 1)]^{-1}$ and, finally, put an edge between u and v.

Connected graphs and permutations

Unicyclic component

Sketch proof. G_m may be thought of as follows. Start from a uniform random labelled tree T_m on $[m]$ and pick two independent uniform points, u and v. Let L_m be the length of the path between the two points. Now reweight the distribution of T_m by $[2(L_m + 1)]^{-1}$ and, finally, put an edge between u and v.

$\Pi(G_m)$ has two permutation cycles, one "outside" and the other "inside" the graph-cycle.

Connected graphs and permutations

Unicyclic component

Sketch proof. G_m may be thought of as follows. Start from a uniform random labelled tree T_m on $[m]$ and pick two independent uniform points, u and v. Let L_m be the length of the path between the two points. Now reweight the distribution of T_m by $[2(L_m + 1)]^{-1}$ and, finally, put an edge between u and v.

$\Pi(G_m)$ has two permutation cycles, one "outside" and the other "inside" the graph-cycle. Each subtree hanging off the graph-cycle belongs entirely to one or the other of them. For a single subtree, it is equally likely to be inside or outside.

Connected graphs and permutations

Unicyclic component

Sketch proof. G_m may be thought of as follows. Start from a uniform random labelled tree T_m on $[m]$ and pick two independent uniform points, u and v. Let L_m be the length of the path between the two points. Now reweight the distribution of T_m by $[2(L_m+1)]^{-1}$ and, finally, put an edge between u and v.

$\Pi(G_m)$ has two permutation cycles, one "outside" and the other "inside" the graph-cycle. Each subtree hanging off the graph-cycle belongs entirely to one or the other of them. For a single subtree, it is equally likely to be inside or outside.

A pair of subtrees whose roots are at distance 2 or more along the graph-cycle belong to independent permutation cycles. (It turns out that there is negligible probability of having two adjacent "large" subtrees, or two rooted at the same point, so that we may treat the different subtrees as choosing their permutation cycles independently.)

Connected graphs and permutations

Unicyclic component

We now make use of the scaling limit of G_m. Before reweighting, we have

$$\frac{1}{\sqrt{m}} T_m \xrightarrow{d} \mathcal{T},$$

in the Gromov–Hausdorff–Prokhorov sense, where $\mathcal{T}$ is the Brownian CRT.

Connected graphs and permutations

Unicyclic component

We now make use of the scaling limit of G_m. Before reweighting, we have

$$\frac{1}{\sqrt{m}} T_m \xrightarrow{d} \mathcal{T},$$

in the Gromov–Hausdorff–Prokhorov sense, where $\mathcal{T}$ is the Brownian CRT. In particular, if we look at the (signed) height process of the forest of subtrees hanging off the path between u and v, putting trees above or below the path independently with probability $1/2$, in the limit we see exactly a Brownian bridge.

Connected graphs and permutations

Unicyclic component

The distance between u and v, rescaled by $1/\sqrt{m}$, is encoded in the limit picture as the total local time at 0, for which we write L_1^{br}. Reweighting the law by $(L_1^{\mathrm{br}})^{-1}$, by a theorem of Biane, Le Gall and Yor (1987) we obtain a process with the same distribution as the Brownian pseudo-bridge, $(\tau_1^{-1/2} B(t\tau_1), 0 \leq t \leq 1)$, where B is a standard Brownian motion and $\tau_1 = \inf\{t : L_t > 1\}$.

Connected graphs and permutations

Unicyclic component

The distance between u and v, rescaled by $1/\sqrt{m}$, is encoded in the limit picture as the total local time at 0, for which we write L_1^{br}. Reweighting the law by $(L_1^{\text{br}})^{-1}$, by a theorem of Biane, Le Gall and Yor (1987) we obtain a process with the same distribution as the Brownian pseudo-bridge, $(\tau_1^{-1/2} B(t\tau_1), 0 \leq t \leq 1)$, where B is a standard Brownian motion and $\tau_1 = \inf\{t : L_t > 1\}$.

In particular, there are only $\Theta(\sqrt{m})$ vertices along the cycle. So asymptotically all of the mass is in the pendant subtrees.

Connected graphs and permutations

Unicyclic component

The distance between u and v, rescaled by $1/\sqrt{m}$, is encoded in the limit picture as the total local time at 0, for which we write L_1^{br}. Reweighting the law by $(L_1^{\mathrm{br}})^{-1}$, by a theorem of Biane, Le Gall and Yor (1987) we obtain a process with the same distribution as the Brownian pseudo-bridge, $(\tau_1^{-1/2} B(t\tau_1), 0 \leq t \leq 1)$, where B is a standard Brownian motion and $\tau_1 = \inf\{t : L_t > 1\}$.

In particular, there are only $\Theta(\sqrt{m})$ vertices along the cycle. So asymptotically all of the mass is in the pendant subtrees.

The limiting proportion of the vertices which lie in the "outside" permutation cycle is then given by the proportion of time spent positive by $(\tau_1^{-1/2} B(t\tau_1), 0 \leq t \leq 1)$, which has Beta$(1/2, 1/2)$ distribution by Lévy's arcsine theorem. $\square$

Connected graphs and permutations

Unicyclic component

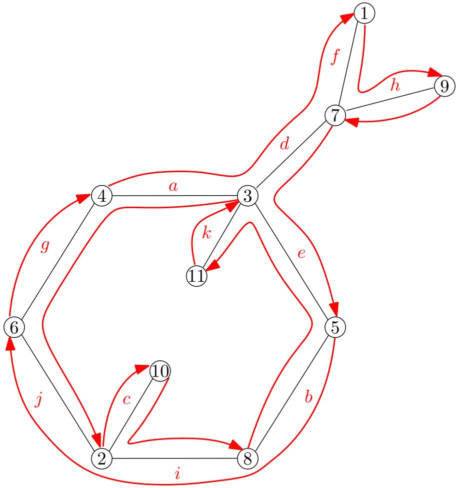

Proportion of mass in the "outside" cycle $\xrightarrow{d}$ Beta$(1/2, 1/2)$.

Connected graphs and permutations

Complex components (surplus ≥ 2)

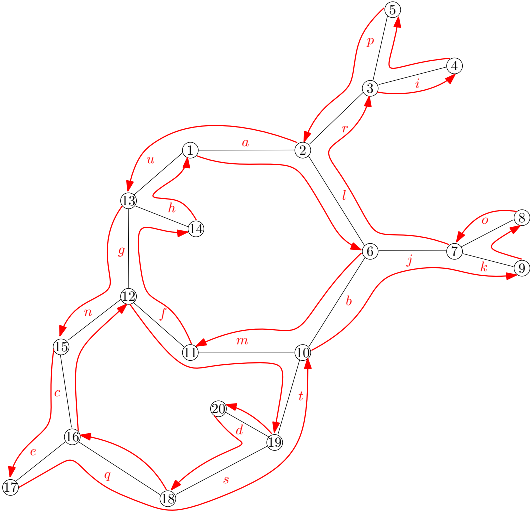

Connected graphs and permutations

Complex components

For the complex case, we make use of a standard decomposition of G into core and pendant subtrees. The core consists of the vertices in cycles, and those in paths joining cycles.

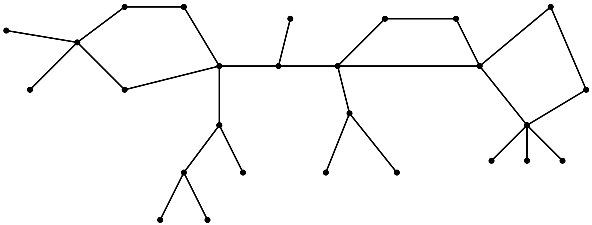

For the complex case, we make use of a standard decomposition of G into core and pendant subtrees. The core consists of the vertices in cycles, and those in paths joining cycles.

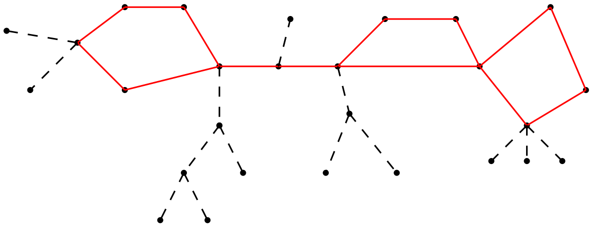

For the complex case, we make use of a standard decomposition of G into core and pendant subtrees. The core consists of the vertices in cycles, and those in paths joining cycles. The kernel is the multigraph obtained by contracting paths of vertices of degree 2 in the core.

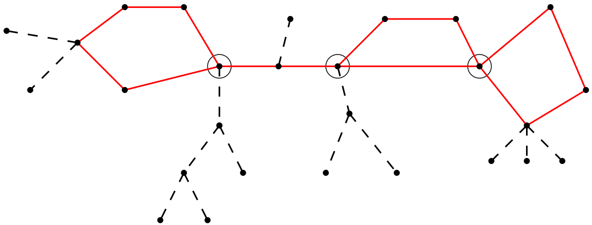

For the complex case, we make use of a standard decomposition of G into core and pendant subtrees. The core consists of the vertices in cycles, and those in paths joining cycles. The kernel is the multigraph obtained by contracting paths of vertices of degree 2 in the core.

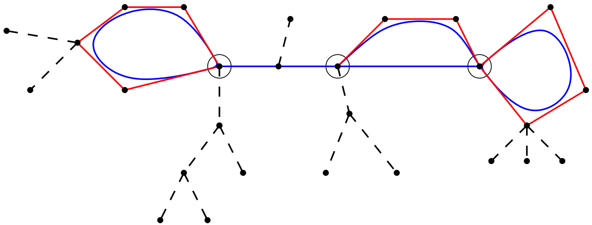

For the complex case, we make use of a standard decomposition of G into core and pendant subtrees. The core consists of the vertices in cycles, and those in paths joining cycles. The kernel is the multigraph obtained by contracting paths of vertices of degree 2 in the core.

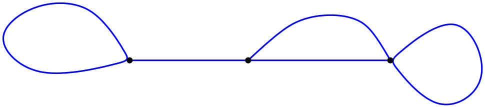

Connected graphs and permutations

Complex components

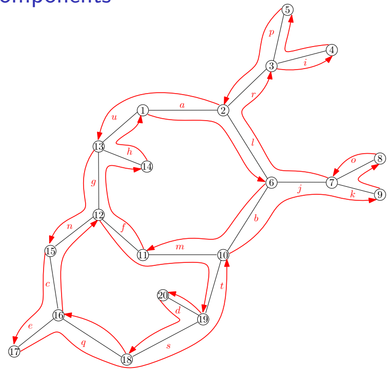

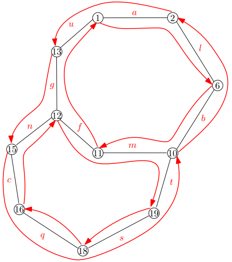

Now consider $\Pi(G)$. Surplus edges either split or merge cycles, depending on where they occur in the ordering.

Connected graphs and permutations

Complex components

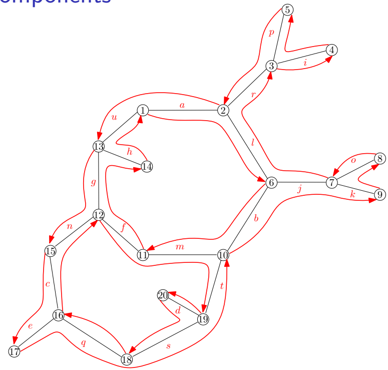

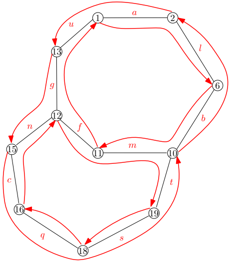

Now consider $\Pi(G)$. Surplus edges either split or merge cycles, depending on where they occur in the ordering. So, in general, the number of permutation-cycles is random, at most equal to surplus $+ 1$.

Connected graphs and permutations

Complex components

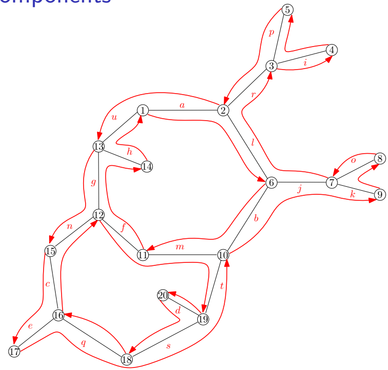

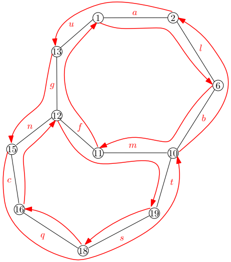

Now consider $\Pi(G)$. Surplus edges either split or merge cycles, depending on where they occur in the ordering. So, in general, the number of permutation-cycles is random, at most equal to surplus $+ 1$. Each pendant subtree again belongs entirely to one of the permutation-cycles. So let us first consider what happens on the core.

Connected graphs and permutations

Complex components

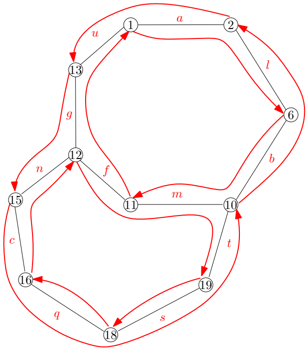

Take G to be a connected graph of minimum degree 2. Let us think of each edge $e = \{u, v\}$ of G as made up of two directed edges, (u, v) and (v, u).

Connected graphs and permutations

Complex components

When we transpose u and v, we think of the label currently at u as traversing the directed edge (u, v) and the label currently at v as traversing the directed edge (v, u). The permutation cycles thus partition the directed edges of G. Let us call the parts of this partition the trajectories, and think of them as directed cycles.

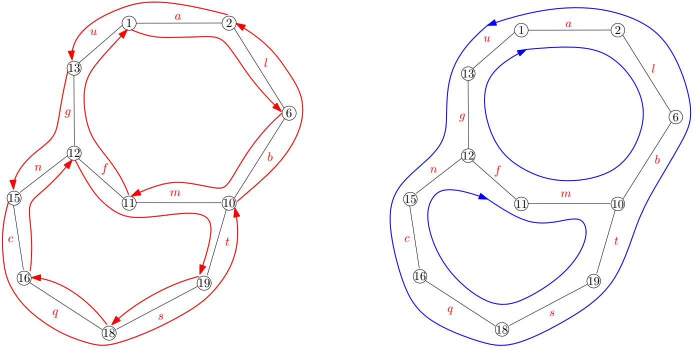

Connected graphs and permutations

Complex components

Nothing happens to the trajectories along paths of degree-2 vertices, so everything is determined by the kernel and, in particular, by what happens at its vertices (which have degree at least 3).

Connected graphs and permutations

Complex components

Nothing happens to the trajectories along paths of degree-2 vertices, so everything is determined by the kernel and, in particular, by what happens at its vertices (which have degree at least 3). Let us restrict attention to the case where the kernel vertices have degree precisely 3, which is the only one that will show up in our results.

Connected graphs and permutations

Complex components

Nothing happens to the trajectories along paths of degree-2 vertices, so everything is determined by the kernel and, in particular, by what happens at its vertices (which have degree at least 3). Let us restrict attention to the case where the kernel vertices have degree precisely 3, which is the only one that will show up in our results.

In this case, the ordering of the neighbouring edges gives the vertex an orientation: for each incoming edge, the orientation specifies which outgoing edge follows it. This cannot be the incoming edge with the opposite direction, so there are precisely two possibilities:

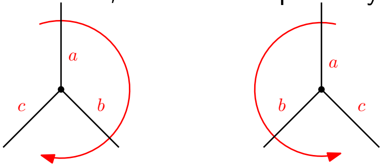

Connected graphs and permutations

Complex components

So, finally, the trajectories are generated by taking the kernel $K = \ker(G)$, assigning an independent uniform orientation to each vertex and connecting them up to one another:

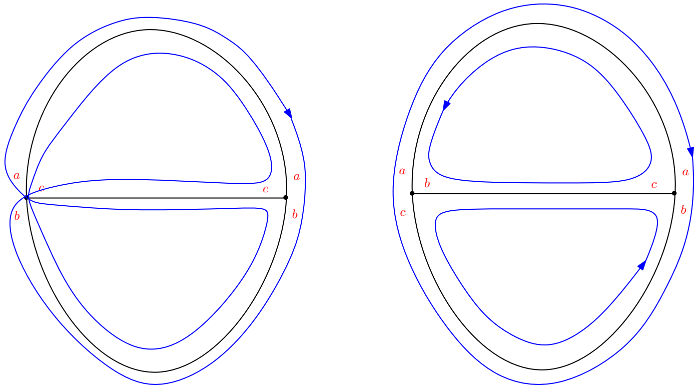

Critical complex Erdős–Rényi components

Let us now consider the complex components which arise in the critical Erdős–Rényi random graph. For $k \geq 2$, let $G_{m,k}$ be a uniform connected graph with vertex-set $[m]$ and $m + k - 1$ edges (i.e. surplus k).

Proposition. (Addario-Berry, Broutin & G., 2010, 2012)

$$\frac{1}{\sqrt{m}} G_{m,k} \xrightarrow{d} \mathcal{G}_k$$

as $m \to \infty$, in the Gromov–Hausdorff–Prokhorov sense, where $\mathcal{G}_k$ is constructed as on the next slide.

The scaling limit of a surplus k critical complex Erdős–Rényi component, $\mathcal{G}_k$

First sample K_k, a connected 3-regular multigraph generated according to the configuration model:

The scaling limit of a surplus k critical complex Erdős–Rényi component, $\mathcal{G}_k$

First sample K_k, a connected 3-regular multigraph generated according to the configuration model: take vertices labelled $1, 2, \ldots, 2(k-1)$.

The scaling limit of a surplus k critical complex Erdős–Rényi component, $\mathcal{G}_k$

First sample K_k, a connected 3-regular multigraph generated according to the configuration model: take vertices labelled $1, 2, \ldots, 2(k-1)$. Assign each vertex 3 (distinguishable) half-edges.

The scaling limit of a surplus k critical complex Erdős–Rényi component, $\mathcal{G}_k$

First sample K_k, a connected 3-regular multigraph generated according to the configuration model: take vertices labelled $1, 2, \ldots, 2(k-1)$. Assign each vertex 3 (distinguishable) half-edges. Sample a uniform matching of the half-edges, conditioned to yield a connected multigraph, which has $3(k-1)$ edges.

The scaling limit of a surplus k critical complex Erdős–Rényi component, $\mathcal{G}_k$

First sample K_k, a connected 3-regular multigraph generated according to the configuration model: take vertices labelled $1, 2, \ldots, 2(k-1)$. Assign each vertex 3 (distinguishable) half-edges. Sample a uniform matching of the half-edges, conditioned to yield a connected multigraph, which has $3(k-1)$ edges.

Sample $(\Theta_e, e \in E(K_k)) \sim \mathrm{Dir}(1/2, 1/2, \ldots, 1/2)$ and, independently, $\mathcal{T}_e, e \in E(K_k)$, independent Brownian CRT's. Rescale distances in $\mathcal{T}_e$ by $\sqrt{\Theta_e}$ and the mass by Θ_e.

The scaling limit of a surplus k critical complex Erdős–Rényi component, $\mathcal{G}_k$

First sample K_k, a connected 3-regular multigraph generated according to the configuration model: take vertices labelled $1, 2, \ldots, 2(k-1)$. Assign each vertex 3 (distinguishable) half-edges. Sample a uniform matching of the half-edges, conditioned to yield a connected multigraph, which has $3(k-1)$ edges.

Sample $(\Theta_e, e \in E(K_k)) \sim \mathrm{Dir}(1/2, 1/2, \ldots, 1/2)$ and, independently, $\mathcal{T}_e, e \in E(K_k)$, independent Brownian CRT's. Rescale distances in $\mathcal{T}_e$ by $\sqrt{\Theta_e}$ and the mass by Θ_e.

Pick two independent uniform points in each tree, and replace the kernel edge e by the scaled copy of $\mathcal{T}_e$, glued at the uniform points.

The scaling limit of a critical complex Erdős–Rényi component

Sample K_2.

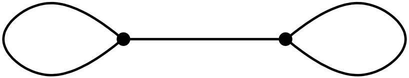

The scaling limit of a critical complex Erdős–Rényi component

Sample independent Brownian CRT's $\mathcal{T}_1, \mathcal{T}_2, \mathcal{T}_3$, each with two uniform marked points.

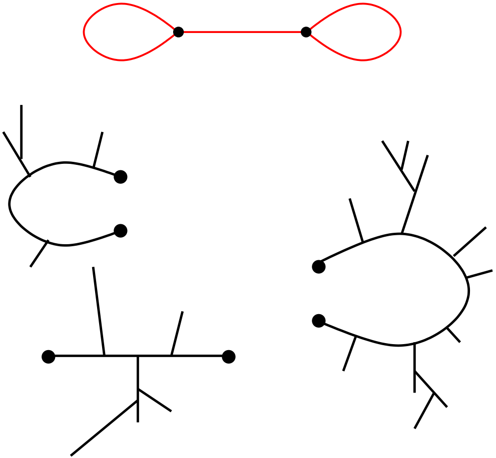

The scaling limit of a critical complex Erdős–Rényi component

Randomly rescale to $\sqrt{\Theta_e}\mathcal{T}_e$, so that the mass of becomes Θ_e, where $(\Theta_1, \Theta_2, \Theta_3) \sim \mathrm{Dir}(1/2, 1/2, 1/2)$.

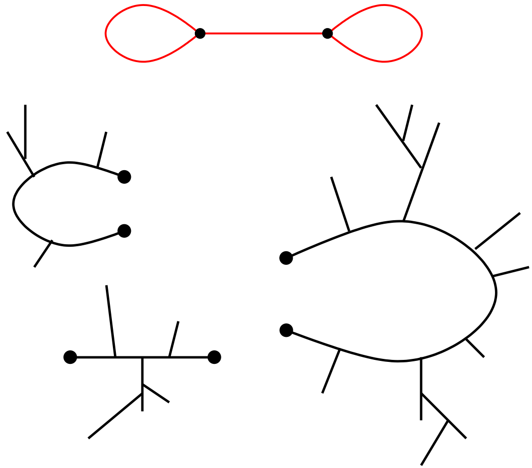

The scaling limit of a critical complex Erdős–Rényi component

Glue the trees to the kernel.

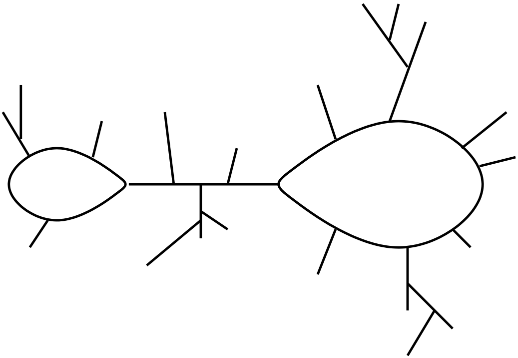

The scaling limit of a critical complex Erdős–Rényi component

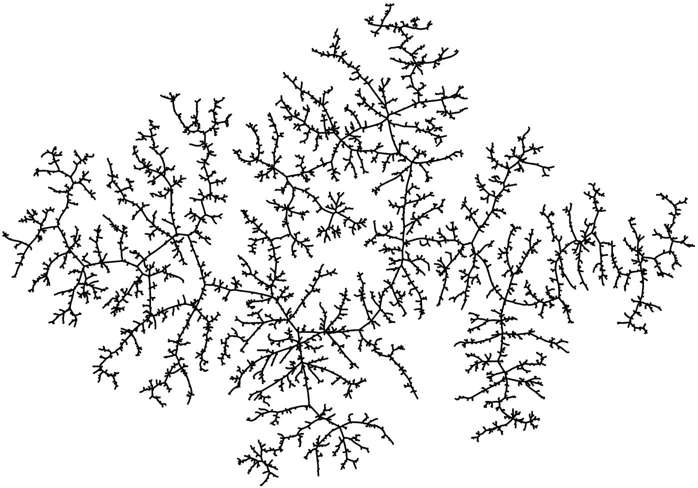

The trajectories

Recall that K_k is a connected 3-regular multigraph with $2(k-1)$ vertices and $3(k-1)$ edges, sampled according to the conditioned configuration model.

The trajectories

Recall that K_k is a connected 3-regular multigraph with $2(k-1)$ vertices and $3(k-1)$ edges, sampled according to the conditioned configuration model.

Let $K_k^{\circlearrowright}$ be K_k where each vertex is endowed with an independent uniform random orientation. This yields a random number $N(K_k^{\circlearrowright})$ of trajectories, $\tau_1(K_k^{\circlearrowright}), \ldots, \tau_{N(K_k^{\circlearrowright})}(K_k^{\circlearrowright})$.

The permutation cycles

Now sample $3(k-1)$ independent Brownian CRT's, $(\mathcal{T}_e, e \in E(K_k))$, each with two uniform points.

For each $\mathcal{T}_e$, along the path between the uniform points, we have (countably infinitely many) pendant subtrees. Independently for each subtree, put it above the path with probability $1/2$ or below the path otherwise.

The permutation cycles

Now sample $3(k-1)$ independent Brownian CRT's, $(\mathcal{T}_e, e \in E(K_k))$, each with two uniform points.

For each $\mathcal{T}_e$, along the path between the uniform points, we have (countably infinitely many) pendant subtrees. Independently for each subtree, put it above the path with probability $1/2$ or below the path otherwise.

The (signed) height process of this forest is precisely a Brownian bridge, for which another theorem of Lévy tells us the time spent positive is distributed as $\mathrm{U}[0,1]$.

The permutation cycles

Now sample $3(k-1)$ independent Brownian CRT's, $(\mathcal{T}_e, e \in E(K_k))$, each with two uniform points.

For each $\mathcal{T}_e$, along the path between the uniform points, we have (countably infinitely many) pendant subtrees. Independently for each subtree, put it above the path with probability $1/2$ or below the path otherwise.

The (signed) height process of this forest is precisely a Brownian bridge, for which another theorem of Lévy tells us the time spent positive is distributed as $\mathrm{U}[0,1]$.

Finally, sample $(\Theta_e, e \in E(K_k)) \sim \mathrm{Dir}(1/2, 1/2, \ldots, 1/2)$, Brownian rescale the trees and glue them to the kernel by the uniform points. The trajectories "pick up" the subtrees which lie to the appropriate side of the kernel edge.

Scaling limit for the permutation cycles

Proposition. For $k \geq 2$, conditionally on K_k, let

$$(\Theta_e, e \in E(K_k)) \sim \mathrm{Dir}(1/2, 1/2, \ldots, 1/2)$$

and, independently, let $(U_e, e \in E(K_k))$ be i.i.d. $\mathrm{U}[0,1]$ random variables. Then

$$\frac{1}{m}\Big(C_1(\Pi(G_{m,k})), C_2(\Pi(G_{m,k})), \ldots, C_{N(\Pi(G_{m,k}))}(\Pi(G_{m,k}))\Big)$$

$$\xrightarrow{d} \left(\sum_{\substack{e=\{u,v\}\\ \in E(K_k)}} \Theta_e \left(U_e \mathbb{1}_{\{(u,v)\in\tau_i(K_k^{\circlearrowright})\}} + (1-U_e)\mathbb{1}_{\{(v,u)\in\tau_i(K_k^{\circlearrowright})\}}\right), \quad 1 \leq i \leq N(K_k^{\circlearrowright})\right)^{\downarrow}$$

as $m \to \infty$.

Scaling limit for the permutation cycles

Proposition. For $k \geq 2$, conditionally on K_k, let

$$(\Theta_e, e \in E(K_k)) \sim \mathrm{Dir}(1/2, 1/2, \ldots, 1/2)$$

and, independently, let $(U_e, e \in E(K_k))$ be i.i.d. U$[0,1]$ random variables. Then

$$\frac{1}{m}\Big(C_1(\Pi(G_{m,k})), C_2(\Pi(G_{m,k})), \ldots, C_{N(\Pi(G_{m,k}))}(\Pi(G_{m,k}))\Big)$$

$$\xrightarrow{d} \left(\sum_{\substack{e=\{u,v\} \\ \in E(K_k)}} \Theta_e \left(U_e \mathbb{1}_{\{(u,v)\in\tau_i(K_k^{\circlearrowright})\}} + (1-U_e)\mathbb{1}_{\{(v,u)\in\tau_i(K_k^{\circlearrowright})\}} \right), \quad 1 \leq i \leq N(K_k^{\circlearrowright}) \right)^{\downarrow}$$

as $m \to \infty$.

Write $\mathbf{\Gamma}_k$ for a random vector with the same distribution as the right-hand side.

Putting the different components together

Let $Z_1^n(\lambda), Z_2^n(\lambda), \ldots$ be the sizes of the components of $\mathbb{G}(n, (1 + \lambda n^{-1/3})/n)$ and let $S_1^n(\lambda), S_2^n(\lambda), \ldots$ be the corresponding numbers of surplus edges.

Putting the different components together

Let $Z_1^n(\lambda), Z_2^n(\lambda), \ldots$ be the sizes of the components of $\mathbb{G}(n, (1+\lambda n^{-1/3})/n)$ and let $S_1^n(\lambda), S_2^n(\lambda), \ldots$ be the corresponding numbers of surplus edges.

Let $B_t^\lambda = B_t + \lambda t - t^2/2$, where B is a standard Brownian motion, and let $\underline{B}_t^\lambda = \inf_{0 \le s \le t} B_s^\lambda$. Conditionally on $B^\lambda - \underline{B}^\lambda$, let $(N_t^\lambda, t \ge 0)$ be an inhomogeneous Poisson process of intensity $(B_t^\lambda - \underline{B}_t^\lambda)dt$ at time t. Let $\zeta_1(\lambda), \zeta_2(\lambda), \ldots$ be the ordered excursion lengths of $B^\lambda - \underline{B}^\lambda$ above 0, and let $\sigma_1(\lambda), \sigma_2(\lambda), \ldots$ be the numbers of Poisson points falling in each of those excursions.

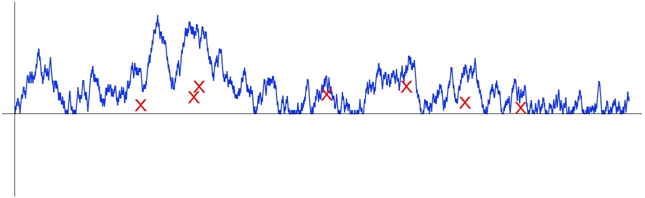

Putting the different components together

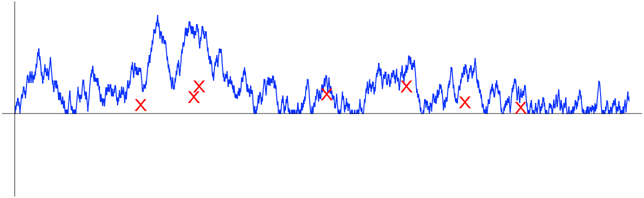

Theorem. (Aldous, 1997) We have the joint convergence

$$n^{-2/3}(Z_1^n(\lambda), Z_2^n(\lambda), \ldots) \xrightarrow{d} (\zeta_1(\lambda), \zeta_2(\lambda), \ldots)$$
$$(S_1^n(\lambda), S_2^n(\lambda), \ldots) \xrightarrow{d} (\sigma_1(\lambda), \sigma_2(\lambda), \ldots)$$

as $n \to \infty$.

The critical random transposition random walk

Recall that for $k \geq 2$, the random vector $\mathbf{\Gamma}_k$ is the scaling limit of the cycle-lengths of the permutation associated with a uniform connected graph on m vertices with $m + k - 1$ edges. Define $\mathbf{\Gamma}_0 = 1$ and $\mathbf{\Gamma}_1 = (B, 1 - B)^{\downarrow}$ where $B \sim \text{Beta}(1/2, 1/2)$.

The critical random transposition random walk

Recall that for $k \geq 2$, the random vector $\mathbf{\Gamma}_k$ is the scaling limit of the cycle-lengths of the permutation associated with a uniform connected graph on m vertices with $m + k - 1$ edges. Define $\mathbf{\Gamma}_0 = 1$ and $\mathbf{\Gamma}_1 = (B, 1 - B)^{\downarrow}$ where $B \sim \text{Beta}(1/2, 1/2)$.

For $k \geq 0$, let $\mathbf{\Gamma}_k^{(i)}, i \geq 1$ be i.i.d. copies of $\mathbf{\Gamma}_k$.

For fixed $\lambda \in \mathbb{R}$, let ν_λ be the distribution of the decreasing rearrangement of all of the terms of

$$\zeta_i(\lambda) \cdot \mathbf{\Gamma}^{(i)}_{\sigma_i(\lambda)}, \quad i \geq 1.$$

The critical random transposition random walk

Recall that for $k \geq 2$, the random vector $\mathbf{\Gamma}_k$ is the scaling limit of the cycle-lengths of the permutation associated with a uniform connected graph on m vertices with $m + k - 1$ edges. Define $\mathbf{\Gamma}_0 = 1$ and $\mathbf{\Gamma}_1 = (B, 1 - B)^{\downarrow}$ where $B \sim \text{Beta}(1/2, 1/2)$.

For $k \geq 0$, let $\mathbf{\Gamma}_k^{(i)}, i \geq 1$ be i.i.d. copies of $\mathbf{\Gamma}_k$.

For fixed $\lambda \in \mathbb{R}$, let ν_λ be the distribution of the decreasing rearrangement of all of the terms of

$$\zeta_i(\lambda) \cdot \mathbf{\Gamma}^{(i)}_{\sigma_i(\lambda)}, \quad i \geq 1.$$

Then as $n \to \infty$,

$$n^{-2/3} \left(C_1(P^n_{1+\lambda n^{-1/3}}), C_2(P^n_{1+\lambda n^{-1/3}}), \ldots \right) \xrightarrow{d} \nu_\lambda.$$

The critical random transposition random walk

Recall that $Y^*(\lambda)$ has distribution ν_λ.

Putting this together with the Feller property for the split-merge process essentially yields the process convergence in the critical window

$$\left(n^{-2/3}\left(C_1(P^n_{1+\lambda n^{-1/3}}), C_2(P^n_{1+\lambda n^{-1/3}}), \ldots\right), \lambda \in \mathbb{R}\right)$$
$$\xrightarrow{d} (Y^*(\lambda), \lambda \in \mathbb{R}).$$

A random map perspective

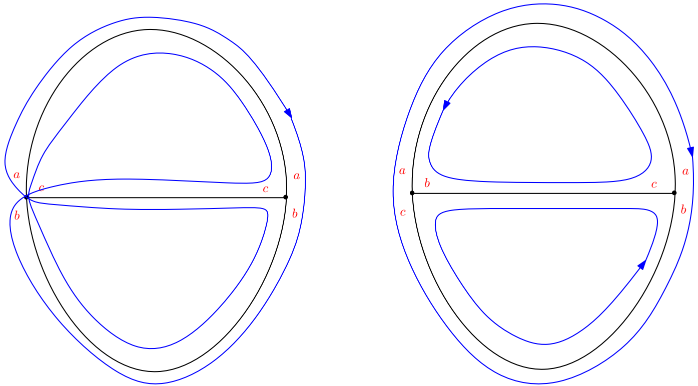

A random map perspective

One can think of the trajectories in a component as the face boundaries in one of the canonical combinatorial descriptions of a map. The number of faces of the map is precisely the (random) number of trajectories, and we can find the genus via Euler's formula, $2 - 2g = v - e + f$.

For example, $v = 19$, $e = 20$ (surplus 2), single permutation cycle ($f = 1$), so $g = 1$.

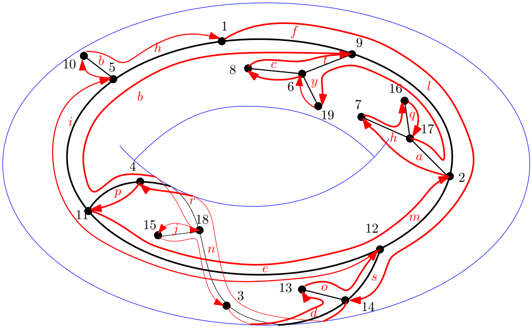

A random map perspective

Our results then be interpreted as a limit for the proportions of vertices lying in each of the faces of a corresponding random map model.

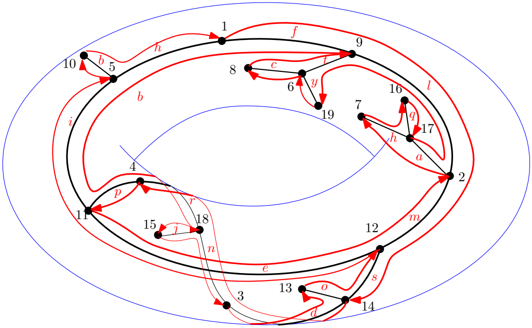

PD(0, 1) limit

We claimed that $\frac{1}{2\lambda} Y^*(\lambda) \xrightarrow{d} \mathrm{PD}(0,1)$ as $\lambda \to \infty$.

PD(0, 1) limit

We claimed that $\frac{1}{2\lambda} Y^*(\lambda) \xrightarrow{d} \mathrm{PD}(0,1)$ as $\lambda \to \infty$.

A caricature of $B_t^\lambda = B_t + \lambda t - t^2/2$ for large λ:

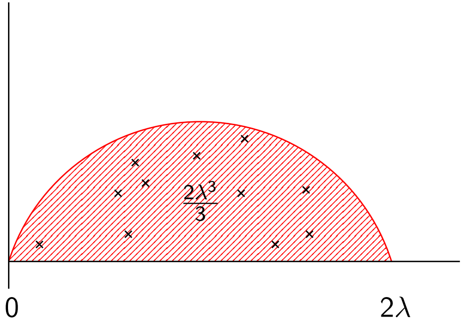

PD(0, 1) limit

We claimed that $\frac{1}{2\lambda} Y^*(\lambda) \xrightarrow{d} \mathrm{PD}(0,1)$ as $\lambda \to \infty$.

A caricature of $B_t^\lambda = B_t + \lambda t - t^2/2$ for large λ:

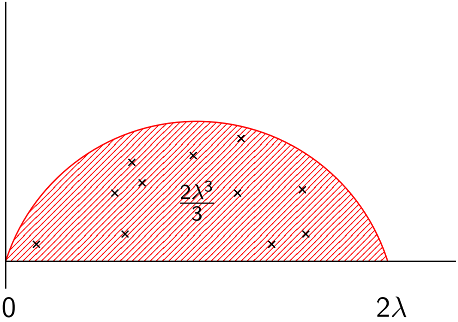

As $\lambda \to \infty$,

$$\zeta_1(\lambda) = 2\lambda + o_{\mathbb{P}}(1), \quad \zeta_2(\lambda) = o_{\mathbb{P}}(1), \quad \sigma_1(\lambda) = \Theta_{\mathbb{P}}(\lambda^3).$$

So on the scale of λ, we may ignore cycles coming from any component except the largest, which has (random) surplus on the order of λ^3.

$\mathrm{PD}(0,1)$ limit

So it should be sufficient to deal with the permutation-cycles arising from a single component with diverging surplus k.

Lemma. $\boldsymbol{\Gamma}_k \xrightarrow{d} \mathrm{PD}(0,1)$ as $k \to \infty$.

PD(0, 1) limit

So it should be sufficient to deal with the permutation-cycles arising from a single component with diverging surplus k.

Lemma. $\mathbf{\Gamma}_k \xrightarrow{d} \mathrm{PD}(0,1)$ as $k \to \infty$.

As $k \to \infty$, if $(\Theta_1, \ldots, \Theta_{3(k-1)}) \sim \mathrm{Dir}(1/2, \ldots, 1/2)$ and $U_1 \sim \mathrm{U}[0,1]$ then

$$\mathbb{E}\left[\Theta_1 U_1\right] = \frac{1}{6(k-1)}$$

and

$$\max_{1 \le i \le 3(k-1)} \Theta_i = o_{\mathbb{P}}(k^{-1/2}), \quad \mathbb{E}\left[\max_{1 \le i \le 3(k-1)} \Theta_i^2\right] = o(k^{-1}).$$

$\mathrm{PD}(0,1)$ limit

So it should be sufficient to deal with the permutation-cycles arising from a single component with diverging surplus k.

Lemma. $\mathbf{\Gamma}_k \xrightarrow{d} \mathrm{PD}(0,1)$ as $k \to \infty$.

As $k \to \infty$, if $(\Theta_1, \ldots, \Theta_{3(k-1)}) \sim \mathrm{Dir}(1/2, \ldots, 1/2)$ and $U_1 \sim \mathrm{U}[0,1]$ then

$$\mathbb{E}\left[\Theta_1 U_1\right] = \frac{1}{6(k-1)}$$

and

$$\max_{1 \le i \le 3(k-1)} \Theta_i = o_{\mathbb{P}}(k^{-1/2}), \quad \mathbb{E}\left[\max_{1 \le i \le 3(k-1)} \Theta_i^2\right] = o(k^{-1}).$$

So we have a large connected 3-regular multigraph K_k decorated with well-behaved small masses. This suggests that we should be able to treat the masses as deterministic, and that the distribution will be driven simply by the lengths of the trajectories.

PD$(0,1)$ limit

Recall that $K_k^{\circlearrowright}$ is a 3-regular configuration multigraph with $2(k-1)$ vertices, $3(k-1)$ edges and independent uniform orientations at its vertices, and that $\tau_1(K_k^{\circlearrowright}), \ldots, \tau_{N(K_k^{\circlearrowright})}(K_k^{\circlearrowright})$ are the associated trajectories. For a trajectory τ, write $|\tau|$ for the number of edges it contains. Note that

$$\sum_{1 \le i \le N(K_k^{\circlearrowright})} |\tau_i(K_k^{\circlearrowright})| = 6(k-1).$$

PD(0, 1) limit

Recall that $K_k^{\circlearrowright}$ is a 3-regular configuration multigraph with $2(k-1)$ vertices, $3(k-1)$ edges and independent uniform orientations at its vertices, and that $\tau_1(K_k^{\circlearrowright}), \ldots, \tau_{N(K_k^{\circlearrowright})}(K_k^{\circlearrowright})$ are the associated trajectories. For a trajectory τ, write $|\tau|$ for the number of edges it contains. Note that

$$\sum_{1 \leq i \leq N(K_k^{\circlearrowright})} |\tau_i(K_k^{\circlearrowright})| = 6(k-1).$$

Theorem. (Gamburd, 2006) As $k \to \infty$,

$$\frac{1}{6(k-1)}(|\tau_i(K_k^{\circlearrowright})|, 1 \leq i \leq N(K_k^{\circlearrowright}))^{\downarrow} \xrightarrow{d} \mathrm{PD}(0,1).$$

PD(0, 1) limit

Recall that $K_k^{\circlearrowright}$ is a 3-regular configuration multigraph with $2(k-1)$ vertices, $3(k-1)$ edges and independent uniform orientations at its vertices, and that $\tau_1(K_k^{\circlearrowright}), \ldots, \tau_{N(K_k^{\circlearrowright})}(K_k^{\circlearrowright})$ are the associated trajectories. For a trajectory τ, write $|\tau|$ for the number of edges it contains. Note that

$$\sum_{1 \leq i \leq N(K_k^{\circlearrowright})} |\tau_i(K_k^{\circlearrowright})| = 6(k-1).$$

Theorem. (Gamburd, 2006) As $k \to \infty$,

$$\frac{1}{6(k-1)}(|\tau_i(K_k^{\circlearrowright})|, 1 \leq i \leq N(K_k^{\circlearrowright}))^{\downarrow} \xrightarrow{d} \mathrm{PD}(0,1).$$

This is known as the Brooks–Makover conjecture, and was proved using representation theory. We have an alternative purely probabilistic proof.

PD(0, 1) limit

Putting everything together, we obtain

$$\frac{1}{2\lambda} Y^*(\lambda) \xrightarrow{d} \mathrm{PD}(0,1)$$

as $\lambda \to \infty$, as claimed.

PD(0, 1) limit

Putting everything together, we obtain

$$\frac{1}{2\lambda} Y^*(\lambda) \xrightarrow{d} \mathrm{PD}(0, 1)$$

as $\lambda \to \infty$, as claimed.

Work in progress: we show that the PD(0, 1) limit also holds in the barely supercritical regime, where $t = 1 + \epsilon(n)$ with $n^{-1/3} \ll \epsilon(n) \ll 1$.

Joyeux anniversaire Jean-François !